



HAL
open science

Scenario Discovery in a Complex Economy : Exploring the Parameter space of Agent-based Models

Karl Naumann-Woleske

► **To cite this version:**

Karl Naumann-Woleske. Scenario Discovery in a Complex Economy : Exploring the Parameter space of Agent-based Models. Economics and Finance. Institut Polytechnique de Paris, 2024. English. NNT : 2024IPPAX005 . tel-04648545

HAL Id: tel-04648545

<https://theses.hal.science/tel-04648545v1>

Submitted on 15 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Scenario Discovery in a Complex Economy

Exploring the Parameter space of Agent-based Models

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°626 École doctorale de l'Institut Polytechnique de
Paris (EDIPP)

Spécialité de doctorat: Sciences Économiques

Thèse présentée et soutenue à Palaiseau, le 5 Mars 2024, par

KARL NAUMANN-WOLESKE

Composition du Jury :

Antoine Mandel Universite Paris 1 Pantheon Sorbonne	Président
James P. Sethna Cornell University	Rapporteur
Francois Geerolf OFCE & Sciences Po	Examineur
Irene Monasterolo Utrecht University	Examinatrice
Michael Benzaquen CNRS & Ecole Polytechnique	Directeur de thèse
Jean-Philippe Bouchaud Capital Fund Management	Co-directeur de thèse
Alan Kirman EHESS	Invité

To Gabrielle and Yara

Acknowledgments

I would like to extend my sincerest gratitude to Professors Michael Benzaquen, Jean-Philippe Bouchaud and Alan Kirman for the opportunity to complete my doctoral journey under their guidance. Throughout these past three years they always made time to guide me along my research path since I started as an intern at the Econophysix Lab in 2020. They gave me the autonomy to explore my interests in various directions, whether fruitful or not. It has truly been a blessing to be their student and learn from their vast experience. In particular, I want to thank Michael for his trust in taking me on as an intern and student, for his patience, kindness and pragmatic approach. For his trust in my independence and in allowing me to supervise, guide and teach several interns along the way. Thank you to Jean-Philippe for pointing me in the direction of the fascinating research project of parameter-space exploration, for his availability and generosity with his time, sharing his excitement for research in complex systems, his patience in explaining concepts from statistical physics, and our regular debates about macroeconomics. This thesis is in many ways the culmination of three years of work, but also only a starting point for many more research avenues and I am excited to continue our collaboration in the future.

Many thanks to the members of my jury, James Sethna, Antoine Mandel, Irene Monasterolo and Francois Geerolf for committing their time and attention to the fruits of my research projects. In particular, James and Antoine for agreeing to act as rapporteurs and reading through the entirety of this thesis in detail.

I would like to thank the members of the Chair of EconophysiX and Complex Systems for making it a welcome and wonderful environment, so thank you Jerome, Max, Cecilia, Samy, Salma, Swann, Fabian, Ruben, Elia, Michele, Rudy, Victor, Nirbhay, Antoine, Natasha, and Jutta. Thank you to Federico and Jose for welcoming me to the Econophysix Lab when I first started, and for your guidance in matters of navigating a PhD (and introducing me to the ways of hyperbullet chess). In particular, I would like to thank Max, it was a great pleasure to welcome you to the Econophysix Lab as an intern and see you go on to pursue your own research streams. Thank you for your patience

as we slowly sifted through the research that has become the first part of this thesis. Thank you also to Elia, with whom inspiration from a workshop in Trieste turned into a brand new research avenue. I am excited to see where both of your research journeys take you!

I would like to extend my thanks to William Hynes and the New Approaches to Economic Challenges initiative at the Organisation for Economic Cooperation and Development (OECD) for generously funding this doctorate, and for allowing me to explore the world of complexity economics during my 2019 internship. Likewise to Sebastian Barnes for inviting me to present the fruits of my research at the OECD.

Finally, I would like to thank Gabrielle and Yara, to whom this manuscript is dedicated. Gabrielle has been at my side for the entirety of this long adventure, and her encouragement and belief in me throughout this process has been unfailing both in times of rapid progress and times of existential questions. She has brought so much joy to my life, and now in 2023 the birth of Yara has brought on one of the most beautiful, difficult and rewarding new journeys just at the start of writing the final manuscript. Thank you both for making life wonderful.

Publications

There are several publications and preprints that have been included as a part of this thesis:

1. **Naumann-Woleske**, Knicker, Benzaquen & Bouchaud (2021) “Exploration of the Parameter Space in Macroeconomic Models” *Forthcoming in Handbook of Complexity Economics*, *arXiv:2111.08654*
2. **Naumann-Woleske** (2023) Agent-based Integrated Assessment Models: Alternative Foundations to the Environment-Energy-Economics Nexus *arXiv:2301.08135*
3. **Naumann-Woleske**, Benzaquen, Gusev, Kroujiline (2022) “Capital Demand Driven Business Cycles: Mechanism and Effects” *Review of Behavioral Economics* Vol. 9 No. 4
4. Knicker, **Naumann-Woleske**, Bouchaud, & Zamponi (2023) “Post-COVID Inflation & the Monetary Policy Dilemma: An Agent-Based Scenario Analysis” *Review & Resubmit at the Journal of Economic Interaction and Coordination*. *arXiv:2306.01284*
5. Morelli, **Naumann-Woleske**, Benzaquen, Tarzia, Bouchaud (2021) “Economic Crises in a Model with Capital Scarcity and Self-Reflexive Confidence” *arXiv:2109.09386*

Résumé en Français

L'économie se caractérise par un réseau de millions d'entreprises connectées par des liens physiques interagissant avec un réseau de milliards de ménages à travers une toile compliquée de trillions de contrats. L'écologie des agents hétérogènes interagissant localement conduit à des dynamiques agrégées émergentes telles que des successions de crises et de reprises ou de croissance économique, des régularités statistiques et des structures institutionnelles. Comme l'a dit Anderson (1972): "more is different" (plus c'est différent) lorsqu'il s'agit de comprendre comment les dynamiques macroéconomiques émergent des micro-agents et de leurs interactions. Un débat persistant a eu lieu sur les approches de modélisation visant à saisir ces dynamiques (voir par exemple Vines and Wills, 2018, 2020, pour une revue). L'utilisation de modèles multi-agents (ABM) est une approche récente en macroéconomie qui génère ces phénomènes en simulant une multiplicité d'agents hétérogènes en interaction. En simulant une économie "from the bottom up" (du bas vers le haut), les ABM commencent par une population d'agents hétérogènes (par exemple, plusieurs ménages, entreprises, institutions financières) dotés de leurs règles comportementales pertinentes, ainsi que d'un ensemble de protocoles pour les interactions entre agents. Ce système est ensuite simulé, en suivant l'état et les décisions de chaque agent, à partir desquels les dynamiques agrégées peuvent ensuite être directement calculées. L'évolution du système basée sur les règles comportementales des agents donne lieu aux dynamiques complexes qu'un modèle macroéconomique devrait récupérer. De plus, la modularité de ces modèles implique qu'une variété de différentes politiques publiques peuvent être explorées. La non linéarité hors équilibre de ces modèles signifie que des solutions analytiques sont généralement irréalisables, et les modélisateurs s'appuient sur des simulations numériques du système pour obtenir un aperçu des phénomènes générés (Fagiolo and Roventini, 2017). Cette charge computationnelle, associée au grand nombre de paramètres et de conditions initiales, rend le calibrage et l'exploration complète de l'espace des paramètres presque irréalisable.

Dans cette thèse, je propose une nouvelle approche pour aborder le problème de l'exploration de l'espace des paramètres dans les modèles macroéconomiques multi-agents en posant la question suivante :

Quel est l'ensemble de phénomènes qualitativement différents qu'un Modèle Multi-Agent Macroeconomique (MABM) peut générer, et qu'est-ce qui régit leurs transitions ?

Cette thèse est divisée en plusieurs parties distinctes couvrant mes recherches sur la question de la génération de scénarios au cours des trois dernières années. Dans la Partie I, je développe une approche algorithmique pour explorer différents scénarios dans les modèles multi-agents et montre comment elle peut récupérer l'espace des phases du modèle multi-agents Mark-0. La Partie II de la thèse se concentre sur deux modèles d'agents interagissant intégrés dans des cadres macroéconomiques néoclassiques pour montrer comment les phénomènes émergents de ces interactions peuvent enrichir la phénoménologie de ces modèles, ainsi qu'une application du MABM Mark-0 à la poussée d'inflation post-COVID.

Une critique des modèles multi-agents est qu'ils sont perçus comme des *boîtes noires* où les mécanismes conduisant aux phénomènes émergents sont peu clairs, l'ensemble des dynamiques qu'un modèle peut générer est illimité, et en raison du grand nombre de paramètres, l'ajustement des phénomènes empiriques est futile. La Partie I de cette thèse propose une méthode pour résoudre ces préoccupations et montre qu'il existe un ensemble fini de phases qualitativement distinctes dans un ABM, et que les dynamiques locales dépendent uniquement de quelques paramètres. Pour ce faire, le Chapitre 2 décrit une approche pour ajuster et comprendre les modèles à grand nombre de paramètres développés dans la littérature biophysique. Cette approche, appelée *sloppiness* par ses initiateurs, met en évidence que les dynamiques générées par les modèles de grande dimension sont souvent sensibles à seulement quelques combinaisons de paramètres bien contraintes. En revanche, de nombreuses autres combinaisons de paramètres sont mal contraintes, elles peuvent varier sur plusieurs ordres de grandeur sans changement significatif dans les observables. Identifier et classer ces combinaisons de paramètres peut éclairer l'environnement local de sensibilité des paramètres, nous permettant d'exploiter les directions les plus bien contraintes pour explorer l'espace des phases d'un modèle en modifiant au maximum les dynamiques observables. Passant de la théorie à la pratique, le Chapitre 3 quantifie ce concept pour les ABM, en prenant deux mesures distinctes pour quantifier un changement dans les observables : premièrement, une approche basée sur l'erreur quadratique moyenne qui compare directement les réalisations de séries temporelles et est destinée à fonctionner sur des échelles de temps plus courtes intéressant les décideurs politiques, et deuxièmement, la divergence de Kullback-Leibler symétrique comme une approche probabiliste pour comparer la distribution des observables qui peut être plus robuste pour les systèmes stochastiques non linéaires comme les ABM. Pour illustrer le fonctionnement de cette approche, le modèle de "fourmis de Kirman" est utilisé comme exemple théorique dans le Chapitre 4. Les approches basées sur l'EQM et sur la divergence KL sont évaluées analytiquement et numériquement sur l'ensemble des paramètres, et permettent de récupérer la combinaison de paramètres représentant une transition de phase dans le modèle. Dans le Chapitre 5, je me tourne vers l'analyse du modèle multi-agents Mark-0 qui a fait l'objet de notre premier article (Naumann-

Woleske et al., 2023). Le modèle Mark-0 est un MABM simple qui affiche néanmoins un ensemble de phénomènes qualitativement riches. En particulier, il existe quatre phases distinctes du taux de chômage, allant du plein emploi au chômage résiduel, aux oscillations endogènes et au chômage total. L'évaluation des directions de paramètres bien contraintes de Mark-0 à différents points de ces phases permet généralement de récupérer les paramètres responsables des transitions de phase, ce qui implique que les perturbations dans ces directions peuvent être un moyen efficace de découvrir les différents scénarios qu'un ABM peut générer. En exploitant ces paramètres dans un algorithme de montée de gradient simple, je montre dans le Chapitre 6 qu'on peut explorer de manière efficace et efficiente l'espace des phases du modèle Mark-0, non seulement en récupérant les différentes phases mais aussi en acquérant une compréhension des sensibilités localisées de ces dynamiques. En particulier, on peut avoir une idée de certaines des combinaisons de paramètres qui déclenchent différentes transitions de phase.

La Partie II de cette thèse adopte une approche différente pour explorer l'économie en tant que système complexe adaptatif. Plus précisément, je considère l'hétérogénéité et le comportement non rationnel dans deux cadres principaux, un modèle simple de Solow (Chapitre 7) et l'approche de l'Équilibre Général Dynamique Stochastique Néo-Keynésien (Chapitre 8), ainsi que les effets de la confiance dans le modèle multi-agents Mark-0 (Chapitre 9).

Les deux premiers modèles de cette partie représentent des investigations sur la manière dont les décisions d'investissement motivées par le sentiment peuvent conduire à des phénomènes agrégés cycliques, inspirées par l'idée que "the subjective evaluation of prospects over aq time horizon is the major proximate basis for investment and portfolio decisions, and these subjective estimates are changeable" (l'évaluation subjective des perspectives sur un horizon temporel est la principale base immédiate pour les décisions d'investissement et de portefeuille, et ces estimations subjectives sont changeables) (Minsky, 1976). Sur le plan méthodologique, cette exploration vise à réunir une approche multi-agents et une approche équilibrée classique pour mettre en évidence comment même de petites irrationalités et interactions des agents peuvent induire toute une série de phénomènes qualitativement différents qui seraient autrement négligés. Il convient de noter que ce sont des approches *simples*, elles ne sont pas destinées à être calibrées, mais à démontrer que même de simples interactions peuvent conduire à des phénomènes émergents pertinents pour l'utilisation des modèles.

Le Chapitre 7 développe un *Modèle Solow Dynamique*, introduisant le capital et l'investissement dans un modèle de croissance de Solow simple, reproduisant notre travail dans Naumann-Woleske et al. (2022). L'investissement dépend d'un grand nombre d'investisseurs individuels, qui forment une attente de rendements futurs en fonction de l'état de l'économie ainsi que des attentes des autres investisseurs. S'appuyant sur des outils de la physique statistique, ce système peut être écrit comme un ensemble d'équations différentielles non linéaires des résultats agrégés sur la base de micro-interactions, où l'effet du retour d'information du sentiment des autres investisseurs conduit à des cycles d'enthousiasme et de pessimisme de durée hétérogène. Bien qu'il

s’agisse d’un modèle simple, nous identifions un cycle économique de fluctuations quasi-périodiques autour d’un chemin de croissance d’équilibre stable dans une phase où l’économie est menée par la demande d’investissement. Ces cycles peuvent être identifiés comme un mécanisme de résonance de cohérence, où le processus aléatoire exogène conduit à des fluctuations quasi-périodiques, comme dans l’effet “small shocks, large cycles” (petits chocs, grand cycle) de Bernanke et al. (2019). À son tour, le Chapitre 8 reproduit notre travail dans Morelli et al. (2021), en développant davantage le modèle néo-keynésien étendu de Morelli et al. (2020), qui a développé un modèle néo-keynésien adapté où un grand nombre de ménages sont caractérisés par un sentiment sur l’état de l’économie. Tout comme dans le Modèle Solow Dynamique, le sentiment des ménages est influencé par leurs voisins, ce qui conduit à un effet de “keeping up with the Joneses” (garder le rythme avec les Jones), et, encore une fois en utilisant des outils de la physique statistique, peut être écrit comme une approche de représentant agent modifiée avec des attentes non linéaires. Dans ce modèle, nous introduisons également le capital comme facteur de production, permettant au ménage comportemental d’investir l’épargne dans un actif risqué fournissant du capital aux entreprises. Cette addition conduit à un modèle de cycle réels (RBC) comportemental avec une riche phénoménologie, où la pénurie de capital peut entraîner des périodes prolongées de faible production jusqu’à ce que la confiance soit restaurée, mettant en évidence l’importance de la communication par les gouvernements. Plus précisément, le rôle de la confiance dans l’investissement ici est similaire au paradoxe de l’épargne de Keynes : une confiance plus faible conduit à plus d’épargne, exacerbant les crises économiques, ce qui signifie qu’en plus de la communication pour restaurer la confiance, les décideurs politiques peuvent envisager des plans d’investissement directs. Le modèle est, encore une fois, un modèle simpliste destiné à mettre en évidence la richesse de la phénoménologie macroéconomique émergeant des interactions des agents individuels. Dans le Chapitre 9, nous examinons ensuite de plus près la confiance dans le contexte de l’augmentation de l’inflation suite à la pandémie de COVID-19 en 2020. Avec une inflation culminant à 9,1% aux États-Unis en juin 2022, il existe des récits concurrents pour expliquer ce phénomène : de “too much money chasing too few goods” (trop d’argent pour trop peu de biens) aux prix de l’énergie et de la nourriture, en passant par l’inflation motivée par les bénéfices et l’inflation conflictuelle. Pour comprendre certains des scénarios inflationnistes qui pourraient survenir et les lier à différentes explications et réponses politiques, nous étendons le modèle Mark-0 avec un secteur énergétique stylisé et introduisons trois chocs pour simuler des moteurs exogènes et des récits concurrents.¹ Nous constatons que la reprise économique après COVID-19 est lente en l’absence de politique publique d’atténuation, prenant au moins plusieurs années. Cependant, il n’existe qu’un nombre restreint d’options pour une politique monétaire efficace, une réponse avec une forte hausse des taux d’intérêt entraîne de manière disproportionnée le chômage. D’autre part, le déracinement de la confiance des agents dans la capacité de la Banque Centrale à contenir l’inflation influence les dynamiques de l’inflation elle-même en influençant le processus de négociation des

¹Y compris un choc COVID affectant la productivité et la demande, un choc de la chaîne d’approvisionnement sur la productivité, et un choc de prix de l’énergie.

salaires, le processus de fixation des prix et les décisions de consommation des ménages, entraînant une inflation plus élevée et de moins bonnes performances économiques.

Contents

Acknowledgments	iv
Publications	vii
Résumé en Français	viii
List of Figures	xix
List of Tables	xxi
Introduction	1
1 Introduction	3
1.1 The Economy as a Complex System	5
1.2 The State of Macroeconomic Modeling	6
1.3 Agent-based Models: Alternative Foundations	8
1.4 This Thesis	10
I Exploring the Parameter Space in Macroeconomic Models	15
2 How Complex Are Complex Models	17
2.1 Predictive Power Despite Uncertainty	18
2.2 Are Macroeconomic Agent-Based Models Sloppy?	20
3 Quantifying Sloppiness	23
3.1 Identifying Stiff and Sloppy Directions	23
3.2 Non-linear Least Squares	25
3.3 Probabilistic Loss	28

4	Kirman’s Ants	31
4.1	The Recruitment Model	32
4.2	Identifying Stiff and Sloppy Directions	32
4.3	Steps into the stiffest direction	35
5	Mark-0: Sloppiness in a MABM	39
5.1	The Mark-0 Model	40
5.2	Separating Signal and Noise: Estimating the MSE-Hessian	45
5.3	Disentangling the Parameter Hierarchy of Mark-0	47
5.4	Dynamic Stochastic General Equilibrium Models: A Note	53
6	Exploration of the Mark-0 Agent-Based Model	57
6.1	Phase Transitions in Two Dimensions	58
6.2	A Simple Algorithm for Parameter space Exploration	60
6.3	Understanding the Map from Parameters to Phases	62
6.4	Recovering Mark-0’s Phases and their Transitions	63
6.5	What Drives the Phase Transitions?	69
6.6	Distinguishing Explorations Quantitatively: the $pMSE(\mathcal{E}_Q)$	71
II	Confidence and Collective Behavior in Macroeconomic Models	73
7	Investment-Driven Business Cycles	75
7.1	The Dynamic Solow Model	78
7.2	Demand and Supply: Two Limiting Cases	86
7.3	Business Cycles and Long-Term Growth in the General Case	92
7.4	Conclusion	101
8	Economic Crises and Self-Reflexive Confidence	103
8.1	A Behavioral Business Cycle Model	106
8.2	Crises & Phase Diagrams	118
8.3	Discussion and Conclusions	126
9	Post-COVID Inflation & the Monetary Policy Dilemma	129
9.1	The Mark-0 Model	134
9.2	Policy Channels in Mark-0	140
9.3	Stationary Dynamics: the Role of the Central Bank	143
9.4	The COVID Shock and its Aftermath in the Absence of Monetary Policy	145
9.5	Monetary Policy Response to Inflationary Shocks	152
9.6	Fiscal Stimulus	156
9.7	Model Sensitivity: The Dangers of a Wage-Price Spiral	159
9.8	Summary & Conclusions	161

Conclusion	165
10 Summary of Results	167
11 Conclusions and Perspectives	171
11.1 Diving Deeper into the Estimation and Exploration of MABM	171
11.2 Applications: Agent-based Integrated Assessment Models	173
11.3 Inflationary Dynamics and their Drivers	185
Bibliography	187
Appendices	217
A Derivations for Quantifying Sloppiness	219
A.1 Derivation of the Hessian for the MSE Loss	219
A.2 Hessian for the sKL divergence	220
A.3 Hessian for the Hellinger divergence	221
B Derivations for Kirman’s Ants	223
B.1 The Kullback-Leibler Divergence	223
B.2 Hessian Matrices for Kirman’s Ants	224
C Numerical Estimation of MSE-based Hessians	227
D Additional Information for the Exploration of the Parameter-space	233
D.1 Best $pMSE(\mathcal{E}_Q)$ Explorations	233
D.2 Transition Drivers	238
E ABIAM Model Structures	243
E.1 Macroeconomic Overview	243
E.2 Energy and Resource Modules	252
E.3 Climate Modules	257
E.4 Policy Studies	261
F The Dynamic Solow Model	265
F.1 Parameterisation of the Dynamic Solow Model	265
F.2 Model Variables and Parameters	270
F.3 Approximate Solution to the Supply-Driven Regime	272
F.4 Asymptotic Analysis of Long-Term Growth	273
G The Self-Reflexive DSGE Model	277
G.1 Table of Parameters	277
H Post-COVID Inflation	279
H.1 Table of Parameters	280

H.2 Counterfactual Simulations	281
H.3 Monetary Policy	282
H.4 The Effect of Anchoring	286
H.5 Monetary Policy with Stronger Central Bank Reaction	287
H.6 Sensitivity of Monetary Policy to α_c and α_Γ	288
H.7 Helicopter Money	289
H.8 Windfall Tax	292
H.9 Sloppiness	294
H.10 Sensitivity to the Indexation Parameters g_p and g_w	296

List of Figures

2.1	Schematic interpretation of a two-dimensional sloppy model	19
4.1	Dynamics of Kirman’s ant model	33
4.2	Eigenvalues of the MSE and SKL Hessian in Kirman’s Ants	35
4.3	Loss from Perturbations in Eigenvector Directions in Kirman’s Ants	36
5.1	Unemployment Dynamics for 3 Phases in the Mark-0 model	45
5.2	Overview of Mark-0 MSE Hessians at Three Points	48
5.3	Eigenvalue Dynamics	49
5.4	Eigenvectors across Random Realizations for three points	50
5.5	Eigenvector Cosine Similarity across Random realizations for three points	51
5.6	Eigenvector Perturbations at Three Points	52
5.7	Eigenvalue Spectra for six DSGE models	53
5.8	Eigenvectors for six DSGE models	54
6.1	Mark-0 unemployment phase diagram in two dimensions	59
6.2	Sobol Baseline Exploration	64
6.3	Best runs based on number of phases recovered	65
6.4	New Phases in Mark-0	66
6.5	Explorations projected into the Mean-Std space	67
6.6	Number of Recovered Phase Transitions	68
6.7	Eigenvectors related to the FE phase	70
6.8	Overview of the $pMSE(\mathcal{E}_Q)$	71
7.1	A conceptual flowchart of the Dynamic Solow model	79
7.2	Output in the Supply-driven case	87
7.3	3D phase portrait in the demand-driven case	89
7.4	Projected phase portrait in the demand-driven case	90
7.5	Sample dynamics in the 3D s, h, z -space	91
7.6	Sample dynamics of sentiment and output in the demand-driven case	91

7.7	Long-term dynamics in the supply-driven regime	94
7.8	Long-term dynamics in a limit-cycle regime	95
7.9	Long-term dynamics with a stochastic limit-cycle	96
7.10	Long-term dynamics with coherence resonance	96
7.11	Long-term dynamics in the general case	97
7.12	Medium-term dynamics of the general case	98
7.13	Distribution of business cycle duration	99
7.14	Distribution of sentiment cycle duration	100
8.1	Trajectories of the OECD confidence index	105
8.2	Schematic illustration of household income	107
8.3	Behavior of consumption, wages and capital in the Leontief limit	110
8.4	Phase diagrams of the DSGE model	117
8.5	Phase dynamics of the adapted DSGE	120
8.6	Effect of memory timescales on dynamics	123
9.1	Economic dashboard for the three Central Bank scenarios described in sections 9.3, 9.3 in absence of any policy and shock	144
9.2	Empirical Shocks	146
9.3	Economic dashboard for the three shocks in the Inactive Central Bank scenario and in the absence of any policies	148
9.4	Counterfactual COVID shock in the Inactive Central Bank scenario and in the absence of any policies	150
9.5	Economic dashboard for a stronger COVID shock in the Inactive Central Bank scenario, with and without Easy-Credit policy	151
9.6	Unemployment and Inflation with all shocks, Easy-Credit policy, and three distinct monetary policy scenarios	154
9.7	Helicopter Money in the Reactive Central Bank with Floating Trust sce- nario, Easy-Credit policy, and all shocks	157
9.8	Windfall Tax in the Reactive Central Bank with Floating Trust scenario, Easy-Credit policy, and all shocks	158
9.9	Hyperinflation tipping points	160

List of Tables

3.1	Classification of eigenvector relation to model dynamics through eigenvalue dynamics as a function of the total number of timesteps based on Francis and Transtrum (2019, Table 1)	27
5.1	Parameterisation of the Mark-0 model	44
E.1	Overview of the General Structure of four Agent-based Integrated Assessment Models. Details on the energy and resource sector are in Table E.2, the climate module in Table E.3, and policy institutions in Table E.4 . . .	243
E.2	Overview of the Energy and Resource Module of four Agent-based Integrated Assessment Models	252
E.3	Overview of the Implementation of Climate Modules in four Agent-based Integrated Assessment Models	258
E.4	Overview of the Policy Experiments of four Agent-based Integrated Assessment Models	261
F.1	Parameters and Notation of the Dynamic Solow Model	271
G.1	Parameters and Notation of the Adapted DSGE model	278
H.1	Inactive Central Bank parameter set for Mark-0 Model	280

Introduction

Introduction

The macroeconomy can be seen as a complex adaptive system (Dosi and Roventini, 2019), where a network of millions of firms connected by physical ties interacts with a network of billions of households through a complicated web of trillions of contracts. The ecology of heterogeneous agents interacting locally leads to emergent aggregate dynamics such as successions of crises and recovery or economic growth, statistical regularities and institutional structures. In the words of Anderson (1972): “more is different” when it comes to understanding how macroeconomic dynamics emerge from micro-agents and their interactions. There has been an enduring debate on modeling approaches to capture these dynamics (see e.g. Vines and Wills, 2018, 2020, for a review), and relating them to microscopic properties of agents. One promising approach to capture the heterogeneity and multiplicity of agents, their interactions, and the resulting complex dynamics is found in agent-based computational economics (Tesfatsion, 2002; LeBaron and Tesfatsion, 2008; Fagiolo and Roventini, 2017; Dawid and Delli Gatti, 2018; Leijonhufvud, 2006; Haldane and Turrell, 2019; Caverzasi and Russo, 2018).

Agent-based models (ABMs) form the methodology of agent-based computational economics and generative social science (Epstein, 1999). By simulating an economy from “the bottom up”, ABMs start with a population of heterogeneous agents (e.g. multiple households, firms, financial institutions) endowed with their relevant behavioral rules, and a set of protocols for inter-agent interactions. This system is then simulated over time, tracking each agent’s state and decision, from which aggregate dynamics can then be directly computed. The evolution of the system based on agents’ behavioral rules gives rise to the complex dynamics a macroeconomic model should recover. Moreover, the modularity of these models implies that a variety of different macroeconomic policies can be explored, including the effects of policy mixes, which is often not possible using other modeling approaches (Van Den Bergh et al., 2021). This usefulness of ABM for policy analysis is demonstrated in the uptake of these models in policy institutions such as the recent CANVAS model at the Bank of Canada (Hommes et al., 2022), projects

at the Central Banks of England (Carro et al., 2022), Italy (Catapano et al., 2021) and Hungary (Mérő et al., 2022). However, there remain many open questions with respect to Macroeconomic ABM (MABM hereafter). In particular their calibration to data remains an active research area (Lamperti et al., 2018b) and there is an ongoing debate around using MABM for timeseries forecasting (see Poledna et al., 2023, for a first example) versus as an exploratory tool to uncover different scenarios and policy effects (Polhill et al., 2021). The non-linear out-of-equilibrium nature of these models means that analytical solutions are generally infeasible, and modelers rely on numerical simulations of the system to derive insight into the generated phenomena (Fagiolo and Roventini, 2017). This computational burden, together with the large number of parameters and initial conditions makes calibration and the full exploration of the parameter space an almost intractable problem.

In this thesis, I provide a new approach to addressing the problem of parameter space exploration in Macroeconomic Agent-based Models by asking:

*What is the set of phenomena that a Macroeconomic Agent-based Model can generate?
And what determines their dynamics?*

Inspired by research in physics on the ability to make predictions even when parameter uncertainty is high (see Quinn et al., 2023, for a review), I show that also for MABMs a given observable dynamic depends only on a handful of well-constrained parameter combinations, while any remaining parameter combinations can vary over multiple orders of magnitude without significantly changing the MABM's predictions. In a second step, I exploit these sensitive directions to sequentially perturb a MABM, thus generating a set of maximally different dynamics (the set of possible scenarios) in an efficient and informative way. Applying this method to the Mark-0 MABM (Gualdi et al., 2015, 2017; Bouchaud et al., 2018; Sharma et al., 2020; Knicker et al., 2023), I show how we can recover the full set of possible unemployment dynamics in the model.

The conclusion of these exercises is that despite their apparent complexity in terms of emergent phenomena, the dimension of the parameter space and initial values, the dynamics of a given observable can be efficiently explored. Moreover, the structure of the sensitive parameter-combinations can lead to insights on potential policy levers, such as understanding which model parameters can push the dynamics from endogeneous crises to full employment as demonstrated in Knicker et al. (2023). Additionally, the sensitivity of a given phenomena to only a handful of parameter-combinations means that these models can make effective predictions if these directions are well-constrained by the data. Individual parameter estimates can have large confidence intervals, but the model can nonetheless make accurate forecasts.

In the remainder of this chapter, I outline what it means to view the economy as a complex system (Section 1.1) and how this relates to current macroeconomic modeling efforts grounded in the workhorse general equilibrium approach (Section 1.2). Section 1.3 then outlines how an agent-based approach is a fruitful alternative that has already yielded multiple insights, yet also remains an open area of research with multiple gaps

that should be addressed. Finally, I outline how the elements presented in this Thesis address the overarching research question on scenario exploration in three distinct parts (Section 1.4), thus contributing to closing the research gaps in the methodology of agent-based models.

1.1 The Economy as a Complex System

Describing the economy as a complex system dates back to Prigogine and Stengers (1988), though it is often attributed to discussions at the Santa Fe Institute (Anderson et al., 1988; Arthur et al., 1997). There are multiple definitions of what it means to be complex, from structural interpretations based on the existence of intricate institutional structures and relationships (Pryor, 1996; Stodder, 1995), to computational ones based on the difficulty of determining solutions to optimization problems (Leijonhufvud, 1993; Stodder, 1997; Albin and Foley, 1998; Sargent, 1993), and philosophical perspectives where complexity cannot be “deductively defined but can only emerge inductively” from modeling (Rosser, 1999). Here I consider the definition based on Arthur et al. (1997) and developed at the Santa Fe Institute. In their view, summarized by Rosser, complexity implies six characteristics:

- (1) dispersed interaction among heterogeneous agents acting locally on each other in some space;
- (2) no global controller that can exploit all opportunities or interactions in the economy even though there might be some weak global interactions;
- (3) cross-cutting hierarchical organization with many tangled interactions;
- (4) continual adaptation by learning and evolving agents;
- (5) perpetual novelty as new markets, technologies, behaviors, and institutions create new niches in the “ecology” of the system;
- and (6) out-of-equilibrium dynamics with either zero or many equilibria existing and the system unlikely to be near a global optimum.

Rosser (1999)

Out of these attributes, it is also noted that the system will display an “emergent global structure from strictly local effects” (Rosser, 1999), such as the processes of economic growth and the evolution of the institutional framework that embeds the economy. The reason here for elaborating on the definition of what it entails for system to be complex system is not for the purpose of demonstrating that the economic system is a complex system, as I would venture this to be a sufficiently widely held belief, whether one uses the terminology of complexity theory or not. Rather, it is to pose the question of how one approaches the modeling of such a system in order to understand and explain its dynamics, design policy, communicate a story about the systems, act on different policy choices, predict its future path, or explore the set of possible behaviors.¹ The reason for this question being that “it is much more useful to characterize the economics profession

¹These modeling purposes follow the REDCAPE principle of Page (2018), but one can also refer to Epstein (2008).

as a diverse evolving set of ideas, loosely held together by its modeling approach to economic problems” (Colander et al., 2004).

1.2 The State of Macroeconomic Modeling

In many ways, the dominant approach to understanding, explaining and predicting macroeconomic dynamics is the use of the Dynamic Stochastic General Equilibrium (DSGE) model (Woodford, 2009). While a full description of a DSGE model is beyond the scope of this thesis, one might simplify to say that DSGE models are models of macroeconomic fluctuations based on *micro-foundations*, encapsulating the idea that “macro-economic behavior should be build up from the aggregation of the individual actions of self-interested, typically optimizing, agents” (Haldane and Turrell, 2019).² In the simplest sense the current wave of DSGE models are based on a real business cycle (RBC) core with New Keynesian features. The core real business cycle model generally features an “infinitely-lived representative household that seeks to maximize the utility from consumption and leisure, subject to an inter-temporal budget constraint” (Galí, 2015), together with a continuity of firms with identical technology and subject to exogenous shocks. The solution to this model is then a series of equilibria at each time step, whose values are fixed by the exogenous shock process.³ The modern New-Keynesian DSGE version adds to this monopolistic competition, nominal rigidities (also often called sticky-prices and/or sticky-wages), and the short-run non-neutrality of monetary policy (due to nominal rigidities, there is a role for monetary policy via changes in the short-term nominal interest rates).

This modeling approach emerged from the critique of Lucas (1976) that structural econometric models make implausible assumptions to fit the data, most famously, that they are invariant with respect to policy interventions. Specifically, agents in these models cannot adapt their behavior when policy—and thus their incentives—change, which is clearly important when one wishes to explain and understand the effects of a policy. To address this, the thinking goes that models should be micro-founded, which has since implied that they should be based on self-interested optimizing agents, more commonly known as “rational expectations” (Muth, 1961; Lucas, 1972, 1987; Lucas and Sargent, 1979; Kydland and Prescott, 1982). To this end, DSGE models were developed out of earlier Real Business Cycle models, such that monetary policy questions could be answered.

Not only is the use of these models standard among academics, but also among Central Banks (Yagihashi, 2020). Indeed, their wide-spread adoption might suggest that the issue of macroeconomic modeling has been resolved, and there is no role for agent-based

²In this thesis, I refer by DSGE to the New Keynesian variant, which is well-introduced in Galí (2015), with some seminal works of reference by Christiano et al. (2005) and Smets and Wouters (2004, 2007).

³The equilibrium solution here, representing the “Stochastic” and “General Equilibrium” parts of DSGE, is a series of equilibria, which makes the assumption that the system is adiabatic, i.e. coordination occurs faster than the exogenous process of shocks (Dessertaine et al., 2022).

approaches. However, the DSGE approach has faced heavy critique, implying room for alternative and complementary methods.

A Crisis in Macroeconomics?

The economics discipline itself is divided on the future prospects of DSGE models. Some believe that DSGE is the only game in town (Christiano et al., 2018)⁴, with authors such as Kehoe et al. (2018) stating: “[...] a disciplined debate rests on communication in the language of dynamic general equilibrium theory”. By contrast, there is a plethora of authors who view DSGE-approaches as a “dead end” (Dosi and Roventini, 2019) and believe that “for more than three decades, macroeconomics has gone backwards” (Romer, 2016) by focusing on this approach. Colander et al. (2004) make the argument that the *elite* of the field has been either suspicious or plainly rejected the neo-classical tenets of rational expectations with perfect foresight, a representative agent, and equilibrium solution, for a while or actually never accepted it but worked within the confines of the methodology. Early examples include Kenneth Arrow, Thomas Sargent, and Leonard Rapping, all of whom were involved in the creation of these models, later rejecting their premises.⁵ What is important to note here is that this approach to macroeconomic modeling has repeatedly *failed* to the point where Kirman (2010) deemed the inability of these models to predict the 2008 financial crisis as a crisis for economic theory.⁶

There are multiple avenues along which these models have been critiqued, both old and new.⁷ It is beyond the scope of this thesis to give a full accounting of all critiques and their responses, for which I refer to the Rebuilding Macroeconomics project of Vines and Wills (2018, 2020) and the works of Colander et al. (2008), Fair (2012), Romer (2016), and Dosi and Roventini (2019). It is illustrative however to note some of the frequently repeated criticisms because Agent-based approaches can address these well (Farmer and Foley, 2009; Dosi and Roventini, 2019; Fagiolo and Roventini, 2017; Leijonhufvud, 2006; Cincotti et al., 2022). Most notably these include the representative agent (Kirman, 2010, 1992, 2006) and its rationality (Sargent, 1993; Shiller, 2005), the absence of a financial system (Vines and Wills, 2018, and references therein), their weak empirical fits (Korinek, 2017; Romer, 2016; Fukac and Pagan, 2006)⁸, the use of equilibria as a solution concept (Hendry and Muellbauer, 2018; Benhabib and Farmer, 1999; Cass and Shell, 1983; Hirano and Stiglitz, 2022), and the ensuing contention that they have little usefulness in policy analysis (Stiglitz, 2018; Chari et al., 2009). Practitioners of DSGE

⁴Christiano et al. (2018) are incidentally also known for calling non-DSGE users dilettantes in the first draft of their paper, as recorded by Merler (2017).

⁵In fact Kenneth Arrow was one of the early organizers of the Santa Fe Institute workshops that launched the field of complexity economics (Anderson et al., 1988). Current examples include Paul Krugman, Robert Solow, George Akerlof, Joseph Stiglitz, and Larry Summers.

⁶Curiously, though understanding and preventing crises or *taming the business cycle* was a key focus of macroeconomics, these models’ inability to forecast the 2008 great financial crisis is occasionally defended by claiming that it was not their purpose. Note too though that within the model construction, a crisis of such magnitude was ruled out by design (Stiglitz, 2018).

⁷Summers (1986) is an early critique of the RBC foundation to New Keynesian DSGE.

⁸Korinek (2017) actually conclude that the scientific rigor of the method is questionable.

models may suggest that many of these critiques have been addressed, whether that be in the form of bounded-rationality (Branch and McGough, 2018), or Heterogeneous Agent New Keynesian (HANK) Models (Kaplan et al., 2018) and firm-network approaches (Baqaei and Farhi, 2019), yet despite these modifications it “seems likely that some features of economic systems will remain difficult to reproduce in a DSGE setting – for example, crisis dynamics” (Haldane and Turrell, 2019).⁹ Against the backdrop of mounting critiques and failures to predict the 2008-09 crises and post-COVID inflation, many have called for new approaches to modeling macroeconomic phenomena (Colander et al., 2008; Vines and Wills, 2018, 2020), or at the very least for DSGE models to become “less imperialistic” (Blanchard, 2018) and step away from the “micro-foundations hegemony” (Wren-Lewis, 2018). The aim being, at the very least, to improve accuracy of prediction by increasing model diversity (Page, 2007; Timmermann, 2006; Stock and Watson, 2006; Silver, 2013).

1.3 Agent-based Models: Alternative Foundations

Agent-based computational economics (ACE) involves the computational modeling of an economy considering all aspects of a complex system, as defined in Section 1.1 (Tesfatsion and Judd, 2006). Specifically, while a Macroeconomic Agent-based Model (MABM) is concerned with the dynamics of aggregate variables such as production, consumption or inflation, they “explicitly capture the micro-level interaction of different types of heterogeneous economic agents and allow computing the aggregate variables ‘from the bottom up’” (Dawid and Delli Gatti, 2018) in the spirit of “if you didn’t grow it, you didn’t explain it” (Epstein, 1999).

To construct a MABM, one specifies a large set of (multiple) types of agents (typically firms, households, banks, and a government). These agents are heterogeneous in their attributes, such as their endowments or their behavioral rules. Furthermore, their behavior is based on boundedly-rational heuristics using limited information (Sims, 1980; Gigerenzer and Brighton, 2009; Tversky and Kahneman, 1974).^{10,11} These agents then interact based on prescribed protocols influenced by the institutional structure into which their interactions are embedded, thus being “necessarily limited to locally constructive actions, that is, to actions constrained by their interaction networks, information, beliefs and physical states” (Sinitskaya and Tesfatsion, 2015).

⁹For example, in HANK models agents generally only interact through prices as opposed to the richer direct interactions in ABMs. Furthermore, many of the additions and adaptations seem to address one or two of the issues, but it appears that none offer a comprehensive response.

¹⁰One may argue that ABM is exposed to the “wilderness of bounded rationality” (Sims, 1980) as there is no axiomatic basis for behavior, but this can be counteracted by considering evolving rules, as well as experimental (Hommes, 2013), and empirical observation of actual micro-economic behavior (Dawid and Delli Gatti, 2018), or the comparison of behavioral rules within a fixed framework (e.g. Sinitskaya and Tesfatsion, 2015).

¹¹In a world of knightian uncertainty (Knight, 2012), these heuristic approaches may actually be more ‘rational’ than the type of self-interested optimization of a DSGE model (Gode and Sunder, 1993; Hommes, 2006).

The MABM is then numerically simulated over a series of discrete time steps, and the aggregate properties calculated directly from the summation over the individual agents. A numeric Monte Carlo approach is necessary here as one needs to simulate a non-linear stochastic process for each agent. The numeric nature also implies leaving the restriction on complexity enforced by needing an analytical solution (Colander et al., 2008), and does not rely on an equilibrium solution concept.¹² As suggested in the properties of a complex system, this allows for dis-equilibrium dynamics such as the emergence of endogenous crises. A complete (agent-level) equilibrium as in DSGE remains one possible solution, but it is much more likely that there are only system-level equilibria wherein the system as a whole has a steady state but each agent may be continually evolving. The importance of these disequilibrium dynamics in understanding and predicting crises has been raised by authors such as Krugman (2011), inspired by earlier work such as that of Minsky (1976, 2008) and Leijonhufvud (2000). Furthermore, the direct simulation of all agents and only thereafter aggregating leads to emergent phenomena, such as self-organization towards a long-term growth path, endogenous crises and business cycles (no longer relying on large exogenous shocks), and persistent heterogeneity such as in the distribution of firm sizes or incomes (Dawid and Delli Gatti, 2018).

Comparing the description of MABM from the previous paragraph to the attributes of a complex system proposed by Rosser (1999) (see Section 1.1), one can see that each of the six critical avenues is, to varying degrees, captured in the construction and execution of an MABM. This is in contrast to a DSGE approach that, as highlighted in Section 1.2, covers only a small subset of these properties at a time, and whose solutions provide insight only when the system itself is “dealing with close to equilibrium fluctuations” (Haldane and Turrell, 2019).

In a recent review of macroeconomic modeling by Vines and Wills (2018, 2020) following the failures in predicting and addressing the 2008-09 Financial Crises with the tools of the time, Ghironi (2018) and Vines and Wills (2018) identify four important areas of macroeconomics that should be addressed (see also Haldane and Turrell, 2019): (1) Financial Systems¹³, (2) Heterogeneity in households¹⁴, (3) Granularity and Networks¹⁵, and (4) Policy Interdependence.¹⁶ Farmer et al. (2015), Hafner et al. (2020) and Castro et al. (2020) added the ability to do interdisciplinary modeling with respect to climate

¹²All ABMs that can be computed numerically have an explicit mathematical representation, but this generally too complicated to be useful (Leombruni and Richiardi, 2005; Epstein, 2006).

¹³see Ashraf et al. (2017); Assenza and Delli Gatti (2019); Assenza et al. (2015); Banwo et al. (2019); Botta et al. (2020); Cardaci and Saraceno (2019); Cincotti et al. (2010); Delli Gatti et al. (2010); Gabbi et al. (2015); Popoyan et al. (2020); Raberto et al. (2019); Riccetti et al. (2021); Reissl (2021) for some examples.

¹⁴see Caiani et al. (2019b); Mellacher and Scheuer (2020); Palagi et al. (2017, 2021); Papadopoulos (2019); Rengs and Scholz-Wackerle (2019); Russo (2017); Russo et al. (2016) for some examples.

¹⁵see Gualdi and Mandel (2019); Wolf et al. (2013); Stellan et al. (2021); Otto et al. (2017); López et al. (2020); Gatti et al. (2009) for some examples.

¹⁶see Borsato (2020, 2021); Fagiolo et al. (2020); D’Orazio (2019); Dosi et al. (2013, 2015, 2021, 2019, 2018); Lamperti et al. (2021, 2019a, 2018a, 2020); Raberto et al. (2008); Popoyan et al. (2017); Riccetti et al. (2018) for some examples.

change to this list.¹⁷ Agent-based approaches have responded to all of these key points, with some doing so simultaneously. This alone suggests they should be part of the general toolkit of macroeconomics.

While MABMs can model and replicate the dynamics of a complex economy, they are no panacea, after all “all models are wrong, but some are useful” (Box, 1976). When it comes to MABM much of the current debate surrounding the method rests around the estimation and calibration, and their use in forecasting, with many deeming them “black-boxes” with causal mechanisms that are hard to determine and that are extremely sensitive to parameterization. This is complicated by the fact that many of the medium and large MABMs are computationally very expensive, creating a high barrier to effectively estimating their underlying parameters. The estimation and calibration of these models remains an active area of research, with most models following an “indirect calibration approach” (Fagiolo et al., 2019) of replicating a wide range of stylized facts (see Haldane and Turrell, 2019, for a non-exhaustive list). Beyond this, the literature has expanded on the replication of time series (Lamperti et al., 2018b; Barde, 2016) and estimation of parameters using Bayesian approaches (Platt, 2020, 2021). Similarly, in the domain of sensitivity analysis advances have been made using meta-analysis techniques (ten Broeke et al., 2021). It is in this domain of exploration, sensitivity and calibration that the work presented in this thesis adds.

1.4 This Thesis

This thesis is split into multiple distinct but interrelated parts covering my research on the question of scenario generation over the past three years. In Part I, I develop an algorithmic approach to exploring different scenarios in Agent-based Models and show how it can recover the phase-space of the Mark-0 Agent-based Model. Part II of the thesis focuses on two models of interacting agents embedded within neoclassical macroeconomic closures to show how emergent phenomena from these interactions can enrich the phenomenology of these models, together with an application of the Mark-0 MABM to the post-COVID inflation surge.

A major critique of Agent-based Models is that they are seen to be *black boxes* wherein the mechanisms that lead to emergent phenomena are unclear, the set of dynamics a model can generate is unbounded, and due to the large number of parameters fitting empirical phenomena is futile. Part I of this thesis provides a step towards dissolving these concerns and shows that there is a finite set of qualitatively distinct phases in a given ABM, and the local dynamics depend only on a few parameters. To do so, Chapter 2 outlines an approach to fitting and understanding large-parameter models developed in the biophysics literature. This approach, dubbed *sloppiness* by its originators, highlights that the generated dynamics of high-dimensional models are often sensitive to only a handful of well-constrained parameter combinations. On the other hand, many

¹⁷see e.g. Lamperti et al. (2019a, 2021, 2018a, 2020); Wolf et al. (2013); D’Orazio and Valente (2019); Otto et al. (2017); Ponta et al. (2018); Safarzyńska and van den Bergh (2017b, 2022).

other parameter combinations are ill-constrained, they can vary over multiple orders of magnitude without a significant change in observables. Identifying and classifying these parameter-combinations can shed light on the local landscape of parameter-sensitivity, allowing us to exploit the most well-constrained directions to explore a model’s phase space by maximally changing the observable dynamics. Moving from theory to practice, Chapter 3 quantifies this concept for ABMs, taking two distinct measures to quantify a change in observables: first, a mean-squared error approach that directly compares time series realizations and is intended to work on shorter timescales of interest to policymakers, and second, the symmetric Kullback-Leibler divergence as a probabilistic approach to compare the distribution of observables which may be more robust for non-linear stochastic systems like ABMs. To illustrate the mechanics of this approach, Kirman’s *Ants*, a model of agent-herding is used as a theoretical example in Chapter 4. Both the MSE-based and KL-based approaches are evaluated analytically and numerically on the whole set of parameters, and recover the parameter-combination representing a phase transition in the model. In Chapter 5, I turn to the analysis of the Mark-0 Agent-based Model that was the subject of our first paper (Naumann-Woleske et al., 2023). The Mark-0 model is a simple MABM that nonetheless displays a qualitatively rich set of phenomena. In particular, there are four distinct phases of the unemployment rate, ranging from full employment to residual unemployment, endogenous oscillations and total unemployment. Evaluating the well-constrained parameter directions of Mark-0 at different points in these phases generally recovers the parameters responsible for phase transitions, implying that perturbations into these directions may be an effective means of uncovering the different scenarios an ABM can generate. Exploiting these parameters in a simple gradient-ascent algorithm, I show in Chapter 6 that one can effectively and efficiently explore the phase-space of the Mark-0 model, not only recovering the different phases but also gaining an understanding of the localized sensitivities of these dynamics. In particular, one can get an impression of some of the parameter combinations that drive different phase transitions.

Part II of this thesis takes a different approach to exploring the economy as a complex adaptive system. Specifically, I here consider heterogeneity and non-rational behavior in two mainstream settings, a simple Solow model (Chapter 7) and New Keynesian Dynamic Stochastic General Equilibrium approach (Chapter 8), as well as the effects of confidence in the Mark-0 Agent-based Model (Chapter 9).

The first two models in this part represent investigations into how sentiment-driven investment decisions can lead to cyclical aggregate phenomena, as inspired by the notion that “the subjective evaluation of prospects over a time horizon is the major proximate basis for investment and portfolio decisions, and these subjective estimates are changeable” (Minsky, 1976). Methodologically, this exploration aims to bring together an agent-based and classical equilibrium-based approach to highlight how even small non-rationalities and interactions of agents can induce a whole host of qualitatively different phenomena that would otherwise be overlooked. Note that these are *toy* approaches, they are not intended to be calibrated, but to demonstrate that even simple interactions

can lead to emergent phenomena that are relevant for model uses. Chapter 7 develops a *Dynamic Solow Model*, introducing capital and investment into a simple Solow growth model, replicating our work in Naumann-Woleske et al. (2022). Investment depends on a large number of individual investors, who form an expectation of future returns based on the state of the economy as well as the expectations of other investors. Drawing on tools from statistical physics, this system can be written as a set of non-linear differential equations of the aggregate outcomes on the basis of micro-interactions, where the effect of the feedback of other investors' sentiment leads to cycles of exuberance and pessimism of heterogeneous duration. Albeit a simple model, we identify a business cycle of quasiperiodic fluctuations around a steady equilibrium growth-path in a phase where the economy is led by investment demand. These cycles can be identified as a coherence resonance mechanism, wherein the exogenous noise leads to quasiperiodic fluctuations, as in the “small shocks, large cycle” effect of Bernanke et al. (2019). In turn, Chapter 8 replicates our work in Morelli et al. (2021), by further developing the extended New Keynesian model of Morelli et al. (2020), who developed an adapted New Keynesian model where a large number of households are characterized by a sentiment on the state of the economy. Much like in the Dynamic Solow Model, households' sentiment is influenced by their neighbors, which leads to a “keeping up with the Joneses” effect, and, again using tools from statistical physics, can be written as a modified representative agent approach with non-linear expectations. Into this model, we introduce capital as a factor of production, allowing for the behavioral household to invest savings into a risky asset providing capital to firms. This addition leads to a behavioral real business cycle model with a rich phenomenology, wherein capital scarcity can lead to prolonged periods of low-output until confidence is restored, highlighting the importance of messaging by governments. Specifically, the role of confidence in investment here is similar to Keynes' paradox of thrift: lower confidence leads to more saving, exacerbating economic downturns, meaning that in addition to messaging to restore confidence, policymakers may consider direct investment plans (such as the American Rescue Plan and the Infrastructure Investment and Jobs Act in 2021) The model is, again, a bare-bones toy model intended to highlight the richness of macroeconomic phenomenology emerging from the interactions of individual agents. In Chapter 9 we then take a closer look at confidence in the context of the increase in inflation following the COVID-19 pandemic in 2020. With inflation peaking at 9.1% in the US in June 2022, there have been competing narratives to explain this phenomenon: from “too much money chasing too few goods” to energy and food prices, profit-driven inflation and conflict inflation. To understand some of the inflationary scenarios that might arise and link them to different explanations and policy responses, we extend the Mark-0 model with a stylized energy sector and introducing three shocks to simulate exogenous drivers and competing narratives.¹⁸ We find that the economic recovery following COVID-19 is sluggish in the absence of mitigating policy, taking at least several years. Yet, there is only a narrow window for effective monetary policy, responding with a strong interest rate hike disproportionately raises

¹⁸Including a COVID shock affecting productivity and demand, a supply-chain shock on productivity, and an energy-price shock.

unemployment. On the other hand, the de-anchoring of agents' trust in the Central Bank's ability to reign in inflation influences the dynamics of inflation itself by influencing the wage-bargaining process, the price-setting process and households' consumption decisions, leading to higher inflation and worse economic performance.

PART I

Exploring the Parameter Space in Macroeconomic
Models

How Complex Are Complex Models

I remember my friend Johnny von Neumann used to say, with four parameters I can fit an elephant, and with five I can make him wiggle his trunk.

*Enrico Fermi in conversation with Freeman Dyson
(2004)¹*

Confronted with highly parameterized models, and a limited set of observed macro and microeconomic data, economists and other scientists can nonetheless make accurate predictions. One current example of this is the use of large neural networks for classification tasks. Despite adding thousands of parameters per layer, the predictive performance of these models does not appear to decline, as is typically the case when overfitting (Kaplan et al., 2020). One reason for this phenomena may be that these models are *sloppy*: the predicted phenomena depend only on a handful of parameter combinations, while the bulk of the remaining parameter-combinations has little to no effect on the predicted outcomes. This intuition is particularly important for macroeconomic agent-based models, which can generate complex phenomena but also often depend on a large number of parameters, especially if one considers the initialization of a large number of heterogeneous agents. Understanding stiff and sloppy parameter-combinations in these models would allow researchers to develop a deeper understanding of the space of possible model predictions, the sensitivity of these predictions to parameter choices, and methods for the calibration of these models.

The premise of the first part of this thesis is to build a method to approach the question: *what are the different types of dynamics that an agent-based model can generate?* To answer this question, one can exploit the fact that high-parameter MABMs are sloppy with

¹Mayer et al. (2010) actually showed it is possible to generate *elephantine* shapes with four complex numbers.

respect to their predictions. In particular, one can use the critical sensitive parameter-combinations to explore a model’s space of predictions and thereby develop a phase diagram of possible phenomena in this model, and understand what each phenomena depends on. In this chapter, I will first outline what it means for a model to be sloppy (Section 2.1) and what this implies for the predictions and parameterization of a model. Pursuing this, in Section 2.2, I outline how one can use the intuition of sloppy models to approach two key challenges for macroeconomic agent-based models: first, the idea that MABMs can generate any phenomena and obfuscate its causal mechanisms, and second, the difficulties in calibrating MABMs. In doing so, I outline how this thesis addresses the first of these challenges by building up an exploratory tool for MABM, and how this complements existing methods of global sensitivity analysis in MABM.

2.1 Predictive Power Despite Uncertainty

The notion of *stiff* and *sloppy* parameter directions was first introduced by Brown and Sethna (2003) while studying models of biochemical regulation. This type of model has a high number of unknown parameters brought together in a large set of nonlinear differential equations, yet it can nonetheless make accurate predictions even when there are large uncertainties in each of the individual parameters (Brown et al., 2004). For each of the individual parameter estimates, “95% confidence intervals each spanned more than a factor of 50” (Gutenkunst et al., 2007). The reason that a model with extreme individual parameter uncertainty can still provide accurate predictions is that the model is *sloppy*. Specifically, Brown et al. (2004) find that there are a few *stiff* parameter-combinations that are well-constrained by the data, which in turn constrain the predictions of a model. Meanwhile, there is a large variety of *sloppy* parameter-combinations that have little effect on the model’s predictions, and are generally not well constrained, thus allowing them to vary over great ranges without affecting the prediction. The result is that individual parameters may vary to a large degree while predictions are constrained. Not only do the well-constrained stiff directions lead to accurate predictions, they “reveal critical focal points in the signaling network” (Brown et al., 2004), suggesting that their analysis can lead to a better understanding of critical drivers of model outcomes.

Figure 2.1 presents a schematic interpretation of a sloppy model, as shown also by Brown et al. (2004) and Gutenkunst et al. (2007). Ellipsoid lines represent points of equal loss in the model’s prediction with respect to the optimal point Φ^* at the center, and axes of these lines are generally not aligned with the bare parameters. Instead, they are combinations, such as the ones given here by the v_1 and v_2 arrows. In this case, the v_1 vector represents a stiff direction, where little variation is needed to drastically change the model’s prediction as captured by a loss function. Meanwhile, the v_2 direction allows for comparatively larger variation in parameters while not changing the prediction, these directions are deemed sloppy. In the models studied by Brown et al. (2004), it is the stiff directions that are well-constrained by their data, and thus even with a large number of parameters and a low number of observations they can gain accurate predictions from

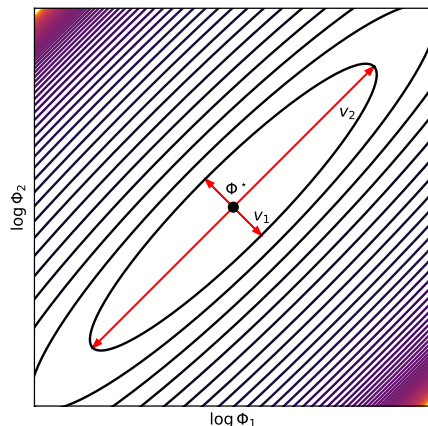


Figure 2.1 – Schematic interpretation of a two-dimensional sloppy model based on Brown et al. (2004, Figure 4) and Gutenkunst et al. (2007, Figure 1A). Colored ellipses present equi-loss lines for a generic loss function by which the model’s quality of fit is measured. The axes are the bare log-parameter axes of a model with parameters $\log \Phi_1$ and $\log \Phi_2$. The red arrows on the axis of the ellipse represent vectors in parameter space that are stiff ($v_1(\Phi)$) and sloppy ($v_2(\Phi)$).

their model. It is important to note here that in practical applications, this ellipse will be of a much higher dimension, non-linear and typically have a ribbon-like structure with wider variety in the radii of the ellipse.

Sloppiness as a feature of models is not limited to the realm of biochemical models. Since Brown and Sethna (2003) introduced this term, research in several fields has indicated that the stiff-and-sloppy parameter space structure appears across a wide range of models. In their review, Quinn et al. (2023) point out that this structure has been found in cell signaling (Brown et al., 2004), radioactive decay and neural networks (Transtrum et al., 2011), Quantum Wave Functions and insect flight (Waterfall et al., 2006), the Ising model (Machta et al., 2013), Meat Oxidation, Cosmic Microwave Background Radiation (Quinn et al., 2019), an Energy Recovery Linear Accelerator (Gutenkunst, 2008), a model of the circadian clock (Daniels et al., 2008), power systems, a gravitational model and transmission loss in an underwater environment. On the flip-side, it should be noted that not all multi-parameter models are necessarily sloppy, for instance Waterfall et al. (2006) point out that multiple linear regression, while it has a high number of parameters does not have a sloppy hierarchical structure. Nonetheless, for most non-linear least-squares type problems, they suggest that this type of hierarchy of parameter combinations emerges.

At this point, one may wonder whether it is possible to *cure* sloppiness by removing sloppy directions to reduce the degrees of freedom, or reparameterizing the model. It is very important to clarify at this point that if a model is sloppy, depending on only a few stiff parameter combinations, this does not necessarily entail that many individual parameters are irrelevant and that one might remove parameters not found in the stiffest direction to reduce a model’s degrees of freedom. As pointed out by Gutenkunst et al.

(2007), stiff directions are often vectors in parameter space that include many non-zero elements, a finding confirmed also in later articles, such that naively removing some parameters is infeasible. Turning instead to parameter space transformations, Waterfall et al. (2006) point out that for some simple model structures one can transform the parameter space to be non-sloppy; however, for many applied models this is actually quite difficult because (a) the transformations are non-linear depending on the point in parameter space, and (b) the natural parameterization of these models is pre-determined by the respective science, such as through rate constants in chemistry or elasticities of consumption in economics, while a linear combination of all model parameters has no economic meaning.

2.2 Are Macroeconomic Agent-Based Models Sloppy?

Macroeconomic Agent-based models fit the bill of a modeling approach that is likely to display a sloppy parameter-structure: (a) these models typically involve a large number of parameters, especially when one also considers the possibility for each of the heterogeneous agents to be individually parameterized or initialized, (b) these models have a very non-linear structure, for instance in the use of Leontief production functions (maximum functions) or agents interacting with threshold based choices. These features of MABMs have led many to critique this modeling strategy as “ad-hoc” black boxes, capable of generating a multitude of phenomena, and where “the causes and mechanism driving results are blurred” (Napoletano, 2018) (see the Introduction for more detail). A key challenge for MABM is thus to understand how sensitive predictions are to the selection or parameter values, and which mechanisms in the model drive the predictions. In addition, it is critical to understand the space of possible predictions for a given MABM (Macal, 2016; Crooks et al., 2008; Magliocca et al., 2018).

The aim of this part of the thesis is to take the sloppy analysis approach outlined above to understand the space of possible phenomena that a MABM can generate, as well as the sensitivity of these phenomena to parameterization. Putting this in terms of sloppy models, my aim is to explore and characterize a given MABM’s model manifold. The concept of a model manifold arises from information geometry, where one can “view the space of all possible model predictions as forming a manifold, whose coordinates are the model’s parameters ” (Quinn et al., 2023). First introduced by Transtrum et al. (2010) and Transtrum et al. (2011), the model manifold is typically a *hyperribbon*, an ellipsoid shape whose size in different dimensions varies of multiple orders of magnitude. Here, sloppy directions are short distances on the manifold, while stiff parameter combinations form long directions along the manifold. Consequently, one could exploit stiff directions to develop an exploratory path that intelligently samples the diversity of predictions on the model manifold.²

²One can also traverse the sloppiest directions to determine effectively simpler theories for a given phenomena. Transtrum (2014) developed this approach through the Manifold Boundary Approximation method, an agnostic data-driven approach to model reduction.

The issue of searching the parameter space of ABMs has received considerable attention in recent years, and several other global sensitivity approaches have already been developed and applied. However, to the best of my knowledge, no other modelers have considered the sloppy analysis approach to the problem. The surrogate-modeling approach proposed by ten Broeke et al. (2021) is the closest alternative to the sloppy approach proposed in this thesis (see also Lamperti et al., 2018b; Zhang et al., 2020; van der Hoog, 2019; Salle and Yıldızoğlu, 2014, for further surrogate approaches).³ In this approach, the parameter space is sampled using quasi-random methods (e.g. Sobol sampling). The model predictions for each of these parameter sets is then fed into a surrogate model, such as a Support Vector Machine (ten Broeke et al., 2021) or neural network (Lamperti et al., 2018b). In doing so, the surrogate is trained to predict the outcomes of the underlying model, and can be used to predict the outcomes for unsampled parameters. The advantage here is computational, as predicting with the surrogate is computationally cheap in comparison to running the original ABM while making reasonably accurate predictions of model outcomes (Lamperti et al., 2018b). However, these approaches require ex-ante insight by the modeler about which phenomena to investigate (e.g. seeking fat-tailed GDP growth distributions in Lamperti et al. (2018b) or fishery survival in ten Broeke et al. (2021)). By contrast, the sloppy approach remains relatively agnostic, allowing for a large exploration of different phenomena before requiring the modeler to assess how to classify these phenomena into different types of phases. An agnostic approach is crucial when some emergent behavior is truly unexpected, as is indeed often the case (see the detailed discussion in Gualdi et al. (2015)). In this respect, starting with the exploratory approach presented in this thesis would complement the meta-modeling approach by suggesting to modelers which phenomena appear across the prediction space and how these relate to the model parameterization. Additionally, MABM will benefit from the continued development of automatic differentiation tools (see Andelfinger, 2023; Chopra et al., 2022; Quera-Bofarull et al., 2023a,b, for inspiration) this means that working on parameter space exploration directly at the model level will become computationally inexpensive itself. I show this with the Mark-0 model in Chapter 5.

A second critique of MABM that may be addressed when considering a sloppiness approach regards model calibration and consequent use for prediction.⁴ While this is not the focus of this thesis, the application of a sloppy approach to experimental design and calibration for MABMs is nonetheless a fascinating avenue for continued research. This is for two reasons, the first being that understanding the sloppy hierarchy of parameter combinations can lead to an efficient design of data-based experiments to constrain all individual parameters. Apgar et al. (2010) develop a method to constrain all parame-

³Other Global Sensitivity Analyses don't scale well to high-parameter methods and may overlook parameter interactions (ten Broeke et al., 2021), and might be why sensitivity analyses are not often performed (Thiele et al., 2014).

⁴I am referring here to prediction of quantitative phenomena in general, though one could also narrow this down to point-forecasts of macroeconomic time series, a point on which MABM practitioners have been cautious as it is also often not the purpose of an ABM (Epstein, 2008).

ters to within 10% of their value in a sloppy model, though they find that high precision (e.g. $\leq 1\%$ may require an infeasible amount of complementary experiments) (Hagen et al., 2013, see also). The second relates to the use of Bayesian methods in MABM calibration. These methods have become more popular in recent years (see Dyer et al., 2022; Grazzini et al., 2017; Platt, 2020, 2021; Lux, 2018, 2022). The performance of these methods relies on their choice of prior distribution, from which the estimation procedure is started. Common priors, such as the Jeffrey’s prior, may lead to a strong bias in models whose effective dimensionality is lower, i.e. in sloppy models (Abbott and Machta, 2023). Instead, one can use the intuition and mechanisms of the sloppy analysis to construct a prior focusing on the most relevant parameters and leading to an unbiased posterior (see Mattingly et al., 2018; Abbott and Machta, 2023, for details).

Conclusion

In conclusion, the phenomena of sloppy models is common to high-parameter models and presents itself through a hierarchy of parameter combinations driving a given model phenomena. This has implications for understanding the mechanisms driving a phenomena, and the sensitivity of a phenomena to the choice of parameters. MABMs, as high-parameter non-linear models, fit this bill. As will be demonstrated in this thesis, one can exploit the sloppy structure of the model to effectively traverse the space of all possible model phenomena, thus clarifying what one can generate with an MABM.

Key Messages

- Sloppiness is common in high-parameter models, and implies there is a hierarchy of parameter combinations that affect a model’s prediction.
- Data generally leads to well-constrained stiff directions, which is sufficient for accurate model predictions. Thus even with a high-dimensional parameter space and little data, accurate predictions can be made.
- Determining the stiff parameter combinations also means one can exploit them to traverse the model manifold, which is the manifold of all possible model predictions, to answer the question: what can this model generate?

Quantifying Sloppiness

A model is sloppy when there exists a hierarchy of parameter-combinations to which the outcomes are successively less sensitive. Chapter 2 introduced the intuition behind this property and its implications. In this chapter, I operationalize these concepts and introduce the notation used in the remainder of Part I. In particular, we are looking for stiff (well-constrained) parameter combinations to which the model is highly sensitive, and sloppy (ill-constrained) parameter combinations to which model predictions are insensitive. Taking a general loss function to quantify the change in model predictions due to a change in parameters, I show that the eigendecomposition of the Hessian matrix of this loss function yields the set of orthogonal parameter combinations and their relative sensitivities (Section 3.1). Pursuing this, I introduce two simple loss functions that will be used going forward: the mean-squared error for comparing directly two timeseries (Section 3.2) and the symmetrized Kullback-Leibler divergence to compare two distributions (Section 3.3). For each of these loss functions, I present their Hessians, and how they are computed numerically, and some intuition about their relative uses.

3.1 Identifying Stiff and Sloppy Directions

Consider a loss function, $\mathcal{L}(\Phi, \Phi + \delta\Phi)$, that quantifies the degree of change in a model prediction \mathbf{x} when moving from a point Φ in P -dimensional parameter space to a second point $\Phi + \delta\Phi$. The second-order Taylor expansion of the loss function for a change in parameters from Φ to $\Phi + \delta\Phi$ evaluated at the point $\delta\Phi = 0$ reads

$$\mathcal{L}(\Phi, \Phi + \delta\Phi) \approx \underbrace{\mathcal{L}(\Phi, \Phi)}_{=0} + \underbrace{\delta\Phi^\top \nabla \mathcal{L}(\Phi, \Phi)}_{=0} + \frac{1}{2} \delta\Phi^\top H(\Phi) \delta\Phi,$$

where I use the fact that by definition the loss function without parameter change is zero to eliminate the first term, and noting that for $\delta\Phi = 0$ we are by definition at a

minimum yields $\nabla\mathcal{L}(\Phi, \Phi) = 0$, and $H(\Phi)$ is the Hessian of the loss function evaluated at point Φ . The i, j -th element of the Hessian matrix, $H(\Phi)$, is defined as

$$H_{i,j}^{\mathcal{L}}(\Phi) := \left. \frac{d^2\mathcal{L}(\Phi, \Phi + \delta\Phi)}{d\log\Phi_i d\log\Phi_j} \right|_{\delta\Phi=0} = \Phi_i\Phi_j \left. \frac{d^2\mathcal{L}(\Phi, \Phi + \delta\Phi)}{d\Phi_i d\Phi_j} \right|_{\delta\Phi=0}, \quad (3.1)$$

where the derivative is taken with respect to the log-parameters to generate relative parameter changes, which avoids issues with different parameter scales and units. The analysis of this Hessian then “corresponds to approximating the surfaces of constant model behavior deviation to be P-dimensional ellipsoids” (Gutenkunst et al., 2007), as schematically presented in Figure 2.1 (Chapter 2). Of course, this represents a quadratic approximation to the true cost surface, which may be very non-linear. But for $\delta\Phi$ being close to zero, this method clearly identifies the critical directions already.

The axes of the equi-loss-ellipsoids are given by the eigenvectors of the Hessian $H(\Phi)$, and their respective widths by the eigenvalues, which can be seen when making the following decomposition. Assuming all values in the Hessian are real (which is the case for all models considered in this dissertation), it can be decomposed as $H(\Phi) = Q\Lambda Q^\top$, with $Q = [v_1(\Phi), \dots, v_P(\Phi)]$ a matrix whose columns are the real eigenvectors $v_i(\Phi)$, and $\Lambda = \text{diag}(\lambda_1(\Phi), \dots, \lambda_P(\Phi))$ the diagonal matrix of their associated eigenvalues sorted such that $\lambda_i(\Phi) > \lambda_j(\Phi) \forall j > i$. This implies the loss can be approximated as

$$\mathcal{L}(\Phi, \Phi + \delta\Phi) \approx \frac{1}{2} \delta\Phi^\top Q \Lambda Q^\top \delta\Phi,$$

from which it becomes clear that the eigenvalue with the largest associated eigenvalue constitutes the stiffest direction as it induces the largest change in the loss-function. By comparison, the eigenvector with the smallest eigenvalue constitutes the sloppiest direction. To see this, consider that a unit vector step can be expressed as a linear combination of $\sum_n \alpha_i v_i$ s.t. $\sum_n \alpha_i = 1$, which implies that the change in loss is $\frac{1}{2} \sum_n \alpha_i^2 \lambda_i$ that is maximized for $\alpha_1 = 1, \alpha_{i \neq 1} = 0$, and minimized for $\alpha_n = 1, \alpha_{i \neq n} = 0$. Consequently, to induce a change l in the loss implies taking a step whose size is proportional to the eigenvalue of the chosen eigenvector direction,

$$\mathcal{L}(\Phi, \Phi + \delta\Phi) \approx l \quad \text{for} \quad \delta\Phi = \frac{\sqrt{2l}}{\sqrt{\lambda_i(\Phi)}} v_i(\Phi) \quad (3.2)$$

in the direction of eigenvector i . This highlights that distance (i.e. $\delta\Phi$) in the metric space formed by the Hessian is “a fundamental measure of distinguishability of stochastic systems” (Machta et al., 2013).

A natural question at this point is how spread out the eigenvalues need to be for a model to be sloppy: How insensitive to sloppy directions vis-a-vis stiff directions should a model be to be considered sloppy? The literature considers a model sloppy if the set of eigenvalues spans multiple decades in a roughly uniform manner on a logarithmic scale. The most common method for evaluation is the eigenvalue spectrum at a point Φ

$$\mathcal{S}_\Phi = \frac{\lambda_P(\Phi)}{\lambda_1(\Phi)} = \frac{\min_i \lambda_i(\Phi)}{\max_i \lambda_i(\Phi)}, \quad (3.3)$$

where the heuristic of calling a model-parameter combination sloppy if $\mathcal{S}_\Phi < 10^{-6}$ is often applied (Waterfall et al., 2006). In principle, one should adjust for the total number of parameters here, analyzing

$$\mathcal{S}_\Phi^{strong} = \left(\frac{\lambda_P(\Phi)}{\lambda_1(\Phi)} \right)^{\frac{1}{P}}, \quad (3.4)$$

for if \mathcal{S}_Φ is small but there are many parameters (P large) then the ratio between the first and second eigenvalue would be of order \mathcal{S}_Φ^{adj} , which may be large. Thus it would be desirable not just for \mathcal{S}_Φ , but also $\mathcal{S}_\Phi^{strong}$ to be small, with the latter being a *strong* criterion for sloppiness vis-a-vis the former. Jagadeesan et al. (2023) propose a more general formulation to quantify the degree of sloppiness that suggests a model is (ε, δ) -sloppy at a point Φ^* with respect to a subset $\mathcal{P}_{sloppy} \subset \mathcal{P}$ of the parameter space if for all points in \mathcal{P}_{sloppy} at a distance of more than δ to Φ^* , the loss is below a threshold ε . This is a natural superset to the definition using \mathcal{S}_Φ (consider $\varepsilon \rightarrow 0$ and $\delta \rightarrow (\min_i \lambda_i(\Phi))^{-\frac{1}{2}}$), that considers an infinitesimal change in parameters by means of the Hessian. For this thesis, I stick with the convention of using \mathcal{S}_Φ .

3.2 Non-linear Least Squares

A simple first loss-function to consider is the mean-squared error (MSE) formulation, which has been used across a variety of applications (e.g. Brown and Sethna, 2003; Brown et al., 2004; Gutenkunst et al., 2007). The MSE loss for a vector of K observables $\mathbf{x}_{s,t}(\Phi)$ at discrete time steps t for random realization s and parameter point Φ in comparison to $\mathbf{x}_{s,t}(\Phi + \delta\Phi)$ is

$$\mathcal{L}^{MSE}(\Phi, \Phi + \delta\Phi) = \frac{1}{2TS} \sum_{t \in T} \sum_{s \in S} (\mathbf{x}_{s,t}(\Phi) - \mathbf{x}_{s,t}(\Phi + \delta\Phi))^\top \Sigma^{-1} (\mathbf{x}_{s,t}(\Phi) - \mathbf{x}_{s,t}(\Phi + \delta\Phi)), \quad (3.5)$$

where, in a slight abuse of notation, we have $T = \{t_1, \dots, t_T\}$ time steps and $S = \{s_1, \dots, s_S\}$ random realizations. Note too that $\mathbf{x}_{s,t}(\Phi)$ is a K -dimensional vector of observations, making the prediction space for most models $\mathbb{R}^{K \times T}$, with a parameter space $P \subset \mathbb{R}^P$, and Σ is a diagonal matrix that leads to a dimensionless loss. In this thesis, I will use $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_K^2)$, the diagonal matrix of the variable specific variances evaluated across T and S for parameters Φ . In a more general case, one might consider the $K \times K$ covariance matrix.

A current limitation of Agent-based Models is their computational burden, in terms of time, memory and processing power (this has also motivated other approaches such as surrogate modeling as in ten Broeke et al., 2021; Lamperti et al., 2018b). Thus estimating $H^{MSE}(\Phi)$ directly would likely be prohibitive at $\mathcal{O}(SP^2)$ model evaluations. However, as we are evaluating the Hessian around Φ , i.e for $\delta\Phi$ very small, the Hessian can be written in terms of first-order derivatives, as in the Levenberg-Marquardt optimization procedure (Levenberg, 1944; Marquardt, 1963). Specifically, the Hessian for the MSE

loss (Eq. (3.5)) at $\delta\Phi = 0$ with respect to the log-parameters is

$$H^{MSE}(\Phi) = \Phi\Phi^\top \frac{1}{TS} \sum_{t \in T} \sum_{s \in S} J_{s,t}^\top(\Phi) \Sigma^{-1} J_{s,t}(\Phi), \quad (3.6)$$

with $K \times P$ Jacobian matrix $J_{s,t}(\Phi) = \left[\frac{d\mathbf{x}_{s,t}(\Phi)}{d\Phi_1}, \dots, \frac{d\mathbf{x}_{s,t}(\Phi)}{d\Phi_P} \right]$. This reduces the computational burden to $\mathcal{O}(SP)$ model evaluations to compute the Jacobian matrix $J_{s,t}(\Phi) = \frac{d\mathbf{x}_{s,t}(\Phi)}{d\log\Phi}$ (see Appendix A.1 for a full derivation). Indeed, this is a commonly used approximation in the literature on sloppy models.

Implementation & Estimation

To estimate the Jacobian matrix numerically, the baseline approach would be to rely on finite difference derivatives, which can be troublesome in sloppy models (Brown et al., 2004). Applying the central difference approach to compute $J_{s,t}(\Phi)$, with a step-size ε , requires $2SP$ runs of a model to estimate $H^{MSE}(\Phi)$ based on Eq. (3.6). A difficulty with this approach is that ABMs are generally path-dependent, meaning one also seeks S or T large to isolate the well- and ill-constrained directions independently of the specific realization of the noise in a simulation. Note that S and T large are not necessarily interchangeable as the models may be ergodic (one can apply the testing procedures of Vandin et al. (2022) to decide whether increases in S or T are warranted).

A recent implementation of the JUNE Epidemiological ABM by Quera-Bofarull et al. (2023a,b) showed that it is possible to implement Agent-based Models in a differentiable manner using computational tools developed in Machine Learning. This is a promising direction for model calibration that might alleviate the burden of computing finite differences because the model’s Jacobian is computed exactly by means of automatic differentiation, which means only one run would be necessary to compute $J_{s,t}(\Phi)$, reducing the computational burden by a factor of $2P$ at the expense of a large increase in memory requirements. The result is that despite the large number of entities and non-linear equations that make up an ABM, the infinitesimal Jacobian, and hence Hessian, is known. The main body of results in Chapters 5 and 6 make use of an implementation of the Mark-0 macroeconomic ABM using the PyTorch library. Interestingly, for the case of MABM, a comparison between finite differences and automatic differentiation has revealed that while one would expect for $\delta\Phi \rightarrow 0$ that the finite difference converge to the autodifferentiated result, the two approaches do not yield equivalent Jacobians (see Appendix C). While the subject of ongoing research, the underlying reason is the degree of path-dependence and noise in the model, which means that for finite difference Hessians to yield valid results requires a large $\delta\Phi$ to separate signal from the underlying noise inherent in the model. This effectively performs a form of coarse-graining on the model that is absent when using automatic differentiation (i.e. an infinitesimally small step size).

Once the numerical Hessian has been estimated, singular value decomposition is used to determine the eigenvalue, eigenvector pairs for that point in parameter space. Singular

Value Decomposition is more numerically stable, and thus preferable in the context of eigenvalues spanning many orders of magnitude. These eigenvalue-eigenvector pairs $(\lambda_i(\Phi), v_i(\Phi))$ reveal the linear log-parameter combinations and their associated stiffness.

Interpretation

Already at this point, one can make inference on the type of dynamic observed at the point Φ and the drivers of its various components. To do so, I consider the classification proposed by Francis and Transtrum (2019) based on the dynamics of $\lambda_i(\Phi)$ as a function of the amount of data T , under the assumption that the model that has a steady state (leaving S fixed for now). One can apply tests derived in Vandin et al. (2022) to ascertain this. Table 3.1 shows the different classifications of model behavior, which Francis and Transtrum (2019) motivate based on the following:

Eigenvalue behavior	Dynamics controlled by the associated eigenvector
$\mathcal{O}(T^{-1})$	Controls behavior of the transient
$\mathcal{O}(1)$	Controls the steady state properties (e.g. level)
$\mathcal{O}(T^2)$	Controls the frequency of an oscillatory regime
$\mathcal{O}(\exp(T))$	Controls the chaotic behavior of the timeseries

Table 3.1 – Classification of eigenvector relation to model dynamics through eigenvalue dynamics as a function of the total number of timesteps based on Francis and Transtrum (2019, Table 1)

1. As T grows, parameter combinations ϕ governing transient dynamics should decay:

$$\frac{d}{d\Phi_\phi} \mathbf{x}_{s,t}(\Phi) \sim 0 \quad \text{for } t \rightarrow \infty$$

2. In turn, parameter combinations governing steady state features should become constant:

$$\frac{d}{d\Phi_\phi} \mathbf{x}_{s,t}(\Phi) \sim \mathcal{O}(1) \quad \text{for } t \rightarrow \infty$$

3. For situations with fixed oscillations, the eigenvalues controlling frequency scale by $\mathcal{O}(T^2)$. To see this, rewrite the steady state as a Fourier series, and again compute the derivative

$$\frac{d}{d\Phi_\phi} \mathbf{x}_{s,t}(\Phi) = \sum_{k=-\infty}^{\infty} \frac{d\mathbf{x}_{s,t}(\Phi)}{d\alpha_k(\Phi)} \frac{d\alpha_k(\Phi)}{d\Phi_\phi} + \frac{d\mathbf{x}_{s,t}(\Phi)}{d\omega(\Phi)} \frac{d\omega(\Phi)}{d\Phi_\phi} \quad \text{for } t \rightarrow \infty,$$

where $\alpha_k(\Phi)$ and $\omega(\Phi)$ are the amplitude coefficients and oscillatory frequency. The first term, $\frac{d\mathbf{x}_{s,t}(\Phi)}{d\alpha_k(\Phi)}$ is bounded by a constant and $\frac{d\alpha_k(\Phi)}{d\Phi_\phi} \sim t$, implying $\mathcal{O}(T^2)$ as a behavior of the eigenvalue whose associated eigenvector controls the frequency of oscillation (see also Kramer et al., 1984; Larter et al., 1984; Wilkins et al., 2009).

Though one should here make a note of caution: this approach assumes that the behavior of the system can be described as a limit cycle, which may not be the case for a stochastic MABM that nonetheless oscillates.

4. Finally, for chaotic series Francis and Transtrum (2019) finds a scaling behavior where parameter controlling chaotic behavior scale exponentially in the length of the time series

$$\frac{d}{d\Phi_\phi} \mathbf{x}_{s,t}(\Phi) \sim e^{\lambda_\phi t} \quad \text{for } t \rightarrow \infty$$

Note that it is assumed here that $\mathbf{x}_{s,t}(\Phi)$ is a stationary series, as e.g. an exponential growth path would also lead to an exponential dependency if not converted into a stationary path of growth rates. Additionally, the exponential sensitivity will depend on the type of dynamic studied, for instance for not perfectly periodic oscillations, a small shift in frequency can lead to an exponential growth in the eigenvalue as a function of T (as shown in Chapter 5). Thus, exponential scaling is not necessarily due to a chaotic system.

3.3 Probabilistic Loss

Macroeconomic Agent-based Models (MABMs) are non-linear stochastic systems, which suggests comparing the predicted distributions instead of the predicted time series. That is, we use the model to produce a distribution $P(\mathbf{x}|\Phi)$ representing the probability that the model \mathcal{M} with parameters Φ generates observable \mathbf{x} . One example in the context of MABMs would be to compare distributions of agents, such as by income or firm size at a point in time t , or alternatively distributions of variables or their growth rates across time, such as the GDP growth distribution. The utility of using a probabilistic approach to the sloppy analysis is that it is less noise-dependent than an individual path $\mathbf{x}_{s,t}(\Phi)$ is, indeed one would expect the distributions across runs to be comparable. This alone may reduce the computational burden.

To compare the distributions between two parameter-sets Φ and $\Phi + \delta\Phi$, I apply the symmetrized Kullback-Leibler divergence (Kullback and Leibler, 1951) (also known as Jeffrey’s divergence (Jeffreys, 1948)), which is defined as

$$\mathcal{L}^{sKL}(\Phi, \Phi + \delta\Phi) = \frac{1}{2} (\mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi) + \mathcal{L}^{KL}(\Phi + \delta\Phi, \Phi)) \quad (3.7)$$

$$\mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi) = \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \log \left(\frac{P(\mathbf{x}|\Phi)}{P(\mathbf{x}|\Phi + \delta\Phi)} \right), \quad (3.8)$$

and is a global measure of distance.¹ Once again, we can derive the Hessian matrix (see Appendix A.2), which is proportional to the Fisher Information Metric (FIM) $g(\Phi)$, which “may be thought of as measuring distance in parameter space in units of standard

¹It is symmetric, $\mathcal{L}^{sKL}(\Phi, \Phi + \delta\Phi) = 0$ for $\delta\Phi = 0$, and can be calculated $\forall \delta\Phi$.

deviations, using the width of the distribution $P(\mathbf{x}|\Phi)$ on the space of possible data” (Quinn et al., 2023). Specifically, we can write

$$H_{i,j}^{SKL}(\Phi) = \Phi_i \Phi_j g_{i,j}(\Phi) d\Phi_i d\Phi_j \quad (3.9)$$

$$g_{i,j}(\Phi) = \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{\partial \log P(\mathbf{x}|\Phi)}{\partial \Phi_j} \frac{\partial \log P(\mathbf{x}|\Phi)}{\partial \Phi_i} \quad (3.10)$$

At this point, it is also worthwhile noting that several other loss-functions such as the Hellinger or Bhattacharyya distances also lead to the same Hessian as many of them fall under the category of f-divergences (Csiszár and Shields, 2004).

Unlike the MSE, there is no dynamics to the eigenvalues as a function of the number of simulated time steps. However, the FIM has other desirable properties. In terms of information geometry, it implies that one can measure the widths of the model manifold in different directions by means of a geodesic (a curve of least distance that connects two points). This provides a first intuition on the exploratory algorithm developed in Chapter 6.

In practical terms, estimating the FIM requires an estimate for the probability distribution $P(\mathbf{x}|\Phi)$, which is done here using Kernel Density Estimation (KDE) as a general purpose tool for getting a smoothed estimate $\hat{P}(\mathbf{x}|\Phi)$ of the probability density $P(\mathbf{x}|\Phi)$ based on the sample $\mathcal{X} = \{\mathbf{x}_{s,t}(\Phi) | \forall s \in S \forall t \in T\}$. To be precise, I apply a KDE with a Gaussian Kernel and $K \times K$ bandwidth matrix \mathbf{H} to estimate the density at K -dimensional point \mathbf{x} , which reads

$$\hat{P}(\mathbf{x}|\Phi) = \frac{1}{2^{|\mathcal{X}|}} \sum_{\mathbf{x}_i \in \mathcal{X}} (2\pi)^{-\frac{K}{2}} |\mathbf{H}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \mathbf{x}_i)^\top \mathbf{H}^{-1} (\mathbf{x} - \mathbf{x}_i) \right\} \quad (3.11)$$

The particular choice of Kernel is not crucial to the fit, allowing us to use the simple Gaussian form. The bandwidth choice \mathbf{H} on the other hand is critical (Epanechnikov, 1969). For univariate cases, I apply the Improved Sheather-Jones bandwidth estimate (Botev et al., 2010), while for multivariate cases of K dimensions I use $\mathbf{H} = \text{diag}(h_1^2, \dots, h_K^2)$ for simplicity, with h_i the Improved Sheather-Jones bandwidth estimate for the i -th variable. This allows for different variances for each variable, though not for cross-correlations as there is only re-scaling but no rotation.²

Within the models considered, it is likely that a given variable will be bounded, such as the unemployment rate being between 0 and 100%. This has implications for the estimate of the KDE, as it is in an unbounded space, it will lead to biased estimates at the boundaries. The simplest way to address this issue is to apply a transformation of variables (Wand et al., 1991). Specifically, following Koekemoer and Swanepoel (2008), I project the relevant variable $x_{k,s,t} \in \mathbf{x}_{s,t}(\Phi)$ onto \mathbb{R} using a simple logit transformation. This means applying $Y = g(X)$ such that $P(x_{k,s,t}|\Phi) = P(g(x_{k,s,t})|\Phi) g'(x_{k,s,t})$ with

²Wand (1994) suggests using at least $\mathbf{H} = \text{diag}(h_1^2, \dots, h_K^2)$ as $\mathbf{H} = h^2 I$ is too restrictive, and a full estimation might be expensive. In principle, one should of course allow for correlated variables.

$g(x) = \log \frac{x-a}{b-x}$ where we have boundaries $[a_k, b_k]$. In a one-dimensional setting (for simplicity), this implies

$$P(x_{k,s,t}|\Phi) = \frac{b_k - a_k}{(x_{k,s,t} - a_k)(b_k - x_{k,s,t})} P\left(\log\left(\frac{x_{k,s,t} - a_k}{b_k - x_{k,s,t}}\right) \mid \Phi\right), \quad (3.12)$$

The KL-divergence is dimensionless and invariant under parameter transformations,

$$KL(P(\mathbf{x}|\Phi) \parallel P(\mathbf{x}|\Phi^*)) = KL(P(g(\mathbf{x})|\Phi) \parallel P(g(\mathbf{x})|\Phi^*))$$

and thus also its symmetrized version. This invariance also implies that the Hessian matrix is equivalent to the non-transformed Hessian, making it a rather simple fix to the case of a bounded variable.

With these formulas and metrics in mind, in the next chapter I present an analysis of a simple two-dimensional model that highlights the power of such an approach. Pursuing this, the Mark-0 model is considered, it is a complete macroeconomic Agent-based model with a 19-dimensional parameter space.

Key Messages

- Stiff and sloppy directions are defined by the eigenvalue-eigenvector pairs $(\lambda_i(\Phi), v_i(\Phi))$ of the Hessian matrix of a given loss function $\mathcal{L}(\Phi, \Phi + \delta\Phi)$ comparing points Φ in parameter space.
- The largest eigenvalue corresponds to the stiffest direction, and the smallest to the sloppiest.
- The degree of sloppiness, \mathcal{S}_Φ , is defined by the ratio of smallest to largest eigenvalue.
- Using a mean-squared error (MSE) loss function allows us to characterize time-series dynamics by the dynamics of the eigenvalues as a function of the amount of time steps.
- Using a Jeffrey's divergence (SKL) loss function, one can compare distributions generated by the model, which depend less on the specific path of the stochastic model.

Chapter 4

Kirman's Ants: Understanding Sloppiness in Two Dimensions

There are two key questions surrounding sloppiness: first, can we identify critical parameter combinations using this parameter space decomposition approach? Second, if we can identify stiff and sloppy directions then how do we relate their degree of stiffness to the degree of change we see in the loss-function when perturbing the model's parameters in this direction. Chapter 3 introduced two loss-functions and how to compute their respective stiff and sloppy directions by means of decomposing the Hessian matrix. In this chapter, I build up some intuition for this approach by applying the sloppy models decomposition to Kirman's model of ant recruitment (Kirman, 1993). The model describes in the simplest terms how a group of interacting agents may collectively randomly switch ("herd") between two equivalent choices. Since its conception, the model has found multiple applications in explaining group behavior, from financial markets (Choi et al., 2022) to fishing (Moran et al., 2021).

Kirman's original model was developed using two parameters. In Section 4.2 I apply the parameter space analysis proposed in Chapter 2 and 3 by deriving the Hessian matrix with respect to the two parameters for both the Mean-Squared as well as the symmetrized KL approaches. Decomposing these matrices shows that the stiffest direction corresponds to the ratio of the parameters, which controls the steady state distribution of the proportion of ants at one food source, while the second combination has an expected zero eigenvalue as only the ratio of the two parameters is relevant. When perturbing the model in the stiffest direction I find that throughout the parameter space taking step sizes proportional to the inverse square root of the eigenvalue leads to equal changes in the loss-function. Using this, it is a simple exercise to *explore* the model by following this stiff direction to go smoothly from a bimodal to a unimodal distribution of Ants, that is, recovering all the possible *phases*, here given by the distribution of ants.

4.1 The Recruitment Model

The setup of the ant model is straightforward: Consider a set of N ants facing two equivalent infinite food sources A and B , where k is the number of ants at source A . Each period of the model corresponds to a draw of a random ant. This ant may switch its choice of food source with a small probability $\rho \in [0, 1]$. If this is not successful, another ant is drawn randomly and the first ant is “recruited” by the second with a probability $\mu \in [0, 1]$.

Define $x_t = \frac{k}{N}$ as the fraction of ants at food source A in period t of the model. Based on Kirman (1993) and Moran et al. (2020) we can write the dynamics of the model in the $N \rightarrow \infty$ case as

$$\dot{x} = \rho(1 - 2x) + \sqrt{2\mu x(1-x)}\mathcal{N}(0, 1) \quad (4.1)$$

which can be discretized (Ito) to

$$x_t - x_{t-\Delta t} = \rho(1 - 2x_{t-\Delta t}) + \sqrt{2\mu x_{t-\Delta t}(1 - x_{t-\Delta t})}\sqrt{\Delta t}\mathcal{N}(0, 1) \quad (4.2)$$

One key consideration here is the degree of noise in the system given by the selection of Δt ,¹ which determines the ability of the system to *visit* all possible outcomes $x \in [0, 1]$, but may also destroy the underlying dynamics of the system (i.e one recovers simply the noise distribution).

Continuing in the large N limit and for $\rho > 0$, Kirman (1993) and Moran et al. (2020) obtain that the normalized stationary distribution $P(x|\Phi)$ is a beta distribution:

$$P(\mathbf{x}|\Phi) = \frac{\Gamma(2\frac{\rho}{\mu})}{\Gamma^2(\frac{\rho}{\mu})} (x(1-x))^{\frac{\rho}{\mu}-1} \sim \text{Beta}\left(\frac{\rho}{\mu}, \frac{\rho}{\mu}\right) \quad (4.3)$$

which displays a critical point at $\rho = \mu$. For $\rho < \mu$ there is a unimodal stationary distribution with a peak at $x = \frac{1}{2}$, and for $\rho > \mu$ there is a bi-modal distribution with the density going to infinity at the boundaries. At the critical point $\frac{\rho}{\mu} = 1$, the density is uniform.

Figure 4.1 shows a realization for each of these cases together with their stationary distribution. We can see clearly the difference in persistence of one dominant location for the same random realization between the different parameter sets. Note here that in the case of the uniform distribution ($\text{Beta}(1, 1)$) it takes more data for the distribution to converge than does the bimodal or unimodal case.

4.2 Identifying Stiff and Sloppy Directions

To show that the sloppy models approach outlined in Chapter 3 correctly identifies the stiff and sloppy directions of this model, I derive the Hessian matrix of the two loss functions introduced: the mean-squared error (Eq. (3.5)) and the symmetric KL

¹This is referred to as the temperature of the system in physics parlance.

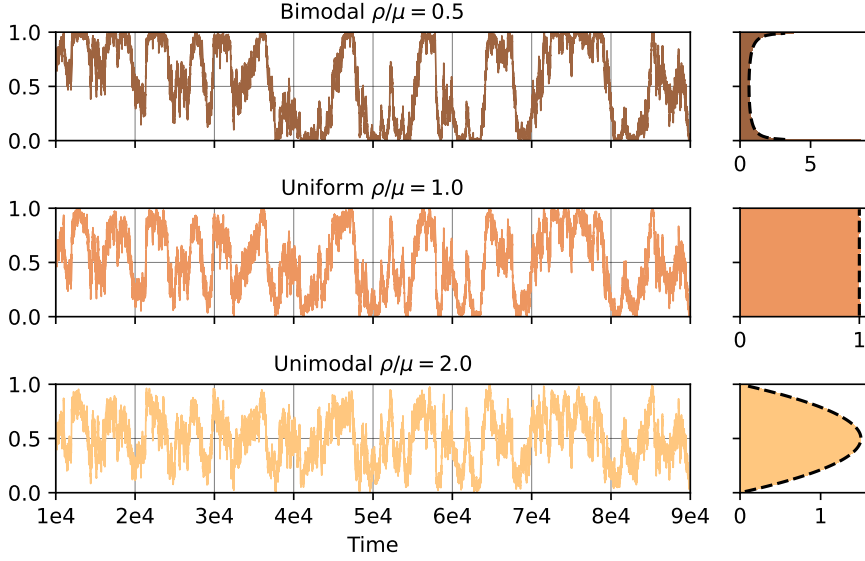


Figure 4.1 – Dynamics of Kirman’s ant model via Eq. (4.2) with $T = 10^7$, $\Delta t = 10^{-4}$, and $x_0 = \frac{1}{2}$. The left panel shows the dynamics for 40,000 time steps, and the right panels show the relevant stationary distributions. I keep $\mu = 1.0$ fixed across the three cases

divergence (Eq. (3.7)). For this model, the distribution of the density of Ants is the natural choice for the sKL-based loss function, while the mean-squared loss is slightly less trivial because irrespective of the point in parameter space, the mean of a symmetric beta distribution is $\frac{1}{2}$, such that the Mean-Squared Error between two realizations would always converge to zero for T and S large. Instead, I consider here the log variance, $\log V(x|\Phi)$, of the beta distribution as the observable for the Mean-Squared Loss, each element of the Hessian will thus refer to relative changes in the variance with respect to relative changes in the parameters. This also allows me to drop the time subscript t , as the variance at each point should be the same as the variance of the distribution as a whole.

The Symmetrized KL-Divergence Approach

Beginning with the symmetric KL based Hessian as the natural approach to this model, the continuous formulation of the ants model (recalling that $x \in [0, 1]$) converges to a beta distribution. Recalling from Eq. (3.9) that the Hessian matrix for the symmetric KL-divergence can be written as

$$H_{i,j}^{SKL}(\Phi) := \left. \frac{d^2 \mathcal{L}^{sKL}(\Phi, \Phi + \delta\Phi)}{d \log \Phi_i d \log \Phi_j} \right|_{\varepsilon_i=0} = \Phi_i \Phi_j \int_0^1 P(\mathbf{x}|\Phi) \frac{d \log P(\mathbf{x}|\Phi)}{d \Phi_i} \frac{d \log P(\mathbf{x}|\Phi)}{d \Phi_j} dx,$$

leads to the closed-form Hessian (see Appendix B.2 for the derivation):

$$H^{SKL}(\Phi) = \frac{\rho^2}{\mu^2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \left(2\psi^{(1)}\left(\frac{\rho}{\mu}\right) - 4\psi^{(1)}\left(2\frac{\rho}{\mu}\right) \right) \quad (4.4)$$

where $\psi^{(1)}(\cdot)$ is the trigamma function. Decomposing Eq. (4.4) yields the eigenvector-eigenvalue pairs

$$\begin{aligned} v_1(\Phi) &= \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix} & \lambda_1(\Phi) &= \frac{\rho^2}{\mu^2} \left(4\psi^{(1)}\left(\frac{\rho}{\mu}\right) - 8\psi^{(1)}\left(2\frac{\rho}{\mu}\right) \right) \\ v_2(\Phi) &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} & \lambda_2(\Phi) &= 0, \end{aligned}$$

where the factor $\frac{1}{\sqrt{2}}$ normalizes the eigenvectors to unit length.

As expected from the derived beta distribution (Eq. (4.3)), there is only one relevant direction: $\frac{\rho}{\mu}$ (eigenvector 1), whereas changes in ρ and μ that retain constant $\frac{\rho}{\mu}$ leave the stationary distribution untouched (eigenvector 2). In this case, the eigenvalue is already at zero by definition, implying that this model simplification would retain the exact dynamics.

The Mean-Squared Error

Turning now to the mean-squared error case, the Hessian (Eq. (3.6)) for the variance $\log V(x|\Phi)$ in the $S \rightarrow \infty$ case reads

$$H_{i,j}^{MSE}(\Phi) = \frac{\Phi_i \Phi_j}{V^2(x|\Phi)} \frac{dV(x|\Phi)}{d\Phi_i} \frac{dV(x|\Phi)}{d\Phi_j},$$

where I drop the t subscript as the variance at any point t is equivalent to the variance of the distribution for $S \rightarrow \infty$. Similarly, k is dropped as we have only one observable $K = 1$, which also eliminates any weighting ($w_k = 1$). Noting that for the symmetric Beta distribution, the variance is

$$V(x|\Phi) = \frac{1}{4} \left(2\frac{\rho}{\mu} + 1 \right)^{-1}$$

the Hessian can be solved for analytically as (see Appendix B.2):

$$H^{MSE}(\Phi) = 2\frac{\rho^2}{\mu^2} \left(2\frac{\rho}{\mu} + 1 \right)^{-2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \quad (4.5)$$

yielding an eigendecomposition of

$$\begin{aligned} v_1(\Phi) &= \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix} & \lambda_1(\Phi) &= \frac{\rho^2}{\mu^2} \left(2\frac{\rho}{\mu} + 1 \right)^{-2} \\ v_2(\Phi) &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} & \lambda_2(\Phi) &= 0, \end{aligned}$$

where we recover exactly the same singular relevant direction: $\frac{\rho}{\mu}$ (Eigenvector 1), while changes maintaining the ratio are irrelevant. In summary, we can note that in such a simplified case, both the sKL and MSE approach correctly identify the critical directions. The only difference in the decomposition of the sKL and MSE Hessian lies in the magnitude of the eigenvalues, which is due to the two different metrics and observables.

4.3 Steps into the stiffest direction

Given that the sloppy methodology identifies correctly the critical directions in the simple Ants model, I turn now to the second important aspect: the notion of distance in the parameter space as measured by the loss-function for two different parameter sets. In Chapter 2 it was shown by a second-order Taylor expansion that a step into the direction of an eigenvector, in this case v_1 , in proportion to the inverse square root of its eigenvalue, $1/\sqrt{\lambda}$, should lead to an equal loss irrespective of the specific point in the parameter space.

Considering now a spectrum of $\frac{\rho}{\mu}$ values, Figure 4.2 (left) shows that as $\frac{\rho}{\mu}$ increases, one observes a non-linear decrease in the first eigenvalue of $H^{SKL}(\Phi)$ as the difference in trigamma functions decreases more rapidly than the increase due to the square term. Indeed for $\frac{\rho}{\mu} \rightarrow 0^+$, $\lambda_1 \rightarrow 2$ and for $\frac{\rho}{\mu} \rightarrow \infty$, $\lambda_1 \rightarrow 1$. For the MSE-based hessian, one observes a saturation in the eigenvalue at $\frac{1}{4}$, and an exponential decrease when moving away from the peak. As the ratio $\frac{\rho}{\mu}$ increases and the eigenvalues decrease, one can see by studying the realizations of the symmetric KL Loss-function (Figure 4.2, right) that one needs a proportionally larger perturbation in the direction of v_1 to achieve a similar magnitude of change in the loss function. The reverse is true for the MSE, where smaller steps generate the same loss as $\frac{\rho}{\mu}$ increases.

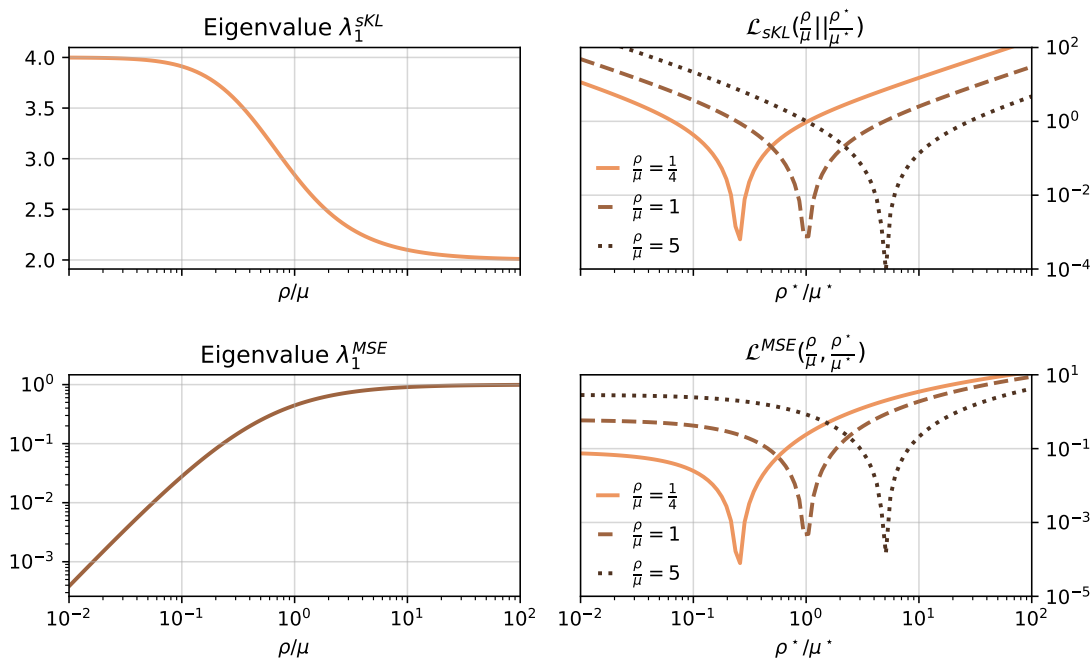


Figure 4.2 – Left: Eigenvalues of the sKL (top) and MSE (bottom) approach at different points in the parameter space. Right: The sKL (top) and MSE (loss) functions using different points of the parameter space as references.

To quantify the step size consider the symmetric KL divergence $KL_{sym}(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*))$ around Φ (see Appendix B.1), which is given by

$$\begin{aligned} KL_{sym}(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) &= KL(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) + KL(P(\mathbf{x}|\Phi^*) || P(\mathbf{x}|\Phi)) \\ &= 2 \left(\frac{\rho}{\mu} - \frac{\rho^*}{\mu^*} \right) \left(\psi \left(\frac{\rho}{\mu} \right) - \psi \left(\frac{2\rho}{\mu} \right) + \psi \left(\frac{\rho^*}{\mu^*} \right) - \psi \left(\frac{2\rho^*}{\mu^*} \right) \right). \end{aligned} \quad (4.6)$$

Given this form, we define $\frac{\rho^*}{\mu^*}$ as the result of perturbing the vector of log-parameters, $\log \Phi$, in the direction of v_1 with a step-size ε :

$$\log \Phi^* = \log \Phi + \varepsilon v_1 \quad \text{with} \quad \Phi = \begin{bmatrix} \rho \\ \mu \end{bmatrix}$$




Figure 4.3 – The correspondence between the magnitude of the loss-function for a perturbation ε in the first eigenvector, v_1 , direction and the magnitude of the eigenvalue. One can see the scaling law $\varepsilon \propto 1/\sqrt{\lambda_1}$ for a constant loss magnitude. Left: the symmetric KL-divergence, Right: the mean-squared error evaluated on the variance

Figure 4.3 (left) shows the magnitude of ε both in the positive and negative directions that would lead to a $KL_{sym}(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) = 10^{-3}$. As $\frac{\rho}{\mu}$ increases, and with it the eigenvalue decreases, the magnitude of ε necessary to achieve an equivalent KL_{sym} loss increases. Indeed, Figure 4.3 (right) shows that the increase can be described by a ratio

$$\varepsilon \propto \frac{1}{\sqrt{\lambda_1}},$$

as shown by the second-order decomposition in Chapter 3.

Key Messages

- Applying the Hessian decomposition to Kirman's Ants correctly identifies critical and non-critical parameter space directions for both an MSE and SKL loss function.
- In the case of Kirman's Ants it allows for a model reduction to one effective parameter.
- Taking steps in the directions of the strongest loss scales with the inverse square root of their respective eigenvalue (its degree of sloppiness), as predicted in Chapter 3.

Mark-0: Sloppiness in a Macroeconomic Agent-Based Model

Adapted from Naumann-Woleske et al. (2023) *Exploration of the Parameter Space in Macroeconomic Agent-based Models*, and significantly extended. The original paper was joint work with M. Knicker, J-P Bouchaud and M. Benzaquen. However, while this chapter is inspired by that work all results are *new* based on the differentiable Mark-0 model, which allows for a deeper analysis.

The underlying hypothesis for the first part of this thesis is that the parameter space of Macroeconomic Agent-based models (MABMs) displays a sloppy structure with a few well-constrained directions to which the model is sensitive, and many ill-constrained directions to which the model is insensitive. In this chapter, I show that one can indeed estimate the Hessian matrix of a MABM, and its eigendecomposition reveals the multi-decade hierarchical parameter-structure characteristic of sloppy models. To do this, I consider three distinct points in the Mark-0 model, each representing a different qualitative dynamic (phase). The estimation at these points reveals the characteristic structure of a sloppy model: there are a few well-constrained directions and many ill-constrained directions. The exact directions are phase-dependent with different dynamics for the associated eigenvalues, representing the degree of sensitivity over time.

Estimating the Hessian is not trivial. Using a C++ and PyTorch implementation of the Mark-0 model, this chapter, and its associated Appendix C, show that finite differences are particularly tricky as they need to deal with noise in estimation and perturbation, while Automatic Differentiation needs to deal with noise only in the perturbation. While finite differences can yield reasonable estimates, the results suggest that future models should strongly consider an implementation that allows for automatic differentiation, as

these results are less computationally intensive and more robust.

Finally, I validate the estimation of the Hessian matrices and understand their robustness with respect to the hyper-parameters T and S i.e. the amount and type of data. I find here that there is little to be gained by increasing the number of Monte Carlo repetitions, S , as with automatic differentiation for two of the phases the estimated eigenvectors are almost invariant to the choice of S . Meanwhile, for the remaining phase, one observes an exponential sensitivity to the most well-constrained direction due to small shifts in the frequency of oscillation, revealing a weakness in the choice of a non-linear least squares approach. In this case, the addition of more random realizations has little value as the exact direction and sensitivity of the most well-constrained parameter combination is dependent on the precise realization of the noise, which implies that a computationally prohibitive S would be necessary to ensure all Hessians agree on the direction and sensitivity. In the T dimension, once $T = 150$ time steps have occurred the first eigenvector-eigenvalue pair has converged to its final form, suggesting that when using automatic differentiation, computational and data requirements are relatively low.

5.1 The Mark-0 Model

The Mark-0 model is a stylized hybrid agent-based model that can generate a multiplicity of different qualitative phenomena it can generate.¹ Mark-0 was created by Gualdi et al. (2015) as a reduction of the Mark family of models (Gaffeo et al., 2008; Gatti et al., 2011) such that Mark-0 can still retain the dynamics of those models. It was later applied to questions of monetary policy (Gualdi et al., 2017; Bouchaud et al., 2018) and policy responses to the COVID-19 pandemic (Sharma et al., 2020). Most recently, we used the model to assess the post-COVID period of high inflation (see Chapter 9 and Knicker et al., 2023). While the aggregate behavior that one can generate with the Mark-0 model is likely not quantitatively precise, it generates a variety of distinct, plausible and generic macroeconomic behaviors. In most papers on Mark-0 the authors have endeavored to provide phase diagrams, showing how the different qualitative behaviors of the model depend on the choice of parameterization. This makes the Mark-0 model the perfect first candidate for a sloppy-based analysis to both use and verify the parameter-sensitivities developed manually by the original authors.

Model Overview

Mark-0 consists of an agentized firm-sector of N firms producing a composite consumption good purchased by an aggregated household. In addition, there is a central bank that sets the baseline interest rate, affecting both the firms' and household's decisions. In the remainder of this section, I briefly outline the model's main equations following the formalism and notation of Sharma et al. (2020) where I use the notation Φ_i to refer

¹Hybrid refers here to the characteristic that only one set of agents (firms in the Mark-0 model) are represented by a large number of agents, while other groups (e.g. households) are modeled as an aggregate dynamic.

to parameters, where i is the original authors' choice of notation. My purpose is not to be mathematically complete but rather provide an intuition on the model's mechanics that aids the later analysis of the well-constrained and sloppy parameter-directions as well as the phase diagrams explored. For a full detailed description, refer to Sharma et al. (2020) and the pseudocodes of Bouchaud et al. (2018).

Firms: There are N firms producing $Y_i(t)$ units of a perishable composite consumption good using a linear production function of only labor $L_i(t)$ at a constant unit productivity. The firms sell their goods to the household at a price $p_i(t)$, and pay a wage rate $W_i(t)$. Each period, a firm updates its production based on the gap between prior production and demand $D_i(t)$ of the household in a heuristic attempt to optimize its income

$$Y_i(t+1) = \begin{cases} Y_i(t) + \min \{ \eta_i^+(t) (D_i(t) - Y_i(t)), u_i^*(t) \} & \text{if } Y_i(t) < D_i(t) \\ Y_i(t) - \eta_i^-(t) (D_i(t) - Y_i(t)) & \text{if } Y_i(t) > D_i(t) \\ Y_i(t) & \text{otherwise} \end{cases} \quad (5.1)$$

where $u_i^*(t)$ is the wage-dependent pool of unemployed workers firm i may access to increase production, representing a maximum output given the unit-productivity (see Sharma et al., 2020). The variable $\eta_i^\pm(t)$ is the speed at which firms hire and fire labor, and depends on a firm's financial fragility, which is represented by its debt-to-payroll ratio $\phi_i(t)$. It is important to note here that a firm will go bankrupt if its fragility rises above a threshold Φ_θ (the full bankruptcy mechanism can be found in Gualdi et al., 2017). The idea is that more fragile firms fire faster and hire slower, and vice versa. Consequently, the hiring and firing rates for firm i can be described by

$$\eta_i^+ = \llbracket \Phi_{\eta_0^+} (1 - \Gamma(t)\phi(t)) \rrbracket \quad \eta_i^- = \llbracket \Phi_{\eta_0^-} (1 + \Gamma(t)\phi(t)) \rrbracket, \quad (5.2)$$

with $\llbracket x \rrbracket$ clipping x to be bound by $[0, 1]$. Here $\Gamma(t)$ represents the sensitivity of firms to their fragility,

$$\Gamma(t) = \max \{ \Phi_{\alpha_\Gamma} (\rho_l(t) - \hat{\pi}(t)), \Phi_{\Gamma_0} \}, \quad (5.3)$$

which is positive in the real loan-interest rate $\rho_l(t) - \hat{\pi}(t)$ with a minimum of Φ_{Γ_0} . The expected inflation, $\hat{\pi}(t)$, is formed by a weighted sum of $\Phi_{\tau_{ema}}$ times an exponentially weighted moving average of the price-level inflation with memory parameter Φ_{τ_π} , and $\Phi_{\tau_{tar}}$ times the target inflation rate Φ_{π^*} of the central bank. Here the underlying price series is a production-weighted price index, $\bar{p}(t)$, computed across all alive firms.

Firms update their relative prices in a multiplicative process taking into account their excess demand and expected inflation:

$$p_i(t+1) = \begin{cases} p_i(t)(1 + \Phi_{\gamma_p} \xi_i(t))(1 + \hat{\pi}(t)) & \text{if } Y_i(t) < D_i(t) \cup p_i(t) < \bar{p}(t) \\ p_i(t) & \text{if } Y_i(t) < D_i(t) \cup p_i(t) \geq \bar{p}(t) \\ p_i(t)(1 - \Phi_{\gamma_p} \xi_i(t))(1 + \hat{\pi}(t)) & \text{if } Y_i(t) > D_i(t) \cup p_i(t) > \bar{p}(t) \\ p_i(t) & \text{if } Y_i(t) > D_i(t) \cup p_i(t) \leq \bar{p}(t) \end{cases} \quad (5.4)$$

where $\xi_i(t) \sim U[0, 1]$ is uniformly distributed and Φ_{γ_p} regulates the size of the price-change.

In a similar manner, firms update their wage offering as

$$W_i(t+1) = \begin{cases} W_i(t) [1 + \Phi_{\gamma_w} \Upsilon^-(1 - u(t)) \xi'_i(t)] [1 + \Phi_{g_w} \hat{\pi}(t)] & \text{if } Y_i(t) < D_i(t) \cup \Pi_i(t) > 0 \\ W_i(t) [1 - \Phi_{\gamma_w} \Upsilon^+ u(t) \xi'_i(t)] [1 + \Phi_{g_w} \hat{\pi}(t)] & \text{if } Y_i(t) > D_i(t) \cup \Pi_i(t) < 0 \end{cases} \quad (5.5)$$

where $\Upsilon^\pm = (1 \pm \Gamma(t)\phi_i(t))$, $\Pi_i(t)$ is the firm's profit, $u(t)$ the current unemployment rate, and $\xi'_i(t) \sim U[0, 1]$. In the case that a positive wage adjustment implies that a firm would have negative profits ceteris paribus, wages are adjusted upward only to the point where $\Pi_i(t) = 0$.

The firms profits are then computed by

$$\Pi_i(t) = p_i(t) \min \{Y_i(t), D_i(t)\} - W_i(t)Y_i(t) + \rho^d(t)\mathcal{E}_i^+(t) - \rho_i^l(t)\mathcal{E}_i^-(t) \quad (5.6)$$

where $\mathcal{E}_i^+(t) = \max \{\mathcal{E}_i(t), 0\}$ and $\mathcal{E}_i^-(t) = \min \{\mathcal{E}_i(t), 0\}$ are a firm's debt (negative assets) and its cash assets, respectively. Here, firms with positive profits and positive cash balance then pay out a dividend on their cash at a rate Φ_δ . Consequently, the evolution of a firm's assets is simply $\mathcal{E}_i(t+1) = \mathcal{E}_i(t) + \Pi_i(t) - \Phi_\delta \mathcal{E}_i^+(t) I(\Pi_i(t) > 0)$ with $I(x) = 1$ if the condition x is met and zero otherwise.

Households: The model features a aggregate household with a consumption budget C_B comprising a fraction $c(t)$ (the consumption propensity) of the households savings $S(t)$, wages $W(t)$, and interest on deposits $\rho_d(t)S(t)$. The consumption propensity $c(t) \in [0, 1]$ depends on a baseline propensity Φ_{c_0} adjusted for real interest rates by Φ_{α_c} :

$$c(t) = \Phi_{c_0} (1 - \Phi_{\alpha_c} (\rho_d(t) - \hat{\pi}(t))). \quad (5.7)$$

Given a budget, the household distributes its demand between all firms i by an intensity of choice function

$$D_i(t) = \frac{C_B(t)}{p_i(t)} \frac{\exp(-\Phi_\beta p_i(t))}{\sum_j \exp(-\Phi_\beta p_j(t))}, \quad (5.8)$$

leading to a realized consumption of

$$C(t) = \sum_i^N p_i \min(Y_i, D_i) \leq C_B(t) \quad (5.9)$$

Finally, the household saving stock evolves based on the sum of flows

$$S(t+1) = S(t) + W(t) + \rho^d(t)S(t) - C(t) + \Delta(t), \quad (5.10)$$

with dividend income $\Delta(t)$ as a fraction Φ_δ of firms' assets (if positive).

Interest Rates: There is a Central bank regulating the interest rates faced by the firms and households. The Central Bank sets the baseline interest rate based on a Taylor-like rule to control inflation (Gualdi et al., 2017, also consider an unemployment rate term)

$$\rho_0(t) = \Phi_{\rho^*} + \Phi_{\phi_\pi} (\hat{\pi}(t) - \Phi_{\pi^*}), \quad (5.11)$$

that depends on an exogenous baseline Φ_{ρ^*} and an inflation target Φ_{π^*} . This rate is then passed on to the deposit and loan rates, with

$$\rho_l(t) = \rho_0(t) + \Phi_f \frac{\mathcal{D}(t)}{\sum_i \mathcal{E}_i^+(t)} \quad (5.12)$$

$$\rho_d(t) = \frac{\rho_0(t) \sum_i \mathcal{E}_i^-(t) - (1 - \Phi_f) \mathcal{D}(t)}{S(t) + \sum_i \mathcal{E}_i^+(t)} \quad (5.13)$$

where $\mathcal{D} = \sum_{j \in \{1, \dots, N\}, \phi_j(t) > \Phi_\theta} \mathcal{E}_j^-(t)$ is the total debt of all defaulting firms, and Φ_f interpolates between the costs of default falling on borrowers ($\Phi_f = 1$) or depositors ($\Phi_f = 0$). To close the model accounting, the total amount of money in circulation, $M = S(t) + \sum_i \mathcal{E}_i(t)$, is kept constant.

Parameterization & Implementation

In total, the model features 19 free parameters that I investigate. The parameters are presented in Table 5.1, which shows the baseline values of parameters based on Knicker et al. (2023), Sharma et al. (2020) and Bouchaud et al. (2018). The table also includes parameter boundaries, which were set either by natural limitations (such as $[0, 1]$ for Φ_f) or by considering a variation of roughly one order of magnitude around the baseline parameters used in previous applications. In addition to the parameters in Table 5.1, the Mark-0 model is subject to several hyper-parameters. The first is the dimension of a step t , which is here taken to be a month. Additionally, for this thesis, I consider the case of $N = 3000$ firms across $T = 300$ periods, which balances size with computational expense, and represents the relevant timescale for policy-analysis (e.g. see analysis in Knicker et al. (2023) or Lamperti et al. (2019a) as examples).

For the purposes of our analysis here, I consider three points in the model's parameter space, each with a qualitatively different behavior hereafter called a *phase*. These points differ only in the choice of parameters Φ_θ and Φ_{ρ^*} , with the remaining parameters as in the baseline shown in Table 5.1. For the purposes of this thesis, I consider only the unemployment rate as a target variable as it has different phases and is representative of the model's overall dynamics. Figure 5.1 shows the dynamics of the unemployment rate for a single realization of the model for each of the three parameter choices. As highlighted by the sample time series, the phases can be classified into four categories: Full Employment (FE, average unemployment below 5%), Full Unemployment (FU, average unemployment above 80%, which is not shown here), Endogenous Crisis (EC, periodically oscillating unemployment mostly for a standard deviation of larger than 10%), and Residual Unemployment (RU, otherwise) by considering their mean and standard

Symbol	Description	Baseline	Lower Bound	Upper Bound
α_c	Sensitivity of consumption propensity to real deposit rates	4.0	0.0	40.0
α_Γ	Sensitivity of firms to real loan rates	50.0	0.0	500.0
β	Intensity of choice for households	2.0	0.0	20.0
c_0	Baseline propensity to consume	0.5	0.001	0.999
δ	Dividend payout rate	0.02	0.0	0.2
η_0^-	Baseline firing propensity	0.1	0.01	1.0
f	Share of bankruptcy burden allocated to firms	0.5	0.0	1.0
Γ_0	Sensitivity of firms to their fragility	0.0	0.0	5.0
γ_p	Firm price adjustment magnitude	0.05	0.005	0.5
ϕ_π	Central-bank reaction to inflation deviations	0.0	0.0	5.0
ϕ	Firm Revival rate	0.1	0.01	1.0
π^*	Inflation Target	0.002	0.0	0.02
r	Ratio of wage to price-adjustment i.e. $\gamma_w = r\gamma_p$	1.0	0.1	10.0
R	Ratio of hiring to firing rate i.e. $\eta_0^+ = R\eta_0^-$	2.0	0.1	10.0
ρ^*	Baseline Central Bank interest rate	0.001	0.0	0.05
τ_π	Exponential Moving Average (EMA) Parameter	0.2	0.01	0.99
τ_{meas}	EMA weight in inflation expectation	0.5	0.0	1.0
τ_{tar}	Central Bank target weight in inflation expectation	0.5	0.0	1.0
θ	Bankruptcy Threshold	4.0	0.0	10.0
g_w	Wage adjustment to inflation	1.0	0.1	10.0

Table 5.1 – Parameterisation of the Mark-0 model. There is no clear *baseline* parameterisation, so the most common parameters were taken from Knicker et al. (2023), Sharma et al. (2020), and Bouchaud et al. (2018). The important part in the table are the bounds, which either follow natural parameter boundaries, or are based around a factor 10 around the baseline value.

deviation as in Gualdi et al. (2015). While Figure 5.1 shows a timescale of ~ 200 periods to converge to a steady state for these three points in parameter space, hereafter the model is initialized to the steady state computed after $T = 1000$ time steps.




Figure 5.1 – Unemployment rate dynamics at three sample points of the Mark-0 model: (FE) a full-employment situation, (RU) a residual unemployment point, and (EC) an endogenous crisis point. All other parameters are taken from Table 5.1

5.2 Separating Signal and Noise: Estimating the MSE-Hessian

Agent-based Models are stochastic generators that generally have a strong path dependence. This is an appealing feature and selling point of ABMs, as it mimics the path-dependence of the real complex economy, however, it also makes estimation with these models difficult as time series can be very noise-dependent. In the case of this thesis, the challenge lies in using numerical finite difference approaches to estimate a MSE-based loss on the time series of the unemployment rate. Intuitively, the estimation and decomposition of the Hessian is a means to approximate the equi-loss surface around point Φ as an ellipsoid whose axes are the eigenvectors and diameters along those axes are proportional to $1/\lambda_i(\Phi)$ (Gutenkunst et al., 2007). The finite difference approach applies minuscule perturbations to parameters to ascertain the first derivative of the unemployment rate at a given time t with respect to the parameter in question. By a small perturbation, I refer here to “standard” perturbation sizes used in finite difference approaches such as 10^{-9} . Generally, when the landscape of the model manifold is smooth, as for instance with an exponential decay, taking small perturbations to parameters is sufficient for an unbiased and low-variance estimate of the derivative. However, for a stochastic ABM like Mark-0, the effect on the time series realization of a given observable from a small perturbation in parameters is almost indistinguishable from the effect of noise: there is a low quality of fit for the ellipsoid on the rough manifold. The reason for this is due to thresholds and if-else directives in the model that, even when noise is fixed, shift the exact timing of different events irrespective of the direction in parameter space that is taken. As shown in Appendix C, estimating the Mark-0 Hessian in three differ-

ent phases using such a small perturbation leads to a set of eigenvector-eigenvalue pairs where the eigenvalues are grouped within a single order of magnitude and perturbations into the directions of the eigenvectors don't yield an expected hierarchy of differences (as suggested by Eq. (3.2)), with worse performance as the degree of variability in the unemployment rate grows.

This leaves three different paths to extracting a meaningful eigenvector-eigenvalue structure when using an MSE approach on the direct time series realization of the path-dependent model. The first, which we applied in Naumann-Woleske et al. (2023) using a C++ implementation of Mark-0, is to use an extremely large data sample.² Specifically, we used $T = 20,000$ time steps across $S = 20$ random realizations, from which an estimation of the first eigenvector-eigenvalue pairs led to a change in observable dynamics larger than that of a change in random realizations or a random direction. More data can slowly pull apart the eigenvalue spectrum, separating eigenvalues across multiple decades sequentially. However, this approach has several drawbacks: most notably, tens of thousands of time steps is multiple orders of magnitude above a policy horizon of interest, and for most models running such a large set of simulations is computationally prohibitive. It also may not effectively resolve the signal-to-noise issue beyond identifying the first eigenvector, as one can observe that eigenvalues remain closely grouped (see Appendix C). A second approach within the numerical finite difference method is to increase the size of the perturbation, thus applying a sort of coarse graining of the model manifold surrounding the point of interest. The effect is to boost the signal from parameter perturbations at the expense of the extracted information perhaps not being an exact reflection of the true local Hessian (in the infinitesimal sense). Nonetheless, this allows us to retain a reasonable timescale ($T = 300$ periods) and a small number of random realizations. As shown in Appendix C, as one increases the finite difference perturbation ϵ to the order of 10^{-2} , the perturbations of the models in the computed eigenvector directions display the characteristic hierarchy postulated in the Taylor expansion of Eq. (3.2). That is, sequentially perturbing the model into directions associated with lower eigenvalues leads to a decreased change in the MSE. However, it appears that even with such a large perturbation only the first handful of eigenvalues have separated from the bulk, suggesting that an infeasibly large perturbation combined with more data would be needed to extract the true underlying eigenvalue distribution.

The third approach, and the one used for the computations in this and the next chapter, is to apply automatic differentiation. Inspired by the differentiable ABM approaches used in Chopra et al. (2022) and Quera-Bofarull et al. (2023a,b), Mark-0 has been ported from C++ to a PyTorch implementation which can make use of the built-in Autograd package (Maclaurin, 2016) to compute the exact Jacobian matrix of a given realization of the Mark-0 model. Automatic differentiation has several advantages over finite differences, in particular, it is both numerically stable and efficient, taking roughly six minutes for the $T = 300$, $S = 1$ case. As Mark-0 is executed, the Autograd package transforms

²A python interface to the code of Sharma et al. (2020) can be found at github.com/KarlNaumann/Mark0.

the Mark-0 model into a (large) sequence of primitive operations captured in a computational graph. Each of the captured primitive operations has an assigned derivative function, thus allowing for the computation of the exact Jacobian by propagating backward through the computational graph from the final leaf nodes.³ The downside of this approach remains its computational expense, not in terms of runtime but in terms of the required memory (either RAM or GPU based), which scales in the number of time steps T and the complexity of operations, though it makes large N possible. For instance, using sparse-tensor calculus Chopra et al. (2022) execute a simulation of “800,000 agents over 133 time steps in 4 seconds on a GPU (and 60 on a CPU)”. This approach leads to the hypothesized multi-decade eigenvalue structure with a hierarchy of perturbation effects, which will be analyzed in the following section.

At this point it is worth pointing out that one could also take a different approach by changing the observable in question to be smoother, such as focusing on the moments of the distribution, such as the mean or standard deviation, or applying a smoothing function such as a moving average. However, this leaves out information, thus obscuring interesting time series dynamics. For example, considering the mean unemployment, the measure may be equivalent for a oscillating and non-oscillating case, or requires some prior knowledge of the model to be effective. Instead, my ambition here is to remain general and agnostic, simply observing directly the time series outcomes without modification.

5.3 Disentangling the Parameter Hierarchy of Mark-0

The simple answer to the question of whether the Mark-0 model’s dynamics depend on only a few well-constrained parameter directions, is yes: the model displays the characteristic parameter hierarchy introduced in Chapters 2 and 3. This section first presents the sloppy structure of Mark-0, as given by automatic differentiation, and then proceeds to test its robustness with respect to the two hyper-parameters T and S . Finally, I confirm the property that perturbations in eigenvector directions lead to a hierarchy in the degree of loss (i.e. the best-constrained direction has the highest loss) as predicted in Eq. (3.1).

Figure 5.2 gives an overview of the eigenvalue spectra, unemployment dynamics and the first nine eigenvectors of the Hessian matrix at three points for the large data case of $S = 50$ random realizations. The interpretation is clear: the eigenvalues of the three different phases span over seven decades in a roughly uniform distribution (Left Panel). In all cases, the first eigenvalue is separated from the remainder by at least one order of magnitude. Turning to the eigenvectors, for each set of dynamics we observe different parameter combinations making up the well-constrained directions, suggesting that the dynamics of each phase are each governed by a different key mechanism (Fig. 5.2 heatmaps). For each of the three chosen points, the first nine eigenvectors represent

³Of course, not all operations truly have a derivative, but one can transform operations such as if-else into differentiable terms by using a *tanh* function, and maxima using an exponential.

combinations of parameters, with only $v_1(\Phi)$ of the RU phase presenting a direction aligned with a bare parameter axis (that of Φ_{c_0}).

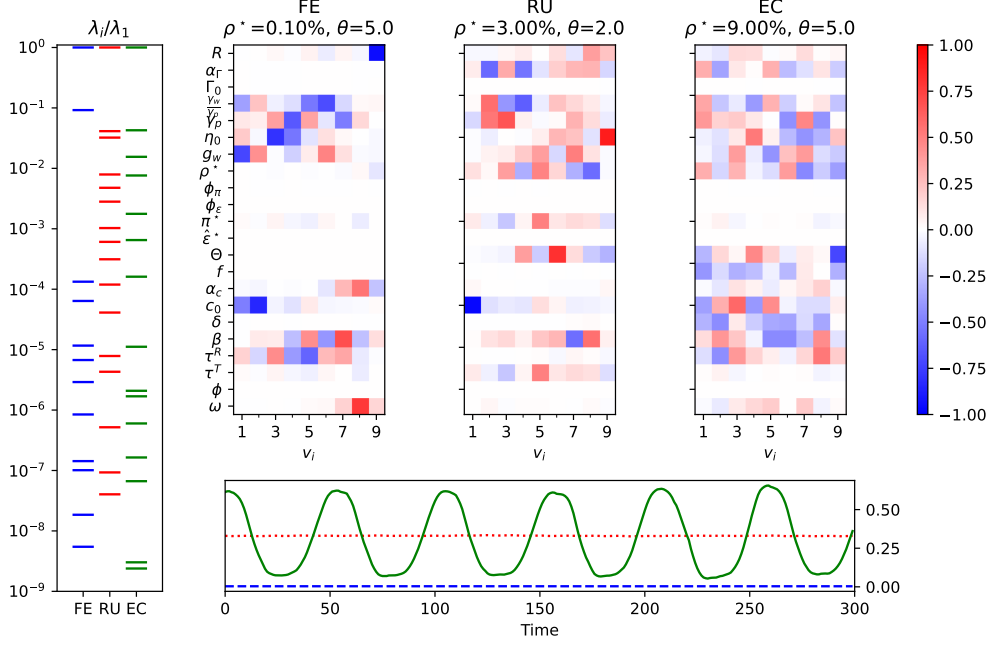


Figure 5.2 – Overview of the Hessian properties of three points in the Mark-0 model. Left: Eigenvalue spectrum $\lambda_i(\Phi)/\lambda_1(\Phi)$ for each point. Bottom: Sample dynamics of the unemployment rate for each point, representing full employment (blue dashed, $\Phi_{\rho^*} = 0.1\%$, $\Phi_{\theta} = 5.0$), residual unemployment (red dotted, $\Phi_{\rho^*} = 3\%$, $\Phi_{\theta} = 2.0$), and endogenous crises (green solid, $\Phi_{\rho^*} = 9\%$, $\Phi_{\theta} = 5.0$). The remaining parameters can be found in Table 5.1. Right: First nine eigenvectors (most well-constrained directions) of the Hessian matrices at each point. Hessians were computed with $S = 50$ and $T = 300$

The Effects of S and T on $\lambda_i(\Phi)$ and $v_i(\Phi)$

Automatic differentiation, as introduced in Section 5.2, is numerically stable and exact in terms of computing the Jacobian. For each execution of the Mark-0 model, the extracted Jacobian and thus Hessian should reflect the same properties of the underlying data generating process, with only the noise terms changing from execution to execution. It thus stands to reason that both the eigenvalues and eigenvectors should be similar across random realizations, with differences between them primarily driven by the finite sample size (here $T = 300$) and in proportion to the degree of variability in the time series as driven by the noise.

Beginning with a study of the eigenvalues $\lambda_i(\Phi)$, for both the FE and RU phase, the dynamics of $\lambda_i(\Phi)$ as a function of T shown in Figure 5.3 suggest that the first eigenvalue of each seed controls the steady state level of the unemployment rate based on the classification of Francis and Transtrum (2019, see also Section 3.2). Both the level and




Figure 5.3 – (Top) Dynamics of the largest 5 eigenvalues of three selected points as a function of the number of time steps T in the model. The solid line represents the median realization, with shaded areas presenting the 95% quantiles over 50 $S = 1$ realizations. (Bottom) Effect of the number of random realizations S on $\lambda_i(\Phi)$ at the final time step $T = 300$. Shaded areas represent the limits of the 95% quantiles for 50 samples, while solid lines represent the median eigenvalue.

This uses the PyTorch implementation with parameters from Table 5.1

dynamic of $\lambda_1(\Phi)$ for these phases is invariant to the number of random realizations over which the Hessian is computed, suggesting that for these phases a $S = 1$ Hessian computation using automatic differentiation is actually sufficient to extract the sensitivities of the underlying data generating process. In turn, the EC phase displays an exponential sensitivity to the first several eigenvector directions. The EC time series in Figure 5.2 reveals the strong periodic oscillations of this phase, with a sample size of $T = 300$ covering slightly more than five cycles. The structure of the MSE as a loss function means that any small shift in the frequency of this oscillation can lead to large changes in the loss, even if the qualitative dynamics are equivalent. Furthermore, such a small shift will increase the loss exponentially over time as the oscillations become more out of phase with one another. Note that here the loss-function is exponentially sensitive, but the underlying model is not chaotic as hypothesized by Francis and Transtrum (2019). Figure 5.3 shows that as one increases S , the spectrum of the eigenvalues converges to what is shown in Figure 5.2, where the eigenvalues are roughly uniformly distributed in log-space, as would be the case for a sloppy model. In this case, ill-constrained parameter combinations take sequentially longer to converge to a steady value.

Turning to the study of the eigenvectors associated with the respective eigenvalues, Figure 5.4 shows the realization of the first eigenvector for each of 50 different random realizations (top) as well as the pairwise cosine similarity (bottom) of the first five eigenvectors across these realizations. One can clearly see that the Hessian structure of the full employment phase is invariant to changes in the noise as there is little variation in the time series itself, while as the degree of variation increases across the residual

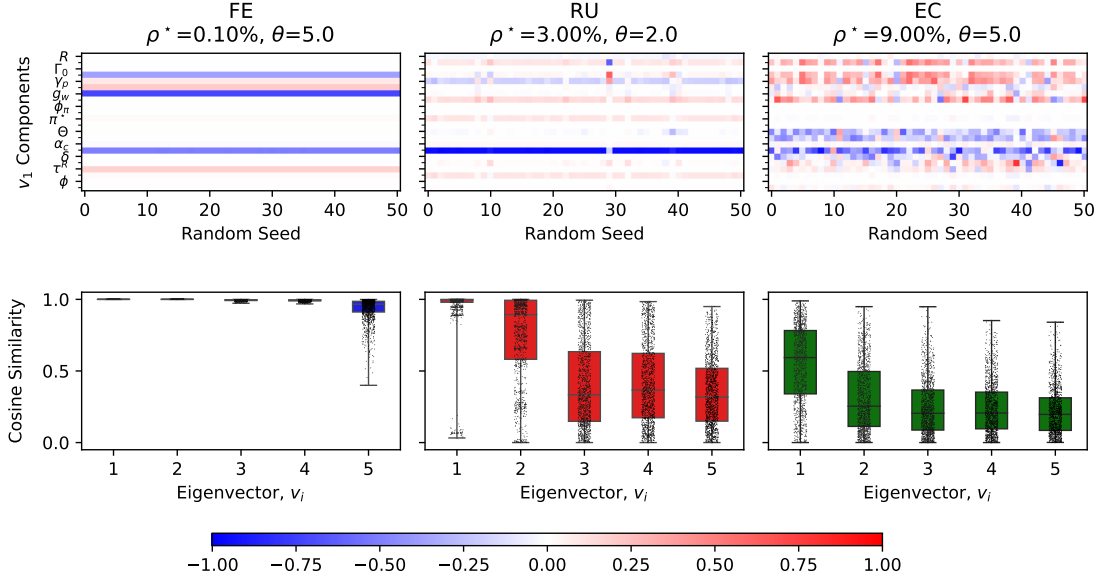


Figure 5.4 – (Top) Eigenvector $v_1(\Phi)$ for 50 different $S = 1$ realizations of Mark-0. (Bottom) Pairwise absolute cosine similarities between random realizations for the first 5 eigenvectors. Box-plot whiskers are the minimum and maximum, while boxes represent the first to third quantile of the data, finally the line represents the median. Parameters are from Table 5.1, with $T = 300$

unemployment and endogenous crises phases finite sample properties kick in. Incidentally, this structure of noise effect on the eigenvectors and eigenvalues is similar in the finite difference approximation (see Appendix C), but significantly less pronounced. The degree of variation in the eigenvectors also matches the variation in the eigenvalues presented in Figure 5.3: for RU only the first direction has a very tight confidence interval, with increasing uncertainty as directions becomes ill-constrained. The same holds for the EC phase, except here also the first eigenvector-eigenvalue pair displays a higher uncertainty between individual realizations of the model due to the structure of the chosen loss-function.

These observations are confirmed by Figure 5.5, which reveals that the first eigenvectors of the FE and RU phase are equivalent irrespective of the number of seeds S , and have converged to a fixed direction after roughly $T = 150$ time steps of the model realization. This too is an important observation as it means that, even if the policy horizon of the model is above $T = 150$ periods one does not need to consider those time steps when estimating the most sensitive parameter combination. The extra data affects only the value of $\lambda_i(\Phi)$, as shown in Figure 5.3, though even there $\lambda_1(\Phi)$ is constant for FE and RU for $T > 150$. For the EC and RU phases, Figure 5.5 suggests that even when aggregating across multiple S , ill-constrained directions may not be in agreement. This, again, is likely due to the higher degree of noise in these particular phases as compared to the FE phase, reflecting the weakness of the MSE as a loss function. In fact, the exponential sensitivity of the EC phase also implies that even for $S = 50$ seeds there

is no perfect agreement on the most well-constrained direction, which for FE and RU occurs almost immediately.

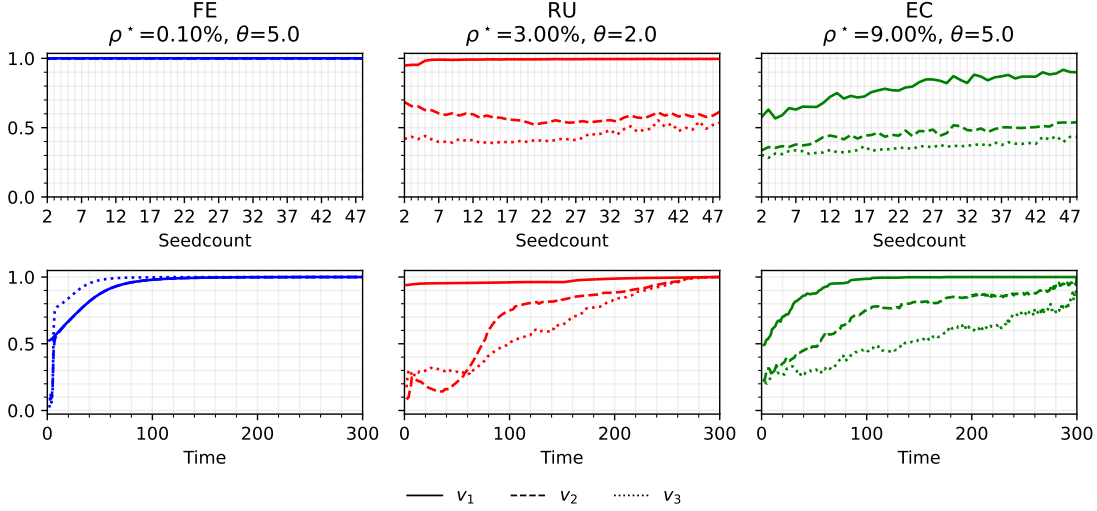


Figure 5.5 – (Top) Mean pairwise cosine similarity of the three most well-constrained eigenvectors at different seed counts. (Bottom) Mean pairwise absolute cosine similarities over time for 50 $S = 1$ Hessians. Parameters are from Table 5.1, with $T = 300$

At this point, I draw two conclusions: the first is that when using automatic differentiation $S = 1$ is a sufficient approximation for fixing $v_1(\Phi)$. In both the FE and RU cases, the $\lambda_1(\Phi) - v_1(\Phi)$ combination is the same across different seeds. Meanwhile, for the exponentially sensitive EC phase it appears not much information is gained from additional S , as the convergence is extremely slow, while the directions are roughly the same, depending on the exact noise in the phase. The second conclusion is that for a given $S = 1$ Hessian, the first eigenvalue-eigenvector combination is fixed already at $T = 150$, implying that if one wanted to compute many Hessians taking $T = 150$ would also be a sufficient approximation to the $T = 300$ Hessian.

Eigenvector Perturbations

To verify that these parameters represent more than random directions, we can study the effect on the MSE of a perturbation of the parameters into each of the eigenvector directions. Figure 5.6 shows the mean-squared error of a fixed perturbation $\varepsilon = 10^{-2}$ in each eigendirection $v_i(\Phi)$ (bottom), as well as examples of the change in the sampled time series when stepping into the most well-constrained and ill-constrained direction (top). The hypothesis developed in Chapter 3 is that well-constrained directions should have a larger effect than ill-constrained directions in proportion to their relative eigenvalues (Eq. (3.2)). One can see here that the hierarchical effects aren't exact, but that they do exist at least for the first eigenvector compared to the rest. The later convergence to similar MSEs in the different phases is due to the stochastic nature of the model, as

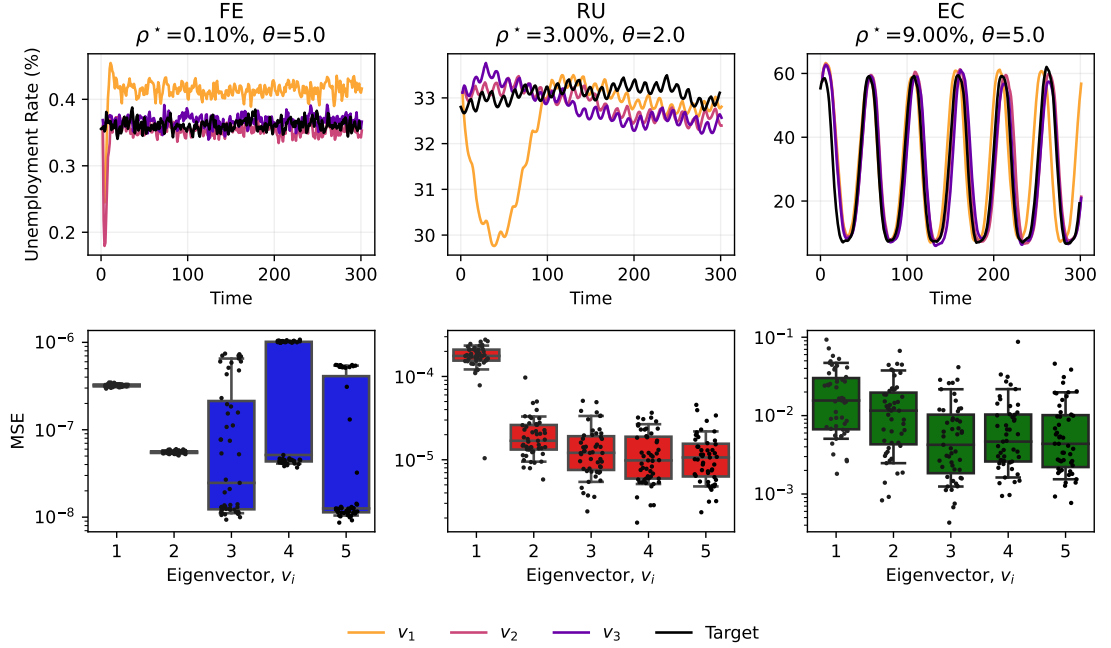


Figure 5.6 – Effect of perturbations into eigendirections of the MSE-Hessian using $T = 300$ and $S = 50$ for each of the three parameter-points. (Top) Time series realization of the unperturbed parameters (black), compared to perturbations into the $v_1(\Phi)$ (orange), $v_2(\Phi)$ (red) $v_3(\Phi)$ (purple) directions. The sign of the eigenvector was left fixed from the numerical singular value decomposition. (Bottom) Mean squared error of the perturbation in $\pm v_i(\Phi)$ (x-axis) over a sample of 50 $S = 1$ Hessians. The step size for all perturbations is $\varepsilon = 0.1$, and $T = 300$.

was the case for the finite difference estimation (Appendix C): any perturbation of the parameters shifts around stochastic effects leading at least some baseline degree of MSE. In the Mark-0 model for instance, firm entry is stochastic which means any change to the parameters changes the exact entry (and exit) times of firms, even if the realizations of the noise are the same. Thus, for phases like RU and EC where there are bankruptcies, this can lead to a baseline MSE irrespective of the direction of the perturbation. A second aspect to this is that the eigenvectors are the infinitesimal ones whereas the step-size used in Figure 5.6 is $\varepsilon = 10^{-2}$, such that the Hessian may not be a perfect approximation at this distance. This is such that one has a perturbation large enough for the parameter hierarchy to become visible among the noise effect mentioned above and to compare the different phases on an even footing. The choice of $\varepsilon = 10^{-2}$ is not necessarily aligned with the natural choice of $1/\sqrt{\lambda_i(\Phi)}$ (Eq. (3.2)), which differs among the different directions.⁴ Combining these two aspects, it is unsurprising that the ill-constrained directions have an effect that one might consider equivalent to the baseline MSE from a random perturbation simply because the step-size needed for them to have

⁴Figure 5.3 suggests selecting $\varepsilon \sim 10^{-1}$ for FE, $\varepsilon \sim 10^{-2}$ for RU, and $\varepsilon \sim 10^{-10}$ for EC for the magnitudes of their MSE to be equivalent.

a similar effect as $v_1(\Phi)$ is not even closely reached.⁵ One needs to pick a ε that strikes this middle ground between boosting signal above the noise from changes in parameters but only so far that the infinitesimal Hessian is still a reasonable approximation of the correct direction to take. What is important for the purposes of this thesis is that there are directions, namely $v_1(\Phi)$, that consistently display a larger loss than can be expected of a random direction.

5.4 Dynamic Stochastic General Equilibrium Models: A Note

This far in the chapter, I have shown that the simple Mark-0 ABM displays a typically sloppy hierarchy of directions in parameter space. Some might argue this to be a downside of ABMs, as their individual parameters are generally not well-constrained and thus fitting such models to data may be futile (the reverse is true, as noted in Chapter 2). Nonetheless, some might wish to fall back to empirically constraining the *standard* approach to macroeconomic modeling: Dynamic Stochastic General Equilibrium (DSGE) models (see Chapter 1 for an introduction). This approach has a large following and a rich literature of different models, as well as how to fit them to data and study exogenous shocks. However, DSGE approaches are not immune to a sloppy parameter space. Figure 5.7 shows the eigenvalue spectra of six DSGE models common to the literature: three models from Galí (2015) that make up the New Keynesian baseline, and three calibrated models including Smets and Wouters (2007), Del Negro et al. (2015), and Carlstrom et al. (2017). As can be seen in the figures, each of these models has a spectrum spanning more than six decades, making them sloppy models as their parameters are also not well-constrained.

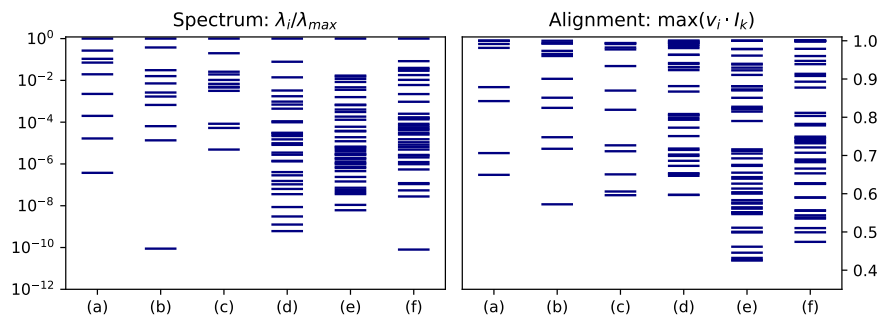


Figure 5.7 – Eigenvalue spectrum (left) and cosine similarity to the closest bare parameter axis (right) for six basic DSGE models: the basic models of (a) Galí (2015) Chpt. 2, (b) Galí (2015) Chpt. 3, (c) Galí (2015) Chpt. 8, and the estimated models of (d) Smets and Wouters (2007), (e) Del Negro et al. (2015), and (f) Carlstrom et al. (2017). Hessians were estimated in a *large data* setting with $S = 100$ random realizations of $T = 5000$ time steps and a numerical perturbation of $1e-3$. Results are similar when using impulse response functions and shorter time series realizations.

Taken from Naumann-Woleske et al. (2023)

⁵Even for the FE case, $v_1(\Phi)$ implies $\varepsilon \sim 10^{-1}$, and $\varepsilon \sim 10^1$ for $v_2(\Phi)$ to show a similar loss. For RU this is $\varepsilon \sim 10^{-2}$ vs. $\varepsilon \sim 10^{-\frac{1}{2}}$. For EC it is even greater at $\varepsilon \sim 10^{-10}$ vs. $\varepsilon \sim 10^{-5}$. All values based on $S = 1$ Hessians.

As with the Mark-0 model studied above, the most well-constrained parameter space directions may not correspond to individual parameter axes (see Figure 5.7 and 5.8). Nonetheless, studying the most well-constrained directions, one notes that the parameters they represent are those controlling the models' noise processes, such as their persistence or covariance. This reflects the adiabatic structure of these models: they are formulated as deviations around a steady state, implying that shocks dominate the effects of any given period. Consequently, fixing the shock process means shifting other parameter combinations has comparatively little effect on the models dynamics, whether one is considering a time series output or an impulse response function.




Figure 5.8 – Eigenvectors for six basic DSGE models: the basic models of (a) Galí (2015) Chpt. 2, (b) Galí (2015) Chpt. 3, (c) Galí (2015) Chpt. 8, and the estimated models of (d) Smets and Wouters (2007), (e) Del Negro et al. (2015), and (f) Carlstrom et al. (2017). Hessians were estimated in a *large data* setting with $S = 100$ random realizations of $T = 5000$ time steps and a numerical perturbation of $1e - 3$. Results are similar when using impulse response functions and shorter time series realizations.

In sum, it thus appears that equilibrium micro-foundations are not immune to the sloppy parameter space structure, in particular as the number of parameters in these models grows (e.g. see models c and d in Figures 5.7 and 5.8).

Key Messages

- The Mark-0 model displays the characteristic parameter-hierarchy of sloppy models: the eigenvalues span multiple decades in a roughly uniform manner.
- Using automatic differentiation, one can obtain a good estimate of the first eigenvalue-eigenvector pair using only one random realization and 150 time steps, with more ill-constrained directions requiring more data, limited more by duration than the number of random realizations.
- Perturbing the model in the most well-constrained directions leads to a consistently high MSE-loss.
- DSGE models also display a sloppy parameter-hierarchy, with eigenvalues spanning many decades. In these models it is the parameters relating to the exogenous noise processes that primarily drive the dynamics.

Exploration of the Mark-0 Agent-Based Model

Adapted from Naumann-Woleske et al. (2023) *Exploration of the parameter space in Macroeconomic Agent-based Models*, and significantly extended. The original paper was joint work with M. Knicker, J-P Bouchaud and M. Benzaquen. However, while this chapter is inspired by that work all results are *new* based on the differentiable Mark-0 model and an original approach to the exploration algorithm.

The observed dynamics of the Mark-0 agent-based model depend only on a few influential parameter combinations, as alluded to in Gualdi et al. (2015) and shown in Chapter 5 by studying the eigen-decomposition of the model’s Hessian (see Chapter 3 for methods). Since the model’s dynamics are sensitive only to some parameter-directions, it stands to reason that perturbing the model’s parameters into these well-constrained directions would uncover a new dynamic that is “as different as possible” to the dynamic of the unperturbed parameters. In models where there are multiple different phases of qualitatively different behavior, iterating this procedure may lead to the discovery of the entire set of scenarios the model is able to generate. Based on our work in Naumann-Woleske et al. (2023), in this chapter I will show that it is possible to systematically explore the phase space of a model by taking a (not so random) walk down the well-constrained directions.

To develop an intuition for the mechanics of such a walk, I consider first the two-dimensional Φ_{ρ^*} - Φ_{θ} plane where all four phases of the model can be recovered (this was originally studied in Gualdi et al., 2015). Studying the vector-field of the first eigenvectors across different points in this field suggests that a walker traversing this field would pass through each of the different phases present in the Mark-0 model.

Building on this, the remainder of the chapter proceeds as follows: first, I outline a simple heuristic algorithm to follow the well-constrained directions through the model’s parameter space. This algorithm requires only a few key decisions, and I elaborate on two extensions. Applying this algorithm to the Mark-0 model introduced in Chapter 5, I show that it is possible to recover all of the phases in the Mark-0, as defined by their original authors (see Chapter 5 and Gualdi et al., 2015). More generally, it is possible to recover at least two phases in most realizations of the algorithms. However, there are weaknesses to such a simple heuristic. The primary weakness being “flip-flopping” back and forth in the parameter space, as the sign of the eigenvector is undefined. I show how this can be addressed both by probing both directions and by restricting the angle of travel. There are also weaknesses due to the choice of the MSE-loss function, such as an exponential sensitivity to small shifts in the frequency of oscillation (see Chapter 5). While for Mark-0 I know a priori what the different phases are, thanks to Gualdi et al. (2015), in general this is not the case, and the objective is to ascertain the phases by exploring. To do so, I introduce the natural measure of pairwise MSE between the realizations at each step of the algorithm. This measure can be used as a consistent guidance tool for separating explorations with more variety in phases from those with less, though it is subject to the same weaknesses as the MSE. The results in this chapter are a proof that one can indeed effectively traverse the phase-space of a MABM using a simple heuristic following the most well-constrained direction. Ideally one would like to try more modularities, such as automatic step-size selection, to develop a maximally efficient algorithm. However, this is left for future research.

6.1 Phase Transitions in Two Dimensions

The primary reason for choosing the Mark-0 model to develop an exploratory algorithm is that the authors of the original model developed several phase diagrams (see Gualdi et al., 2015, 2017; Bouchaud et al., 2018, for a full set). In particular, consider here the case of the Φ_{ρ^*} - Φ_{θ} plane shown in Figure 6.1, where there are four distinct dynamics of the unemployment rate separated by mostly second-order phase transitions. The left panel of Figure 6.1 shows a vector field populated by $v_1(\Phi)$ overlaid on the mean unemployment rate, while the right panel shows the same vector field overlaid on the standard deviation of unemployment.¹

Considering the overlaid vector field for the unemployment rate dynamics in this two dimensional plane, we find that the stiffest direction consistently points in the direction of a nearby phase transition. One might remark that the majority of eigenvectors here align with either the Φ_{θ} or the Φ_{ρ^*} axes in Figure 6.1. This is for two reasons. The first concerns the economic interpretation. For instance, in the full employment phase there are no firms close to the bankruptcy threshold Φ_{θ} such that perturbing this threshold by only a little would have zero effect, whereas a change in the interest rate Φ_{ρ^*} changes the

¹As noted earlier, picking these moments hides some facets of the dynamics (e.g. EC and RU have similar mean values). It serves here simply to illustrate the phase-space as the classification into these phases is based on the mean and standard deviations.

level of consumption and thus has a non-zero effect. Accordingly, in this case the first eigenvectors are aligned with the Φ_{ρ^*} axis. Conversely, in the RU and EC phases there are always firms close to the bankruptcy threshold such that Φ_{θ} gains in importance as small shifts can lead to more or less firms entering bankruptcy throughout the model's realization. Since the threshold is itself very low, a small change has larger effects on firms' bankruptcy and revival. For the area where the threshold and interest rate are both sufficiently large, one can then observe an interplay between the two regimes (e.g. $\log_{10} \Phi_{\rho^*} > -1.8$, $\log_{10} \Phi_{\theta} > 0$). Despite the apparent simplicity of this diagram, where relative to one another only one of the two parameters has an effect, one can still note that following around the first eigenvector will lead one to cross at least one phase transition, if not multiple.




Figure 6.1 – Phase diagram of the Mark-0 unemployment rate in the Φ_{ρ^*} - Φ_{θ} plane based on Figure 2 of Gualdi et al. (2017). (Left) Mean level of unemployment over $T = 200$ periods, (Right) Standard deviation of the unemployment rate over $T = 200$ periods. Black arrows are the first eigenvectors computed by automatic differentiation, the sign of the vector is such that it has the closest similarity to the $[1, 1]$ vector. Parameterization as given in Table 5.1, except $\Phi_{\tau_{meas}} = 1.0$, $\Phi_{\tau_{tar}} = 0.0$ with $N = 3000$, $T = 200$

An important qualification to make here is that by phase transition I am referring to second-order transitions, where there is a continuity between one phase and another, however small. A priori, first order transitions cannot be discovered in this manner for there is no increase in loss up until the threshold and thus the Jacobian would not be able to capture this as there is no gradient. One can accidentally discover a first order transition using finite differences if one happens to be close enough to the transition line for the small perturbation used in computing the Jacobian to cross the threshold, but this is an edge case upon which one cannot rely.

6.2 A Simple Algorithm for Parameter space Exploration

The local parameter space decomposition at each of the points studied so far indicate that, as expected, well-constrained directions lead to the largest change in the MSE loss-function (Chapter 5), and are likely to cross phase boundaries if one continues traveling along them (Section 6.1). This section outlines a simple heuristic algorithm to walk along the well-constrained directions. The goal of this algorithm is to determine a sequence of parameters $\mathcal{E}_Q = \{\Phi(q)|q = 0, \dots, Q\}$ of length Q , whose observables $\mathbf{x}_{s,t}(\Phi(q))$ at each point are as different as possible. The idea is to obtain a complete coverage of the different dynamics that a model can generate within the user-defined parameter space hypercube \mathcal{B} . One way of measuring the coverage of a given exploration is to continue using the Mean Squared Loss which I use here also for the Hessian. In this case, we can measure the *quality* of a given algorithm with respect to its ability to span the space of dynamics with as few steps as possible by the average pairwise mean-squared error

$$pMSE(\mathcal{E}_Q) = \frac{1}{Q(Q-1)} \sum_{\Phi(q_1), \Phi(q_2) \in \mathcal{E}_Q, q_1 \neq q_2} \mathcal{L}^{MSE}(\Phi(q_1), \Phi(q_2)), \quad (6.1)$$

where I use an average to penalize a large number of steps. While this is self-consistent, as with the MSE Hessian it is not foolproof: periodically oscillating phases like EC might see small frequency shifts that lead to high MSE, even in a pairwise sense.

The simplest way of targeting $pMSE(\mathcal{E}_Q)$ is to take steps into the direction of the first eigenvector. By definition, this would be the vector along which the largest change in the loss function should occur. We can thus write

$$\Phi(q+1) = \Phi(q) \pm \varepsilon v_1(\Phi(q)) \quad q \geq 1 \quad (6.2)$$

from which we can observe two key decisions that must be made:

1. Step size ε : Based on the observations from Chapter 5, one needs to set a step-size large enough to change the dynamics beyond changing the exact realization of the noise, while keeping ε small enough such that the local approximation to the Hessian holds. At this point, one might ask whether the step-size itself should be related to $\lambda_1(\Phi)$, perhaps such that one targets a fixed MSE increase as suggested in Eq. (3.2). There are two reasons I choose not to do this, the primary being that letting $\varepsilon \sim 1/\sqrt{\lambda_1(\Phi)}$ means specifying a targeted amount of MSE loss, which is intuitively more difficult than saying the maximal parameter change is ε in log-space (e.g. $\varepsilon = 0.1$ would mean $\sim 10\%$ at most). Furthermore, in cases like the EC phase this might lead to remaining “stuck” in shifting the frequency because the exponential $\lambda_1(\Phi)$ scaling would imply that the parameters wouldn’t change much. By contrast, fixing ε means that one makes sure to actually shift the parameters significantly in this case. Of course, this approach can also lead one to miss phases or dynamics when $1/\sqrt{\lambda_1(\Phi)}$ is much smaller than the fixed value. For a finite difference approach, ε should be at least the perturbation size used in the finite

difference Hessian, such that one can ensure that the chosen direction is the most well-constrained.

2. Sign of the Eigenvector: Since the sign of the eigenvector is indeterminate one must choose whether or not to traverse it in the positive or negative direction. If the model manifold were symmetric, this would likely be irrelevant since one obtains the same degree of loss irrespective of the sign. However, for models with asymmetric manifolds, such as the Mark-0 model, the choice of sign can make a large difference. Consider the case of being near a phase boundary, in one direction one can find entirely different dynamics (i.e. a high loss), whereas its inverse leads to more of the same (i.e. a lower loss).
3. The starting point $\Phi(0)$: ideally, the starting point of the exploration should be roughly centered on the boundaries \mathcal{B} of the parameter space. A reason for this being that steps occur in log-space, which means it can take a long time to traverse from one boundary of the hypercube to the other even if $v_i(\Phi)$ is perfectly aligned with the parameter axes. One can compute the minimal boundary traversal time by taking the difference in log-bounds and dividing by the chosen step-size ε . However, there is no clear "best" heuristic here when choosing a single starting point. Of one has the option to run multiple explorations then it would be reasonable to spread the starting points uniformly over the hypercube made up by the boundaries.

A final consideration to take is how to determine the total number of steps Q to take from the starting point $\Phi(0)$. The simplest response would be to keep the number of steps fixed based on the computational budget available. A smarter approach might be to define the stopping criterion based on the $pMSE(\mathcal{E}_Q)$, such that when no new information is gained, the algorithm stops exploring and one could use the remaining part of the computational budget to launch another exploration from a different starting point.

A Simple First-Pass Approach

The potentially simplest way of addressing these questions is to: (a) fix the sign of the eigenvector by taking whichever sign the singular value decomposition yields, and (b) fix the step-size to a given ε , such as 10^{-1} as done with the perturbations of Figure 5.6. The exception to keeping a fixed sign is if the perturbation would transgress the boundary of the parameter space hypercube \mathcal{B} defined by the user. Here I thus propose that one takes the sign of the eigenvector that does not transgress the boundary, or terminates the algorithm should both directions do so. For the remainder of the chapter I will call this approach the *SimpleAlgo*, and its pseudocode implementation is given by Algorithm 1 below.

Improving the Search Direction

The *SimpleAlgo* approach provides a solid first-pass approach, but, one can likely do better along the lines of the three decisions outlined. In particular, given a likely asym-

Algorithm 1 *SimpleAlgo* : Maximally Simple Eigenvector Exploration

Require: Initial parameters $\Phi(0)$, maximum step count q_{\max} , step-size ε , boundaries \mathcal{B}

Set $q = 0$

repeat

 Compute the Hessian $H(\Phi(q))$ ▷ S and T based on modeler’s choice

 Decompose to determine $v_1(\Phi)$

 Compute new parameters $\Phi(q + 1) = \Phi(q) + \varepsilon v_1(\Phi)$ ▷ Φ is in Log-terms

if $\exists p \in \Phi(q + 1)$ violating bounds $\mathcal{B}(p)$ **then**

 Set $\Phi(q + 1) = \Phi(q) - \varepsilon v_1(\Phi)$ ▷ Try other direction

if $\exists p \in \Phi(q + 1)$ violating bounds $\mathcal{B}(p)$ **then**

 Abort Exploration

end if

end if

$q = q + 1$

until $q = q_{\max}$

metric loss landscape one can choose the sign of the eigenvector in a smarter way by probing both directions and selecting the sign yielding a better $pMSE(\mathcal{E}_Q)$ outcome. This simple trick should avoid “missing” any interesting thresholds. In particular, for a finite difference approach where Hessians take $2SP$ computations, suggesting that $2S$ computations to determine the $pMSE(\mathcal{E}_Q)$ increasing direction may be prudent when parameters are numerous or computational time is long. Algorithm 2, hereafter referred to as *ProbeAlgo*, outlines the modified algorithm.

6.3 Understanding the Map from Parameters to Phases

The phases of the Mark-0 model are classified using the mean and standard deviation of the unemployment rate based on the work of Gualdi et al. (2015). To gain an intuition of how the parameter space is split into the different phases, Figure 6.2 presents the mean and standard deviation of $N = 2^{15} = 32,768$ simulations of a Sobol sequence in the parameter space, together with their ensuing classifications. From Figure 6.2 emerges the intuitive result that a uniform sampling of the parameter space is not equivalent to a uniform sampling of the prediction space of the model. For the case of Mark-0, one gets a large sampling of the full employment and full unemployment phases, with a comparatively small number of samples in either the EC and RU phase. In particular, since a Sobol sequence uniformly covers the parameter space given by \mathcal{B} , one can take the proportion of points belonging to each phase as an estimate of the volume of the parameter space that is filled by one particular phenomena. In this case, the Full Unemployment phase covers nearly half (49.7%) of the parameter space volume, with the Full Employment covering another 27%. Thus, while Sobol sampling can recover all phases, it might take longer to phenomena occurring for only a few parameters but covering a large area in the prediction space to be retrieved. In the same vein, one obtains

Algorithm 2 *ProbeAlgo* : Improved Eigenvector Exploration

Require: Initial parameters $\Phi(0)$, maximum step count q_{\max} , step-size ε , boundaries \mathcal{B}
 Set $q = 0$
repeat
 Compute the Hessian $H(\Phi(q))$ ▷ S and T based on modeler's choice
 Decompose to determine $v_1(\Phi)$
 Compute candidates $\Phi^\pm(q+1) = \Phi(q) \pm \varepsilon v_1(\Phi)$ ▷ Φ is in Log-terms
 if $\exists p \in \Phi(q+1)$ violating bounds $\mathcal{B}(p)$ **then**
 Set $\Phi(q+1) = \Phi(q) - \varepsilon v_1(\Phi)$ ▷ Try other direction
 if $\exists p \in \Phi(q+1)$ violating bounds $\mathcal{B}(p)$ **then**
 Abort Exploration
 end if
 end if
 Compute probes $\mathbf{x}_{s,t}(\Phi^+(q+1))$ and $\mathbf{x}_{s,t}(\Phi^-(q+1))$
 Compute pairwise MSE with previous steps: $PairMSE(\Phi^\pm(q+1))$
 Set $\Phi(q+1) = \operatorname{argmax} PairMSE(\Phi^\pm(q+1))$
 $q = q + 1$
until $q = q_{\max}$ or $PairMSE(q) < PairMSE(q-1) < PairMSE(q-2)$

a large sampling of potentially less valuable phases, in this case the full unemployment phase. Thus, one key characteristic of a successful algorithmic exploration would be to recover the less likely EC and RU phases.

6.4 Recovering Mark-0's Phases and their Transitions

Experimental Setup

Taking here the conclusions of Chapter 5 on the estimation of $v_1(\Phi)$, I consider the case of $S = 1$ and $T = 150$ for the exploratory algorithm. This is a small sample but as Chapter 5 showed, these are sufficient to estimate the first eigenvectors when using automatic differentiation. Using the parameter boundaries defined in Table 5.1, the starting points are chosen as a Sobol sequence of length 64 to fill the hypercube determined by the bounds with a low discrepancy.² This approach should reduce the sensitivity of the ensuing results to the exact choice of initial parameters, such as starting near a phase boundary. However, it will also likely imply that there are some realizations of the algorithm that fail immediately as the model may crash, as with the NAN phases already noted in Figure 6.2. I choose to exclude these starts from the ensuing analysis. Finally, I apply a step size of $\varepsilon = 0.2$, whose interpretation is that the maximal perturbation for a single parameter is ε in log-space, which, while large should make traversing the large hypercube of parameter bounds quicker.

²See Sobol (1967) for a definition of the sequence, and Niederreiter (1988) for a definition of discrepancy.

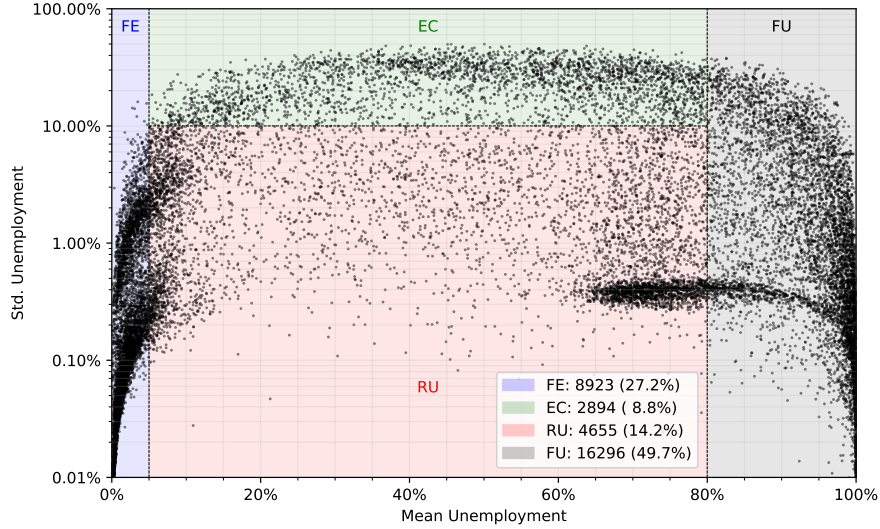


Figure 6.2 – Mean and Standard Deviation of the Unemployment Rate for a Sobol sequence of length $2^{15} = 32,768$ in the parameter space hypercube \mathcal{B} given by Table 5.1. N/A values were assumed to mean 100% unemployment and filled accordingly, there were 7487 simulations containing only N/A values, which implies the economy crashed during the equilibrating period (i.e. in FU). All samples are based on $S = 1$ and $T = 300$, as with the exploration, and a equilibrating period of $T = 1000$

For *SimpleAlgo*, the initial conditions are fixed at the steady state ($T = 1000$ equilibration point) of the initial parameters $\Phi(0)$, therefore, many of the differences in phase also include a transient dynamic as the model adjusts to the shift from $\Phi(q)$ to $\Phi(q + 1)$, which may take longer than the observation period of $T = 150$. This approach is self-consistent since the parameters are perturbed under the same initial conditions as the eigenvectors were calculated, such that the only change occurring is a change in parameters. However, it also means that we explore both the changes in the dynamics of the transients and differences in the model’s steady state. From a model exploration standpoint, one might be interested more in the space of steady states than in the reaction of the model to perturbations in parameters (i.e. moving from the initial conditions’ steady state to the steady state of the new parameters). In this case, one can commit a slight of hand, and re-compute the initial conditions after perturbing the parameters. The idea being that the change in the parameters should lead to a change in steady state (recall for example the perturbation of the FE phase in Figure 5.6), and if one is interested solely in these steady states then one can perturb the parameters and run the model for another $T = 1000$ periods to equilibrate the system. This new steady state then serve as the initial conditions for the next computation of the Hessian. The result is that the Hessian is computed on $T = 150$ periods in the $\Phi(q)$ steady state, rather than having some periods of transition (e.g. Run 15 of *ProbeAlgo* in Figure 6.3 below). Recalling the classification of Francis and Transtrum (2019) (see Chapter 3 Section 3.2), this effectively approach removes transient eigenvectors, whose eigenvalues scale $\mathcal{O}(1/T)$, from consid-

eration as the next direction, leaving only those controlling steady state levels (scaling $\mathcal{O}(1)$) and those controlling oscillations (scaling $\mathcal{O}(T^2)$ or in the EC case also $\mathcal{O}(\exp T)$). Therefore, in the remainder of this Chapter, I will distinguish between two cases of the *SimpleAlgo* approach: one with fixed initial conditions, and one in which they are allowed to vary with the parameters, hereafter *SimpleAlgoFC* for Floating Condition (the same holds for *ProbeAlgo* and *ProbeAlgoFC*).

Can one recover the phases?




Figure 6.3 – Top two runs with the most phases recovered for the *SimpleAlgo* (first row), *SimpleAlgoFC* (second row), *ProbeAlgo* (third row), and *ProbeAlgoFC* (fourth row) approaches with $\varepsilon = 10^{-2}$. Initial parameterizations based on Sobol sequence of length 64. Hessian hyper-parameters are $T = 150$ and $S = 1$. Boundaries from Table 5.1

The simple answer to the question of whether one *can* recover all of Mark-0's phases by following a heuristic like *SimpleAlgo* is *yes*. From the 64 different starting points, I obtain 47 explorations taking more than one step for *SimpleAlgo* and *ProbeAlgo* (38 for *SimpleAlgoFC*, and 40 for *ProbeAlgoFC*). The initial parameter points that did not




Figure 6.4 – Time series of the unemployment rate for two realizations of the *SimpleAlgo* heuristic. The first corresponds to Run 38 Step 30, also shown in Figure 6.3, the second to Run 15 Step

yield explorations were infeasible (NA following the FU phase) immediately. Of the 47 *SimpleAlgo* explorations, 30 explorations completed all 30 steps of the computational budget (22 for *SimpleAlgoFC*, 33 for *ProbeAlgo* and 20 for *ProbeAlgoFC*), with the primary reason for early termination being violations of the boundaries \mathcal{B} , and in some FU cases that the simulation crashed leading to an invalid Hessian matrix.

Figure 6.3 shows the top two realizations of the successful runs of *SimpleAlgo*, *SimpleAlgoFC*, *ProbeAlgo* and *ProbeAlgoFC* with the most different phases recovered (at least 3, and then ranked by $pMSE(\mathcal{E}_Q)$). One can clearly see the wide variety of dynamics that the Mark-0 model can generate across the different parameterizations explored by the simple heuristic of following $v_1(\Phi)$. For instance, Run 41 of *SimpleAlgo* and Run 15 of *ProbeAlgo* both display multiple qualitatively different phases from FE to RU and EC. Indeed, there are several phases that *SimpleAlgo* and *ProbeAlgo* passed through that, to wit, have not been studied explicitly in previous work on the Mark-0 model, or even our own work on exploring the phase space (Naumann-Woleske et al., 2023).³ Two such phases are highlighted in Figure 6.4, wherein there can be alternating high and low spikes in the unemployment rate, with high peaks at $\sim 20\%$ and lower ones at 5%. This phase is described by an active Central Bank with a high baseline interest rate together with firms that react sharply to inflation and interest rates through price-adjustments and the hiring/firing of employees. The right panel of this figure shows that this type of phase also exists with an amplified spike.

Figure 6.5 shows the projection of all of the successfully computed steps in the mean-standard deviation plane. Due to the presence of transient dynamics, the mean and standard deviations of the *SimpleAlgo* and *ProbeAlgo* realizations are based only on the last 50 time steps, which have generally made the transition to the new steady state. However, standard deviations might still be somewhat inflated, which is reflected in the higher prevalence of explorations in the EC phase. Comparing these explorations with the Sobol approach shown in Figure 6.2 suggests that the algorithmic exploration approaches do cover the same prediction space (here in 2D) in roughly the same density as a brute

³See Gualdi et al. (2015, 2017); Bouchaud et al. (2018); Sharma et al. (2020); Knicker et al. (2023).

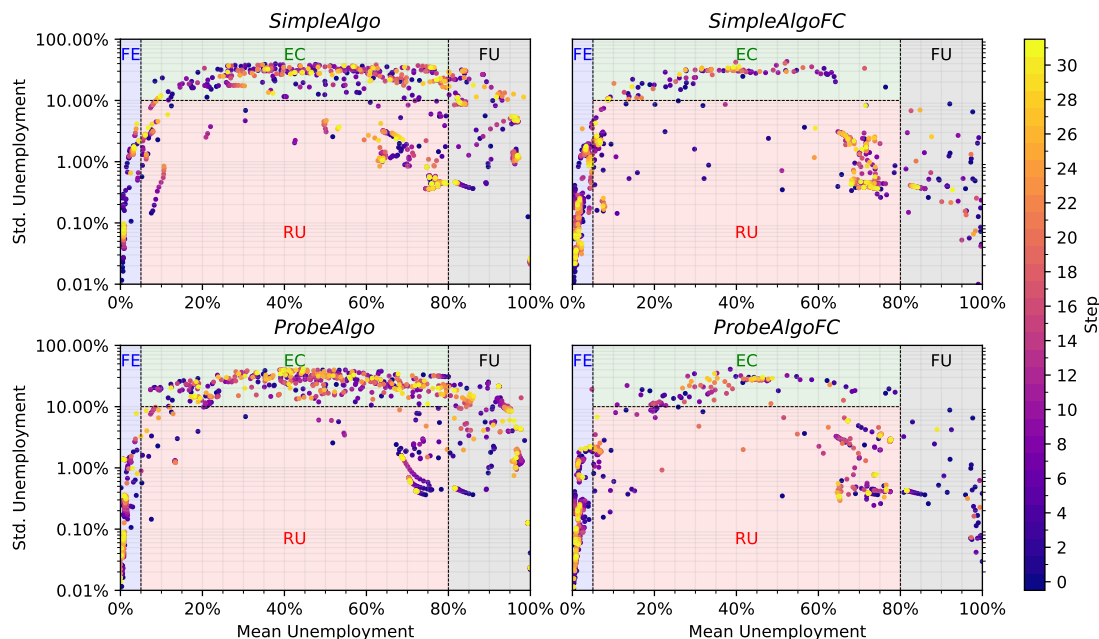


Figure 6.5 – All points traversed by *SimpleAlgo* (top left), *SimpleAlgoFC* (top right), *ProbeAlgo* (bottom left) and *ProbeAlgoFC* (bottom right) projected into the Mean-Std. Deviation space. Color coding indicates the step of the algorithm. Note *SimpleAlgo* and *ProbeAlgo* contain transient dynamics, such that mean and standard deviations were based only on the last 50 time steps to better conform with the phase definitions of Gualdi et al. (2015)

force parameter space sampling. Unlike brute force sampling however, the estimation of the Hessians also provides an understanding of the local sensitivity landscape.

How well can we recover different phases?

The example explorations shown in Figure 6.3 are promising, but also raise questions about the quality of this heuristic. For instance, both figures suggest that it is possible to become trapped in a particular phase with little change, based on the plethora of FU and FE realizations and the oscillatory regimes of *SimpleAlgo* run 34.⁴ The main goal of the algorithm is to recover the different phases of the Mark-0 mode. However, before analyzing the phase-discovery results, one major aspect to note here is that the ability of the exploration to recover the different phases of the Mark-0 model here relies on the definition of phases used in Gualdi et al. (2015) and thereafter, who consider only the mean and standard deviation as a guideline for classifying the dynamics of the unemployment rate. While this is a simple rule of thumb, it is also not entirely satisfying for there are also oscillatory phases classified as RU or FE (e.g. the dynamic of Run 38 Step 30 shown in Figure 6.4 is an RU phase when it should really be an EC phase).

⁴The shifting frequencies is the result of the exponential sensitivity of the MSE to shifts in oscillatory frequency, as noted in Chapter 5.




Figure 6.6 – Number of recovered phase transition by *SimpleAlgo* (top left), *SimpleAlgoFC* (top right), *ProbeAlgo* (bottom left) and *ProbeAlgoFC* (bottom right). The values indicate: (numerator) the number of transitions at step q where $v_1(\Phi(q-1)) \cdot v_1(\Phi(q)) > -0.9$, and (denominator) the total number of transitions found. The color code indicates the fraction of total steps described by each type of transition

In principle, one should explore the space of qualitative dynamics and only thereafter cluster the dynamics to determine the set of qualitative phases.

Figure 6.6 shows the prevalence of different types of transitions recovered by the different explorations. The graph suggests that across all algorithms there is a roughly symmetric structure to the transitions discovered, and in particular for the *SimpleAlgo* approach this is due to flip-flopping. Specifically, some explorations go backwards and forwards across the same phase boundary as $v_1(\Phi(q-1)) \sim -v_1(\Phi(q))$ (see Figure D.1 of Appendix D). Adjusting for this by eliminating transitions where $v_1(\Phi(q-1)) \cdot v_1(\Phi(q)) < -0.9$ (the numerators in Figure 6.6), one can see that this strongly reduces the number of transitions in the *SimpleAlgo* case, whereas *ProbeAlgo* is not strongly affected.⁵ This means that when using the *ProbeAlgo* approach, after a transition one does not immediately transition back into the prior phase, as this decreases $pMSE(\mathcal{E}_Q)$, as opposed to further investigating the new phase. However, Figure 6.6 also shows that within a given phase

⁵This condition is equivalent to saying one places a restriction on the cosine similarity, i.e. the angle between two P-dimensional vectors.

over half the steps within phases violate the above condition. That is, within a phase, both *SimpleAlgo* and *ProbeAlgo* often oscillate back and forth between two areas of the parameter space that lie in the same phase. They may of course be numerically very different: in the EC and RU phase this may be a shift in the frequency of oscillations or a change in the transition dynamics, which leads the exploration to remain in the same phase but still generate a high MSE (see Figure 6.3 top left as an example). While not desirable from a purely exploratory point of view it does suggest a consistency of direction within a given phase: roughly the same parameter-combinations have the largest effect within phases and, generally also lead to phase transitions.

6.5 What Drives the Phase Transitions?

One of the biggest selling points of this type of exploratory approach over other approaches such as meta-modeling, is the information gained on the local sensitivity of the model. Section 6.4 showed that the heuristic algorithms can recover multiple different phases within a single exploration, which implies that we know the direction in parameter space along which this transition occurs. This means one can pose the question of whether there are more “general” parameter combinations that regulate between different phases of the model. In each case, one could see the transition direction akin to the localized normal of the hyper-surface that is the phase transition. In the same vein, the flip-flopping of the *SimpleAlgo* approach also hint at a consistent underlying sensitivity to parameter combinations in each of the respective phases.

In the case of Mark-0 for instance, one might be interested in what drives the economy out of a full-employment phase. Figure 6.7 shows the set of $v_1(\Phi)$ (and their sign) taken by the exploration heuristics when moving within (leftmost panel) and transitioning into and out of the Full Employment phase (right panels). While all of the directions are parameter combinations, implying one should not jump to conclusions about the importance of one parameter over another, one can tentatively point out the prevalence of the firms’ price-setting Φ_{γ_p} , its relation to the wage adjustment Φ_{γ_w} through $\Phi_r = \Phi_{\gamma_w}/\Phi_{\gamma_p}$, as well as the wage-adjustment to inflation Φ_{g_w} as key parameters governing the changes in the unemployment within the FE phase and for transitions into and out of the full employment phase. In line with this observation, our study on the Post-COVID inflation dynamics (see Chapter 9 in this thesis and Knicker et al. (2023)) also highlights the dangers of a wage-price spiral dependent on the balance of Φ_{γ_p} and Φ_{γ_w} . Taking this point further to study also the EC transitions and FU transitions (Appendix D.2), the parameter combinations are more mixed but nonetheless Φ_{γ_p} is frequently a large component of the eigenvector leading to a transition. While these interpretations are tentative it nonetheless suggests that if one does have a larger computing capacity to run a higher number of explorations, one could piece together the critical relations between these parameters that tip the economy from the full-employment phase to other phases, thus building up a phase diagram. But even with such a small sample, one can gain an intuition about some of the stronger drivers of the unemployment dynamics in the Mark-0 model.

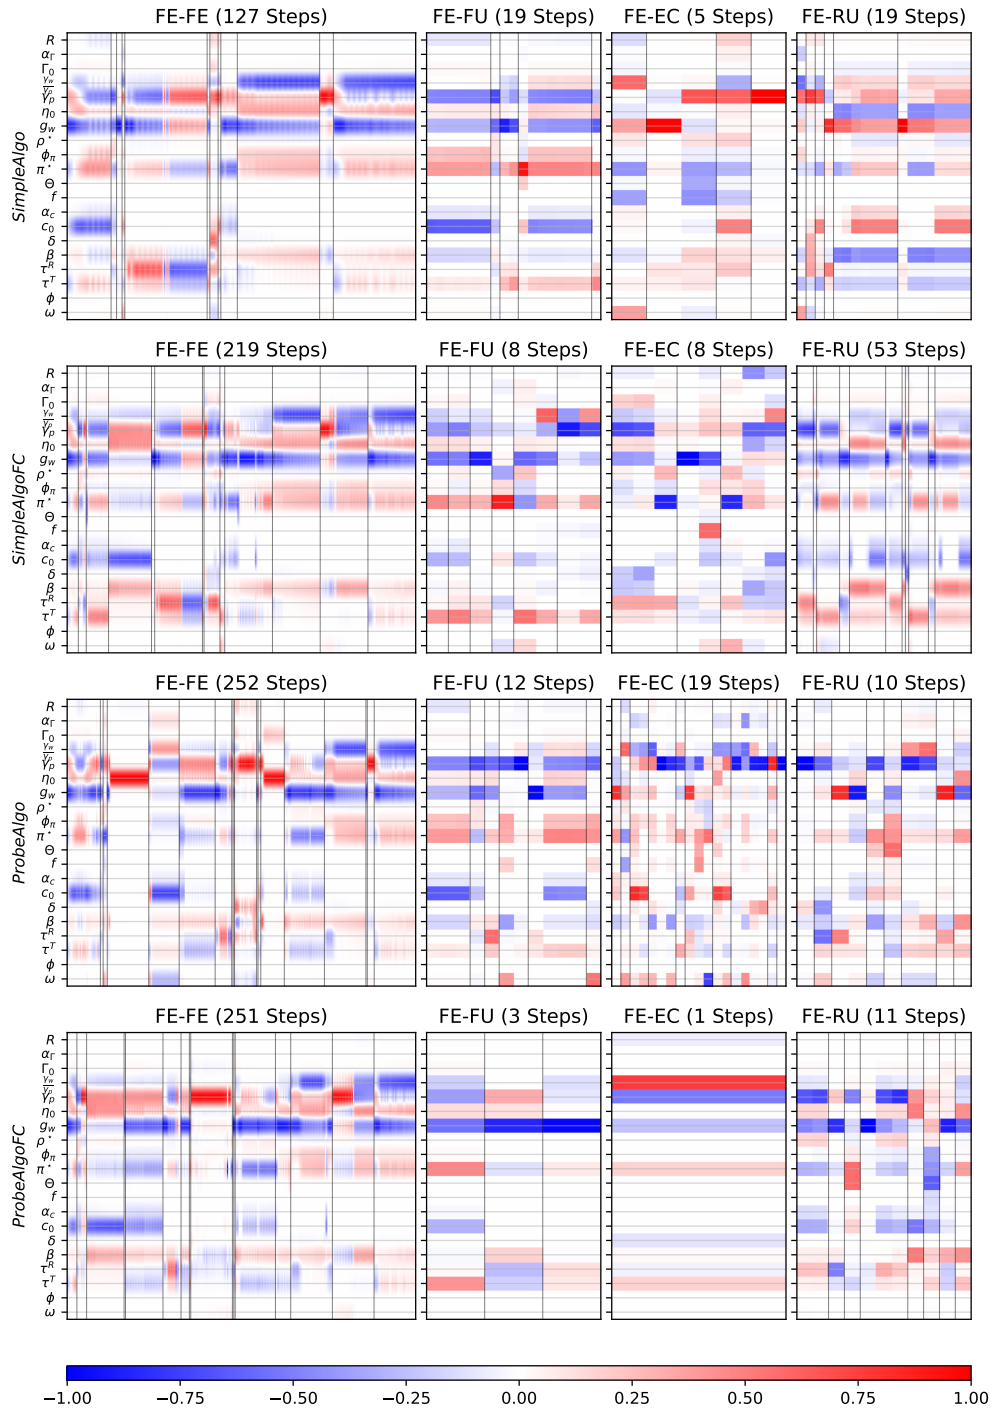


Figure 6.7 – Directions taken by the explorations (vertical) within, into and out of the FE phase in the Mark-0 model, as defined by Gualdi et al. (2015). The vertical lines demarcate different exploration runs.

6.6 Distinguishing Explorations Quantitatively: the $pMSE(\mathcal{E}_Q)$

The previous sections evaluated the explorations on the basis of the phases discovered, using the classifications of Gualdi et al. (2015). But, in a scenario where the phase classifications are unknown ex-ante, it is still important to measure the *quality* of the exploration by a metric capturing the variety of the dynamics that have been explored. The natural metric for measuring an algorithm with Hessians based on the MSE-loss is to consider the pairwise MSE, $pMSE(\mathcal{E}_Q)$, between all steps of the exploration \mathcal{E}_Q , as introduced in Section 6.2. This metric is naturally low when phases are equivalent, and positive should they be different.

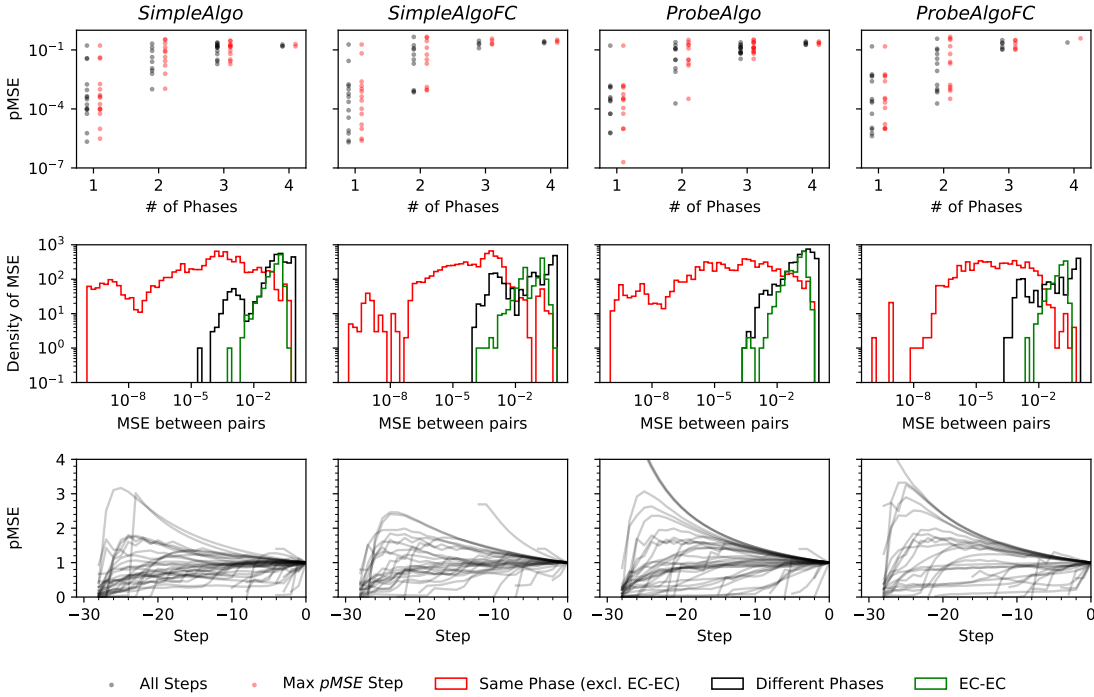


Figure 6.8 – (Top) Pairwise $\mathcal{L}^{MSE}(\Phi(q_1), \Phi(q_2))$ between all steps in an exploration \mathcal{E}_Q . Red shows pairs in the same phase, while black shows pairs in different phases. (Middle) Dynamics of the $pMSE(\mathcal{E}_Q)$ as a function of the number of steps, normalized to 1 for the final step, with indexes denoting distance from final step. (Bottom) $pMSE(\mathcal{E}_Q)$ of final step (black) and maximum point (red) versus the number of phases discovered.

Figure 6.8 shows some of the features of the $pMSE(\mathcal{E}_Q)$. The top panel shows the magnitudes of the $pMSE(\mathcal{E}_Q)$ as a function of the number of phases discovered, both when considering the whole set of steps (black) and when stopping the algorithm at its peak $pMSE(\mathcal{E}_Q)$ value (red). In both cases, the $pMSE(\mathcal{E}_Q)$ is, on average, higher the more different phases were discovered, which is its desired behavior. However, one can see that once at least two phases have been discovered the $pMSE(\mathcal{E}_Q)$ of an exploration with two phases may become indistinguishable compared to one of three or

even four. The reason for this is the inherent weakness of the MSE as a loss function. Chapter 5 already highlighted how small shifts in an oscillation can induce large MSE realizations, which in this case would also translate into a high $pMSE(\mathcal{E}_Q)$ though the phase is qualitatively equivalent. The middle panels of Figure 6.8 highlight this effect, the MSE for a shift within the EC phase (green) has a similar distribution as the MSE of realizations in different phases (black), while the remaining within-phase MSEs are significantly lower. One could apply a shifting approach such as dynamic time warping to compare the different realizations $\mathbf{x}_{s,t}(\Phi)$, which may be more robust to these effects, but such approaches are left to future research.

One thing that can be done however, is to reduce computational burden by preemptively stopping the algorithm once the $pMSE(\mathcal{E}_Q)$ flattens. The bottom panel of Figure 6.8 highlights the development of the $pMSE(\mathcal{E}_Q)$ as a function of the number of steps from the final realization of the exploration. In almost all cases, one can see that there is a peak in the $pMSE(\mathcal{E}_Q)$ early in the exploration, often at less than 10 steps, suggesting that for a given computational budget one should run such a heuristic for up to 10 steps to discover most phases with the benefit of selecting a wider variety of different starting points.

Key Messages

- Using simple heuristic approaches like *SimpleAlgo* and *ProbeAlgo*, one can recover the whole spectrum of different phases (and more) in the Mark-0 model.
- Probing the sign of the eigenvector and placing restrictions on the direction of travel can improve the efficiency of exploration.
- For the Mark-0 model, the dynamics within a phase are consistently sensitive to similar parameter combinations, in particular the speed of adjustment of the prices and wages in response to inflation, suggesting that the explorations traverse a lower-dimensional manifold when it comes to the unemployment rate.
- While the MSE has weaknesses, one can nonetheless distinguish more successful runs from less-successful ones by means of a pairwise approach $pMSE(\mathcal{E}_Q)$.

PART II

Confidence and Collective Behavior in Macroeconomic Models

Investment-Driven Business Cycles: A Sentiment-driven Mechanism

Taken from Naumann-Woleske et al. (2022) *Capital Demand Driven Business Cycles: Mechanism and Effects* without modification.

The field of economics has long been aware of a conceptual dichotomy between studies of short-term dynamics and models of long-term growth. An early distinction was made between the Hicks IS-LM model (1937) and the Solow growth model (1956). The developments in both approaches have captured important dynamics at their respective timescales, such as short-term demand effects and endogenous drivers of long-term growth (e.g. Aghion and Howitt, 1992). Yet it is not well understood how the dynamics at different timescales are interlinked and how medium-term disequilibrium dynamics impact the long-term growth trend of the economy.

Since World War II, the United States of America alone has faced twelve recessions. While the severe short-term consequences of these crises are appreciated, understanding of the long-lasting impact on growth remains underdeveloped. The pervasive recurrence of booms and busts has thus sparked research into the linkages between economic volatility and growth (Cooley and Prescott, 1995; Aghion and Howitt, 2006; Priesmeier and Stähler, 2011; Bakas et al., 2019). Theoretical as well as empirical investigations have turned out to be inconclusive, as authors disagree on both the sign and magnitude of the ultimate effect of volatility on growth.¹ Theoretical literature is divided into two dominant strands that stem from either Schumpeterian notions, in which volatility is good for growth (based on Schumpeter, 1939, 1942), or the learning-by-doing concept (based

¹We suggest Bakas et al. (2019) for a comprehensive review of this literature.

on Arrow, 1962), where volatility is detrimental to growth. The conflicting theoretical frameworks and ambiguous empirical findings indicate that new, alternative approaches may be needed to decipher the genuine nature of the relationship between volatility and growth. Current literature does not generally consider the impact of the interactions among economic agents and their collective dynamics on long-term growth. It is this impact and its underlying mechanisms that we seek to capture and explain.

We are motivated by the micro-to-macro approach of agent-based modeling (LeBaron and Tesfatsion, 2008; Dawid and Delli Gatti, 2018; Hommes and LeBaron, 2018) and, especially, the Keynes-meets-Schumpeter class of models (Dosi et al., 2010, 2015) that study the linkages between endogenous growth and demand policy. While agent-based models successfully capture many complex phenomena, they are generally analytically intractable, making the analysis of the precise mechanics linking volatility and growth difficult. Our approach remains distinct as we aim to derive a tractable system of equations for the aggregate dynamics from micro-level interactions.

This chapter's objective is to develop a model of capital demand driven economic fluctuations, in which interactions among agents to coordinate on economic outcomes lead to periods of booms and busts, and apply it to examine how fluctuations affect the economy across different timescales and possibly shape its long-term growth. Inspired by Keynes (1936), our focus on capital demand is motivated by the observation that firms' investment is both pro-cyclical and volatile (Stock and Watson, 1999), suggesting investment decisions play a key role in business cycles. We treat investment decision-making as an interactions-based process whereby firm managers exchange views and affect each other's opinions. In other words, we emphasize strategic complementarity and peer influence that cause managers to coalign their individual expectations at the micro level. We use the framework developed in Gusev et al. (2015) and Kroujiline et al. (2016) to describe this interaction process mathematically and derive the macroscopic equations governing the dynamics of aggregate capital demand. To close the economy while highlighting the demand-driven effects, we attach these equations to a simple supply side component represented by the Solow growth model (1956).

As a result, we obtain a closed-form dynamical system, hereafter the Dynamic Solow model, which enables us to study a broad range of economic behaviors. The model's primary contribution is the identification of a new mechanism of business cycles that captures their quasiperiodic nature characterized by one or several peaks in a wide distribution of cycle lengths.

We show that, for economically realistic parameters, the Dynamic Solow model admits two attracting equilibria that entrap the economy in either a contraction or expansion.² The equilibria are indeterminate (Benhabib and Farmer, 1999) as both the path to and the choice of equilibrium depend on the beliefs of the agents themselves. The entrapment is asymmetric because technological progress, introduced externally, causes the econ-

²The 2008 crisis gave new impetus to revisiting the single equilibrium framework; e.g. Vines and Wills (2020) recently made the case for moving towards a multi-equilibrium paradigm.

omy to stay on average longer in expansion than contraction, contributing to long-term growth. The flow of exogenous news continually perturbs the economy stochastically and prevents it from settling at either equilibrium. Over time, the economy tends to drift slowly towards the boundary between the contraction and expansion regions, making it easier for a news shock to instigate a regime transition in line with the “small shock, large business cycle” effect (Bernanke et al., 1999). This endogenous mechanism generates quasiperiodic fluctuations as it involves both deterministic dynamics and stochastic forcing.

Such a mechanism, whereby noise applied to a dynamical system leads to a quasiperiodic response, is known as coherence resonance (Pikovsky and Kurths, 1997). It occurs in situations where the system has long unclosed trajectories such that even small amounts of noise can effectively reconnect them and thus create a quasiperiodic limit cycle. Coherence resonance emerges naturally in bi-stable systems³, including our model.

The coherence resonance mechanism differentiates the Dynamic Solow model from preceding research that has often considered limit cycles as the endogenous source of economic fluctuations.⁴ In particular, Beaudry et al. (2020) propose an extended Dynamic Stochastic General Equilibrium model, in which the quasiperiodic character of fluctuations comes from noise acting directly on a periodic limit cycle. We argue, however, that coherence resonance may be the preferred route to generating business cycles as it requires noise only as a catalyst, thus relying much less on random shocks to reproduce regime variability. Furthermore, we show that the fluctuations produced by a noise-perturbed limit cycle, which is as well recovered in a certain parameter range in our model, dampen long-term growth and unrealistically cause capital demand to diverge from supply in the long run.

We note that the Dynamic Solow model nests two limiting cases that match those of previous literature. In the case where capital demand is persistently higher than supply, the model recovers the exponential equilibrium growth of the classic Solow model. In the opposite case, where capital demand is persistently lower than supply, the model exhibits quasiperiodic fluctuations driven by a coherence resonance mechanism similar to that in Kroujiline et al. (2019).

We explore the Dynamic Solow model numerically across multiple timescales, from months to centuries, and identify business cycles as quasiperiodic fluctuations that most frequently last 40-70 years. These fluctuations may be associated with Kondratieff cycles if interpreted as investment driven.⁵ Korotayev and Tsirel (2010) employ spectral

³In dynamical systems, bistability means the system has two stable equilibrium states.

⁴The early literature comprises Hicks (1937), Kaldor (1940) and Goodwin (1951). Later reviews include Boldrin and Woodford (1990), Scheinkman (1990), Lorenz (1993) and Gandolfo (2009). We also note Beaudry et al. (2020) as an influential recent investigation.

⁵Kondratieff himself attributed these cycles to capital investment dynamics. This interpretation was further advanced by a number of papers in the 1980s. Kondratieff cycles are, however, more commonly linked to technological innovation. There have also been attempts to combine investment and innovation explanations. For a review see Korotayev and Tsirel (2010).

analysis to suggest the existence of long-term business cycles. However, the academic community remains divided on this issue and the research has been focused primarily on the fluctuations in the 8-12 year range. These shorter-term cycles cannot emerge in our model because it does not include accelerators such as the financial sector or household debt.

Currently, many macroeconomic models describe an economy in or near equilibrium. Most prominent is the Dynamic Stochastic General Equilibrium class of models (see Christiano et al., 2018; Kaplan and Violante, 2018, for recent reviews). While behavioral limitations and various frictions have been considered, these models operate in an *adiabatic* regime where equilibrium is reached more quickly than the environment changes. In other words, there is some form of perfect coordination (e.g. market clearing where supply and demand equate) among all agents at each point in time. Over long timescales this treatment may be justified, but in the near term coordination failures are inevitable, leading to pronounced fluctuations and persistent spells of disequilibrium.

The Dynamic Solow model enables us to study both the disequilibrium fluctuations and the equilibrium growth. We examine the impact of fluctuations on growth and show that fluctuations can affect economic expansion over extended time intervals. However, the deviations from the balanced growth path disappear with time as demand and supply converge asymptotically in the long run.

The remainder of this chapter is structured as follows. In Section 7.1 we introduce and explain the mechanics of dynamic capital demand and the Solow growth framework within which it rests. Section 7.2 considers two limiting cases: first, we obtain the equilibrium growth path when capital demand exceeds supply; and second, we investigate the demand dynamics and highlight the mechanism underlying fluctuations when capital supply exceeds demand. Section 7.3 formulates and studies the general case of the Dynamic Solow model, focusing on the analysis of mid-term fluctuations and long-term growth. Finally, Section 7.4 concludes by reflecting on the work done and suggests further avenues of research.

7.1 The Dynamic Solow Model

This section develops the Dynamic Solow model.⁶ The modeling framework is set out in Section 7.1 and the equations of the model components are derived in Sections 7.1-7.1.⁷

⁶M. Gusev and D. Kroujiline formulated and presented the basic ideas behind the model at the Macroeconomic Instability seminars of the "Rebuilding Macroeconomics" project (www.rebuildingmacroeconomics.ac.uk) run by the National Institute of Economic and Social Research in London (Gusev and Kroujiline, 2020).

⁷The code for this model is written in Python and is available at github.com/KarlNaumann/DynamicSolowModel.

Model Structure

The Dynamic Solow model is illustrated in Figure 7.1. It consists of a dynamic demand framework that we propose to describe how firms determine capital needs and make investment decisions (right loop), to which we attach the familiar circular income flow of the Solow growth economy (left loop).⁸

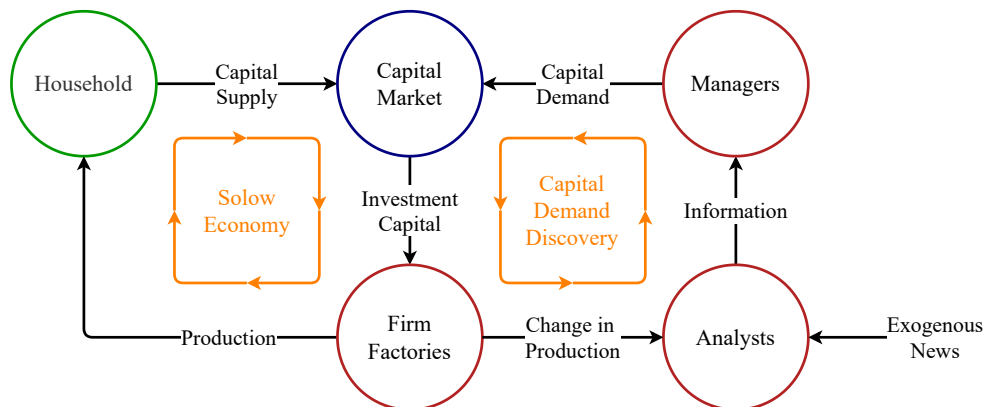


Figure 7.1 – A conceptual flowchart of the Dynamic Solow model. Each individual circle depicts an entity or agent. The circles’ color indicates whether they belong to the same entity (notably, firm managers and firm factories are both parts of the firm). Labeled black arrows define the flows between the respective entities or agents. The orange loops highlight (i) the Solow economy (left loop) and (ii) the dynamic demand decision-making process introduced in this section (right loop).

In the Solow economy, households supply capital through savings and firms convert capital into output for household consumption. Both households and firms are static decision-makers. Households save a fixed share of income and firms convert all supplied capital into production. In contrast, we aim to describe how firms develop a strategic business outlook based on their reading of the current economic situation and accordingly determine their capital needs so as to adjust production capacity. Firms thus become active decision-makers, which results in a dynamically evolving capital demand.

Organizational decision-making is a complex process with competing goals and targets, often based on industry-standard business planning and operating procedures (Cyert and March, 1992; Miller and Cardinal, 1994). Without needing to make firm goals explicit, we posit that corporate decision-making can be viewed as a composite of two distinct processes occurring on different timescales. First, there is information gathering and analysis, characterized by the frequency with which exogenous information such as

⁸We choose this supply-side framework for the following reasons: (i) the capital supply dynamics are less important on the timescales where we expect to find fluctuations and thus can be modeled approximately; (ii) the assumption that households save a constant fraction of income is an appropriate leading-order approximation since it is the first term in the Taylor series expansion of savings as a general function of income; and (iii) the Solow model is a parsimonious representation of economic growth, sharing the basics with many macroeconomic models, which may be helpful to extending our approach to more sophisticated settings.

ad-hoc company news, monthly statistics releases or quarterly earnings reports becomes available. Second, there is the formation of firms' expectations about the future based on the analysis of collected information, which is then translated into investment decisions. Note that the strategic aspect of investment decision-making implies longer timescales than those of information gathering and analysis.

We model this two-tier decision-making on the microscale by introducing two classes of agents: analysts who collect and analyze relevant information and managers who use this analysis to develop a business outlook and make investment decisions. There are industries where these two classes of agents actually exist (e.g. analysts and investors in finance), whereas in other situations this division serves as a metaphor for the different actions completed by an individual participant. Our objective is to derive the macro-level equations for aggregate demand from this micro setting.

External information enters the decision-making process at the analyst level. Initially, we may neglect the cost side and focus solely on revenue generation, elevating in relevance the expectation of future consumption. Motivated by recent work on extrapolative beliefs in finance (Greenwood and Shleifer, 2014; Kuchler and Zafar, 2019; Da et al., 2021), we assume that analysts base their expectations primarily on the current state of the economy by extrapolating the consumption growth into the future. As such, we wish to carve out consumption growth as the most relevant information stream and model all other news as exogenous noise (treating news shocks similarly to Angeletos and La'O, 2013; Angeletos et al., 2018; Beaudry and Portier, 2014). We can further simplify by replacing consumption with production since consumption is approximated as a constant fraction of production in the model. The resulting system acquires a feedback mechanism as higher output growth leads to increasing expectations that cause greater investment, inducing further increases in output growth and starting the process anew.

On the manager level, we emphasize the impact of the opinions and actions of competitors on decision-making, following the growing body of research on peer influence in business (Griskevicius et al., 2008) and strategic complementarity (Cooper and John, 1988; Beaudry et al., 2020). More specifically, we assume that managers exchange views within their peer network with the purpose of coaligning their expectations about the economy.

The Dynamic Solow model employs, as discussed, two different processes for capital demand and supply: firms determine capital needs dynamically via individual interactions and economic feedback while households supply capital in proportion to income. Thus, demand and supply do not necessarily match at each point in time, which brings us to the discussion of capital market clearing on different timescales. The dynamic demand discovery process occurs on timescales much shorter than the timescale of technological growth. At these short and intermediate timescales – relevant to information gathering, investment decision-making and production adjustment – prices are rigid and we expect demand and supply to behave inelastically. However, over long time horizons in which the economy is advancing along the equilibrium growth path, prices become flexible and

the capital market clears via price adjustment. Therefore, we expect that demand and supply converge in the long run.

As such, the conceptual framework behind the model is now complete. The remainder of Section 7.1 is as follows. First, the Cobb-Douglas aggregate production equation is extended to include the shorter timescales at which production capacity adjusts. Thereafter, the representative household and the capital motion equation are introduced. Thereafter, we derive the equations for aggregate capital demand from the micro-level agent-based formulation outlined above. Finally, the conditions for capital market clearing are set out.

Production

We represent aggregate output by a Cobb-Douglas production function that takes invested capital as an input,⁹ generically written as

$$Y = e^{\varepsilon t} K^\rho, \quad (7.1)$$

with output Y , invested capital K , capital share in production ρ and technology growth rate ε . Equation (7.1) implies that output adjusts immediately to any change in capital. In other words, it is only valid on timescales longer than the time it takes to adjust the production capacity (e.g. the construction of a new factory or installation of new machinery). Since we are also concerned with decision-making processes that occur at much shorter timescales than production adjustment, we introduce a dynamic form of production

$$\tau_y \dot{Y} = -Y + e^{\varepsilon t} K^\rho, \quad (7.2)$$

where the dot denotes the derivative with respect to time and $1 \ll \tau_y \ll 1/\varepsilon$ is the characteristic timescale of production capacity adjustment.¹⁰ In the short run, this equation describes the dynamic adjustment of output to new capital levels. In the long run, we recover the Cobb-Douglas production form (7.1) as $\tau_y \dot{Y}$ becomes negligibly small for $t \gg \tau_y$.

Finally, we rewrite equation (7.2) with log variables $k = \ln K$ and $y = \ln Y$ as

$$\tau_y \dot{y} = e^{\rho k + \varepsilon t - y} - 1. \quad (7.3)$$

Households and Capital Supply

We consider a single representative household that is the owner of the firm and thus receives Y as income. A fixed proportion of income, expressed as λY , is saved and the remainder is consumed. This is a convenient simplification that allows us to focus on

⁹For simplicity, we normalize the initial technology level to unity and do not consider labor, which is equivalent to normalizing the labor component to unity or taking the variables in per capita terms.

¹⁰In reality, the adjustment periods are asymmetric because it is easier to reduce capacity than to increase it. For simplicity, we treat capacity increases and decreases as if they were symmetric. This implies that the downturns we observe may be elongated.

the effects of dynamic capital demand. A constant savings rate can also be viewed as a leading-order Taylor expansion of household savings as a general function of income, making it a sensible first approximation.

The total savings are available to firms to invest. We denote them as capital supply K_s . The working capital used in production, K , suffers depreciation at a rate δ . As households are the owners of the capital, the loss δK is attributed to the capital supply. Consequently, the supply dynamics take the form

$$\dot{K}_s = \lambda Y - \delta K. \quad (7.4)$$

Setting $k_s = \ln K_s$, we reformulate equation (7.4) using log variables as

$$\dot{k}_s = \lambda e^{y-k_s} - \delta e^{k-k_s}. \quad (7.5)$$

Dynamic Capital Demand

In this section, we derive the equations for aggregate capital demand. As set out in Section 7.1, this derivation is based on a micro-level framework that divides the firms' investment planning into two processes occurring at different speeds: fast-paced information gathering and analysis; and slow-paced decision-making. We model these processes with two classes of interacting agents: analysts who collect and analyze relevant information; and managers who use this analysis to develop their strategic business outlook and make investment decisions.¹¹

In mathematical terms, we consider two large groups of agents: analysts $i \in \{1, \dots, N_h\}$ and managers $j \in \{1, \dots, N_s\}$, where $N_h \gg 1$ and $N_s \gg 1$. Each analyst and manager has a positive or negative expectation about the future path of production, respectively $h_i = \pm 1$ and $s_j = \pm 1$. The agents interact by exchanging opinions. As a result, the agents influence each other's expectations and tend to coalign them. To stay general, we assume analysts and managers interact among themselves and with each other. These individual interactions drive the evolution of the macroscopic variables: average analyst expectation h (information) and average manager expectation s (sentiment).

At each moment of time t , sentiment s is given by

$$s(t) = n_+(t) - n_-(t), \quad (7.6)$$

where $n_+ = N_s^+/N_s$ and $n_- = N_s^-/N_s$, with N_s^+ and N_s^- representing the respective number of optimists ($s_j = 1$) and pessimists ($s_j = -1$). By construction, s varies between -1 and 1 . At the leading order, we treat interaction as though each s_j is affected by the collective opinions s and h (similarly constructed), each forcing s_j in their respective

¹¹This modeling approach adapts the investor-analyst interaction framework, developed for the stock market in Gusev et al. (2015) and Kroujiline et al. (2016), to the macroeconomic context.

directions.¹² As a result of this simplification, we can introduce the total force of peer influence F_s acting on each manager as

$$F_s(s, h) = \beta_1 s(t) + \beta_2 h(t) + E_s(t), \quad (7.7)$$

where $\beta_1 > 0$ and $\beta_2 > 0$ are the sensitivities and E_s denotes general exogenous influences (to be specified later). Equation (7.7) implies that as the collective expectations of managers and analysts grow more optimistic, the stronger the force exerted on a pessimistic manager to reverse her views (and vice versa).

In addition, managers may be affected by a multitude of idiosyncratic factors causing them to occasionally change opinions irrespective of other participants. We treat them as random disturbances and, accordingly, introduce the transition rates p^{-+} as the probability per unit time for a manager to switch from a negative to positive opinion and p^{+-} as the probability per unit time of the opposite change. We can express the changes in n_+ and n_- over a time interval Δt as

$$n_+(t + \Delta t) = n_+(t) + \Delta t (n_-(t)p^{-+}(t) - n_+(t)p^{+-}(t)), \quad (7.8)$$

$$n_-(t + \Delta t) = n_-(t) + \Delta t (n_+(t)p^{+-}(t) - n_-(t)p^{-+}(t)). \quad (7.9)$$

Noting that $n_+ = (1 + s)/2$ and $n_- = (1 - s)/2$, we subtract (7.9) from (7.8) to obtain in the limit $\Delta t \rightarrow 0$

$$\dot{s} = (1 - s)p^{-+} - (1 + s)p^{+-}. \quad (7.10)$$

To complete the derivation, we must find out how the transition rates depend on peer influence: $p^{-+} = p^{-+}(F_s)$ and $p^{+-} = p^{+-}(F_s)$. It follows from (7.8) that in the state of equilibrium, when $n_{\pm}(t + \Delta t) = n_{\pm}(t)$, the condition $p^{-+}/p^{+-} = n_+/n_- = N_s^+/N_s^-$ holds. Thus p^{-+}/p^{+-} can be interpreted as the ratio of optimists to pessimists. We can assume this ratio changes proportionally to a change in F_s , that is $d(N_s^+/N_s^-)/(N_s^+/N_s^-) = \alpha dF_s$ where α is a positive constant. This interpretation allows us to write $d(p^{-+}/p^{+-})/(p^{-+}/p^{+-}) = \alpha dF_s$, which leads to

$$\frac{p^{-+}}{p^{+-}} = e^{\alpha F_s}. \quad (7.11)$$

Condition (7.11) implies correctly that $p^{-+} > p^{+-}$ for $F_s > 0$, $p^{-+} = p^{+-}$ for $F_s = 0$ and $p^{-+} < p^{+-}$ for $F_s < 0$. To obtain the final condition required to determine p^{-+} and p^{+-} uniquely, we introduce the characteristic time τ_s over which individual expectations change due to random disturbances. Since p^{-+} and p^{+-} are per unit time, τ_s represents the characteristic time over which the total probability for a manager to reverse her expectation is unity:¹³

$$(p^{-+} + p^{+-})\tau_s = 1. \quad (7.12)$$

¹²This treatment, known as the mean-field approach, is the leading-order approximation for a general interaction topology.

¹³Consider the impact of a random disturbance. At its end, the agent's state will remain or change to its opposite. Introduce p'^{++} and p'^{+-} as the probabilities for the agent, in state +1, to end up in

Together conditions (7.11) and (7.12) imply the transition rates:

$$p^{-+} = \frac{1}{\tau_s (1 + e^{-\alpha F_s})}, \quad p^{+-} = \frac{1}{\tau_s (1 + e^{\alpha F_s})}. \quad (7.13)$$

Equations (7.13) allow us to rewrite (7.10) as

$$\tau_s \dot{s} = -s + \tanh(F_s) = -s + \tanh(\beta_1 s + \beta_2 h + E_s), \quad (7.14)$$

where $\alpha/2$ is absorbed into β_1 and β_2 without loss of generality. Note that τ_s acquires a dual meaning: at the micro level, τ_s is akin to the manager's average memory timespan; at the macro level, τ_s is the characteristic time of variation in the aggregate expectation of managers.

Applying this approach to model the dynamics of analyst expectations yields the same form of the evolution equation for information h :

$$\tau_h \dot{h} = -h + \tanh(F_h) = -h + \tanh(\beta_3 s + \beta_4 h + E_h), \quad (7.15)$$

where τ_h represents the analyst's average memory timespan on the micro level and the characteristic time of the variation in the aggregate expectation of analysts on the macro level. Similarly, F_h is the peer influence acting on the analysts' expectations, which is linear in s and h with sensitivities β_3 and β_4 , and E_h denotes general exogenous influences.

Equations (7.14) and (7.15) describe a generalized interactions-based process of decision-making. We now make several assumptions to adapt it to the capital demand discovery mechanism of the Dynamic Solow model (Figure 7.1).

First, we assume managers receive information only via analysts and accordingly set $E_s = 0$. Second, we assume analysts are affected, first and foremost, by the news about economic development and only thereafter by all other news. More specifically, we assume the average analyst projects the output trend forward in time (extrapolative beliefs) and we treat all other relevant news as exogenous noise. Thus we set

$$E_h = \gamma \dot{y} + \xi_t, \quad (7.16)$$

with sensitivity γ and news noise ξ_t acting on the timescale $\tau_\xi \ll \tau_h$. The latter implies that changes to expectations are impacted by short-term shocks with no relation to economic fundamentals (as suggested, for example, by Angeletos et al. (2020)).

Third, we establish separate timescales for information processing and expectation formation. That is, we assume information is received and processed much faster than it

state +1 or -1, respectively, and p'^{-} and p'^{+} for the agent, in state -1, to end up in state -1 or +1, respectively. Thus, $p'^{++} + p'^{+-} = 1$ and $p'^{-+} + p'^{--} = 1$. Assuming that the ending state does not depend on the initial state, i.e. $p'^{++} = p'^{-+}$ and $p'^{-} = p'^{+-}$, we obtain $p'^{-+} + p'^{+-} = 1$. If the disturbances are frequent on the timescale of interest, we can transform the discrete probabilities into continuous transition rates via $p^{-+} = p'^{-+}/\tau_s$ and $p^{+-} = p'^{+-}/\tau_s$ to recover equation (7.12).

takes managers to adapt their long-term outlook and form investment decisions. Therefore: $\tau_h \ll \tau_s$. Fourth, as τ_h is much shorter than τ_s , we assume direct interactions are less important for analysts than for managers and we take $\beta_3 = \beta_4 = 0$ for simplicity.

The final step is to model the link between sentiment and capital demand. Consider a firm whose managers have just decided on capital allocation in line with their collective sentiment. The following day, all else being equal, the managers will not revisit this decision unless their sentiment changes. Therefore, in the short run where $t \ll \tau_s$ (that is, over time horizons where the memory of past sentiment persists), capital demand must be driven by change in sentiment. Conversely, over longer horizons where $t \gg \tau_s$, the connection between previous decisions and sentiment becomes weaker and, therefore, investment decisions must be based on the level of sentiment itself in the long run. For lack of simpler alternatives, we superpose these two asymptotic regimes, $\dot{k}_d \sim \dot{s}$ for $t \ll \tau_s$ and $\dot{k}_d \sim s$ for $t \gg \tau_s$, and, as a result, arrive at a complete system of equations for capital demand:

$$\dot{k}_d = c_1 \dot{s} + c_2 s, \quad (7.17)$$

$$\tau_s \dot{s} = -s + \tanh(\beta_1 s + \beta_2 h), \quad (7.18)$$

$$\tau_h \dot{h} = -h + \tanh(\gamma \dot{y} + \xi_t), \quad (7.19)$$

where $c_1 > 0$ and $c_2 > 0$ represent the capital demand sensitivity to a change in sentiment \dot{s} and the level of sentiment s , respectively; and $\gamma > 0$ represents the sensitivity of information h to the state of the economy or, in other words, the strength of economic feedback.¹⁴

Capital Market Clearing

At the relatively short time horizons relevant to information gathering, investment decision-making and production adjustment, prices are not flexible enough to efficiently match capital demand k_d and supply k_s , which are determined independently from each other. Accordingly, we introduce an inelastic market clearing condition for log invested capital k as

$$k = \min(k_s, k_d), \quad (7.20)$$

to be satisfied at each moment in time. In contrast to the classic framework, in which all household savings are used in production, this condition implies that only a portion of

¹⁴Gusev et al. (2015) derived equations (7.18) and (7.19) in the context of their generalized Ising model. The derivation presented here provides a more intuitive, albeit less rigorous, treatment. Equation (7.18), with h as an exogenous variable, was originally obtained by Suzuki and Kubo (1968) for the classic Ising (1925) model in statistical mechanics. Following Haag and Weidlich (1983), equation (7.18) (often in its stationary form) has frequently appeared in the socioeconomic context. For reviews of the interactions-based approaches reliant on the statistical mechanics methods and, particularly, the applications of the Ising model and its variants to the opinion dynamics problems in economics and finance, we recommend Brock and Durlauf (2001), Bouchaud (2013) and Slanina (2013). We also note that, unlike equations (7.18) and (7.19) that are derived from the micro level, equation (7.17) is obtained as a phenomenological relation between business sentiment and capital demand. Equation (7.17) was initially suggested as a link between investor sentiment and stock market returns in Gusev et al. (2015).

savings will be invested should demand fall short of supply (with the remainder retained in household savings accounts).

Equation (7.20) is a local clearing condition that reflects the short-term price rigidity. Therefore, as was discussed in Section 7.1, this equation cannot remain valid over long-term horizons during which prices become sufficiently flexible to match demand and supply. As such, we supplement (7.20) with an asymptotic clearing condition that holds in the timescale of long-term economic growth:¹⁵

$$k_s \sim k_d \quad \text{for } t \geq O(1/\varepsilon) \gg 1. \quad (7.21)$$

Together, equations (7.20) and (7.21) interlink the supply and demand components and close the Dynamic Solow model.

At this point, it may be useful to discuss the characteristic timescales in the model. The timescales we have encountered are differentiated in length such that $\tau_\xi \ll \tau_h \ll \tau_s \ll \tau_y \ll 1/\varepsilon$. Economically, information gathering occurs on a relatively short timescale, τ_h (with the publication of, for example, monthly and quarterly corporate reports and industry data releases); investment decisions require more time, τ_s (as processed through, for example, annual board meetings); and the implementation of changes to production levels takes much longer, τ_y (the time needed for material adjustments such as infrastructure development). We set $\tau_h = 25$, $\tau_s = 250$ and $\tau_y = 1000$ in units of business days (250 business days = 1 year). We further assume the timespan of exogenous news events to be on average one week and set $\tau_\xi = 5$ as the news noise decorrelation timescale (see Appendix F.1). Finally, we take technology growth rate $\varepsilon = 2.5 \times 10^{-5}$, which implies the timescale of 160 years.¹⁶

7.2 Demand and Supply: Two Limiting Cases

In this section, we inspect two cases that follow from the market clearing condition (7.20): first, the supply-driven case, $k_d > k_s$ such that $k = k_s$, which recovers a Solow-type growth economy; and, second, the demand-driven case, $k_d < k_s$ such that $k = k_d$, in which the economic fluctuations emerge.

Supply-Driven Case $k_d > k_s$

In the supply-driven case, the market clearing condition yields $K = K_s$ (firms use all available capital for production)¹⁷. Consequently, the Dynamic Solow model is reduced to equations (7.2) and (7.4) which can be expressed as a single second-order differential

¹⁵The asymptotic relation (7.21) means that the relative error between k_s and k_d goes to zero as t goes to infinity.

¹⁶In equation (7.1), the term $\exp(\varepsilon t)$, commonly referred to as total factor productivity, yields a growth rate of about 0.6% p.a. for $\varepsilon = 2.5 \times 10^{-5}$, which is not far from the estimates based on the 2005-2016 period using Fernald (2016).

¹⁷It is convenient to use the non-logarithmic variables for this analysis.

equation:

$$\tau_Y \ddot{K} + (1 + \tau_Y \delta) \dot{K} + \delta K = \lambda K^\rho e^{\varepsilon t}. \quad (7.22)$$

For $t \sim 1/\varepsilon$ and longer time intervals, the derivative terms in equation (7.22) become negligibly small and we recover the equilibrium growth path. On shorter timescales, $t \sim \tau_y$, equation (7.22) describes adjustment towards the equilibrium growth path. These two effects can be observed simultaneously by deriving an approximate solution to equation (7.22) for $t \geq O(\tau_y)$ (see Appendix F.3). The resulting production path is given by

$$Y = \left(\frac{\lambda}{\delta}\right)^{\frac{\rho}{1-\rho}} \left(\left(B e^{-\left(\frac{1-\rho}{\tau_y}\right)t} + 1 \right)^{\frac{1}{1-\rho}} + e^{\left(\frac{\varepsilon}{1-\rho}\right)t} - 1 \right), \quad (7.23)$$

where B is the constant of integration.¹⁸ Equation (7.23) explains the output dynamics between intermediate and long-term timescales, capturing both the long-term growth of the classic Solow model (given by the second exponent) and the intermediate relaxation towards the same (given by the first exponent). The approximate analytic solution (7.23) and the exact numerical solution to equation (7.22) are compared in Figure 7.2.




Figure 7.2 – Output $Y(t)$ in the supply-driven case represented by the numerical solution of equation (7.22) (dashed blue line) and the approximate solution (7.23) (solid red line). The precision of the approximate solution is improved with a greater timescale separation $\tau_y \ll 1/\varepsilon$. Parameters: $\rho = 1/3$, $\tau_y = 1000$, $\lambda = 0.15$, $\varepsilon = 10^{-5}$ and $\delta = 0.02$ with an integration constant of $B = 1.5$. The inset box highlights the intermediate adjustment of output from an arbitrary initial value to the equilibrium growth path. Taken from Naumann-Woleske et al. (2022)

Demand-Driven Case $k_d < k_s$

In the demand-driven case, the market clearing condition yields $k = k_d$. The Dynamic Solow model is specified at this limit by equations (7.3) and (7.17)-(7.19) (in this case,

¹⁸This derivation is valid for the parameter values provided in Appendix F.2, subject to the simplifying assumption $\tau_y \delta \gg 1$ which allows us to obtain the solution in a compact form.

equation 7.5 decouples and no longer affects production). To facilitate our analysis, we introduce the variable $z = \rho k_d + \varepsilon t - y$, which makes the model solutions bounded in the (s, h, z) -space (see Appendix F.4). Economically, z represents the direction and strength of economic growth. This follows from rewriting equation (7.3) as $\tau_y \dot{y} = e^z - 1$, noting that for $z > 0$ production expands, for $z < 0$ it contracts and $z = 0$ is a production fixed point. Using z , we re-express the model as a three-dimensional dynamical system that is bounded and autonomous in the absence of exogenous noise ξ_t :

$$\dot{z} = \rho c_1 \dot{s} + \rho c_2 s - \omega_Y (e^z - 1) + \varepsilon \quad (7.24a)$$

$$\tau_s \dot{s} = -s + \tanh(\beta_1 s + \beta_2 h) \quad (7.24b)$$

$$\tau_h \dot{h} = -h + \tanh(\gamma \omega_y (e^z - 1) + \xi_t), \quad (7.24c)$$

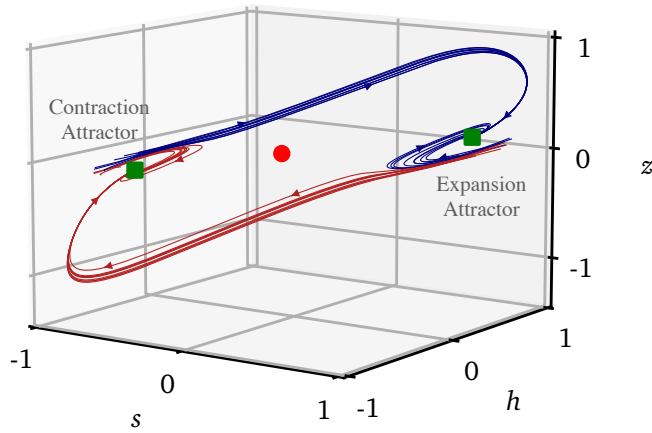
where, for convenience, $\omega_y = 1/\tau_y$.

This dynamical system is parametrized and examined in detail in Appendix F.1. For the relevant range of parameters it has three equilibria: a stable focus where sentiment is positive ($s > 0$) and the economy is expanding ($z > 0$), a stable focus where sentiment is negative ($s < 0$) and the economy is contracting ($z < 0$) and an unstable saddle point in between.¹⁹ The location, basin of attraction and stability of the equilibria are primarily affected by the parameters c_2 (sensitivity to sentiment levels) and γ (sensitivity to economic feedback). In particular, an increasing c_2 strengthens convergence towards the equilibria, so the system acquires greater stability.


If c_2 is below a certain critical value, equations (7.24) generate a periodic limit cycle. The idea that limit cycles provide a mechanism of economic fluctuations dates back to Kalecki (1937), Kaldor (1940), Hicks (1950) and Goodwin (1951). The empirical irrelevance of periodic limit cycles led to a diminished interest in this research direction²⁰; however, Beaudry et al. (2020) have reinitiated the discussion by proposing that cyclicalities can arise from stochastic limit cycles “wherein the system is buffeted by exogenous shocks, but where the deterministic part of the system admits a limit cycle”. In our system, exogenous news noise ξ_t similarly detunes limit cycle periodicity. This mechanism, however, cannot explain the “small shock, large crisis” effect or reproduce the general variability present in the real-world economy. At the other extreme, our system generates noise-prevailing behaviors with weak cyclicalities. Neither extreme accurately reflects empirical observations and thus we seek a sensible balance between these features in a parameter regime that produces significant dynamic effects but precedes the limit cycle formation (Appendix F.1).

¹⁹For convenience, we classify 3D equilibrium points using more familiar 2D terminology. As such: (i) the stable (unstable) node has three negative (positive) real eigenvalues; (ii) the focus has one real and two complex eigenvalues and is stable if the real eigenvalue and the real parts of complex eigenvalues are all negative and unstable otherwise; and (iii) the saddle is always unstable as it has three real eigenvalues that do not all have the same sign. In the figures, the stable points are green and unstable points are red, while the nodes are marked by triangles, foci by squares and saddles by circles.

²⁰Nevertheless, a similar line of research has been pursued in overlapping generations models and innovation cycles. See Hommes (2013) and Beaudry et al. (2020) for references.



(a) Coherence resonance for $c_2 = 7 \times 10^{-4}$ and $\gamma = 2000$. This subcritical regime presents a bi-stable configuration of equilibria: green squares denote the two stable foci and the red circle an unstable saddle. Red trajectories terminate at the $s < 0$ focus in which the economy contracts and blue trajectories terminate at the $s > 0$ focus in which the economy expands. The long trajectories passing near one focus and ending at the other are of a particular interest as they provide the pathway for the economy's regime transitions. Taken from Naumann-Woleske et al. (2022)



(b) Limit cycle for $c_2 = 1 \times 10^{-4}$ and $\gamma = 4000$. In this supercritical regime, only the positive ($s > 0$) equilibrium point survives, having bifurcated into an unstable focus, and a large stable limit cycle emerges that propels the economy between contraction and expansion with a constant frequency. Taken from Naumann-Woleske et al. (2022)

Figure 7.3 – 3D phase portraits ($\xi_t = 0$) in the (s, h, z) -space: (a) the coherence resonance regime and (b) the limit cycle regime. The long trajectories in (a) can be viewed as segments of the limit cycle in (b), which remain unconnected for $\xi_t = 0$. Parameters other than c_2 and γ are from the base case in Appendix F.2. Taken from Naumann-Woleske et al. (2022)




Figure 7.4 – Left: The coherence resonance phase portrait in Figure 7.3a projected on the s, z -plane. The approximate boundary between the contraction and expansion regions is indicated by a dashed black line. When the economy moves along the long trajectories, it traverses quickly the distance in s between the foci. However, the movement speed declines sharply with the proximity to equilibrium. As a result, the economy slowly ascends along the z -axis towards the left focus in the contraction mode and slowly descends to the right focus in the expansion mode. As the economy is undergoing this gradual drift, the distance to the boundary separating contraction and expansion diminishes. Right: The limit cycle phase portrait in Figure 7.3b projected on the s, z -plane. Taken from Naumann-Woleske et al. (2022)

To this end, we consider a subcritical regime with c_2 above but close to its critical value at which the limit cycle emerges. In this situation, henceforth referred to as the coherence resonance regime, the foci are always stable, thus acting as attractors entrapping the economy. In Figures 7.3 and 7.4, we compare the phase portraits (i.e. the system trajectories for $\xi_t = 0$) of the coherence resonance and limit cycle regimes. In the coherence resonance case, we take note of the unclosed largescale trajectories that pass near one attractor and converge to the other. These trajectories, which can be viewed as segments of a limit cycle, are the pathways along which the economy moves between contraction and expansion.

The dynamics of business cycles for nonzero ξ_t are visualized in Figure 7.5. The economy's trajectory displays distinctly bi-stable behavior as it spends most of its time near each focus and transits swiftly between them. When captive to an attractor, the trajectory follows an orbit around the corresponding focus, buffeted by exogenous noise ξ_t , preventing it from settling. Simultaneously, the economy drifts slowly towards the boundary between attracting regions (Figure 7.4(left)), making it easier for a random news shock to thrust it across the boundary to be caught by the other attractor. The news shocks ξ_t thus fulfill a dual purpose: they perturb the economy from equilibrium and provide a trigger that alternates the economic regime between expansions and recessions.

This mechanism can be classified as coherence resonance, a phenomenon whereby noise applied to a dynamical system leads to a quasiperiodic response (Pikovsky and Kurths, 1997). Coherence resonance normally occurs in bi-stable systems that are stochastically forced and in which key variables evolve on different timescales. The Dynamic Solow




Figure 7.5 – A simulated economy’s path for nonzero ξ_t in the coherence resonance case, smoothed by a Fourier filter to remove harmonics with periods less than 500 business days for a better visualization in the 3D phase space. The regions where the trajectory is dense indicate the contraction and expansion attractors, around which the economy spends most of its time. The relatively straight path segments between the attractors correspond to the economic regime transitions that occur on a relatively rapid timescale. Parameters are from the base case in Appendix F.2. Taken from Naumann-Woleske et al. (2022)




Figure 7.6 – Sentiment $s(t)$ (left) and output $y(t)$ (right) on a simulated economy’s path for nonzero ξ_t in the coherence resonance regime. Output undergoes fluctuations within a general growth trend as sentiment evolution exhibits a distinct bi-stable pattern. Parameters are from the base case in Appendix F.2. Taken from Naumann-Woleske et al. (2022)

model satisfies these requirements: (i) news shocks provide a stochastic force; (ii) two stable equilibria emerge in the relevant parameter range; and (iii) the separation of characteristic timescales follows from the dynamics of corporate decision-making processes. The three-dimensionality of equations (7.24) introduces an important novel feature into the classic two-dimensional case of coherence resonance: the above-mentioned slow drift of the economy’s trajectory, which gradually increases the probability of regime transition.²¹ This novel feature nonetheless leaves the basic mechanism unchanged: exogenous noise forces the economy across the boundary separating the regions of different dynamics, effectively reconnecting the trajectories between attractors. As a result, the economy undergoes quasiperiodic fluctuations consisting of alternating periods of expansion and recession punctuated by sharp transitions (as in Figure 7.6).

While both coherence resonance and limit cycle occur in the economically realistic range of parameters (see Appendix F.1), we will show that only coherence resonance is compatible with the long-term economic growth when studying the asymptotic convergence of capital supply and demand in the presence of fluctuations driven by the coherence resonance and limit cycle mechanisms in Section 7.3.

7.3 Business Cycles and Long-Term Growth in the General Case

While the supply- and demand-driven cases have been instructive for highlighting the mechanisms underlying economic dynamics, their applicability as standalone models is limited as supply and demand converge in the long run (equation (7.21)). As such, our primary focus is on the general case in which supply and demand coevolve, potentially leading to an interplay of supply- and demand-driven dynamics. We formulate the general case in Section 7.3, study long-term growth rates in Section 7.3 and examine economic fluctuations in Section 7.3.

Formulation of the General Case

In the general case, invested capital k can alternate between k_d (demand-driven regime) and k_s (supply-driven regime) in accordance with the market clearing condition (7.20). As discussed in Section 7.1, firms’ decision-making processes are influenced by feedback from the economy. However, the supply-driven regime represents a special situation in which firms’ investment decisions do not affect economic output as production is determined in this case solely by capital availability. In other words, the supply-driven regime implies a Solow-type growth economy propelled by expectations of future consumption so high as to induce firms to utilize all capital supplied by households in production. Therefore, \dot{y} , which is positive in this regime, holds no additional information for managers, who are already overwhelmingly bullish about the economy. The idiosyncratic news ξ_t remains the only source of nontrivial information, thereby becoming the focus of

²¹This drift imposes a slow timescale $t \geq O(\tau_y)$ on the frequency of the economy’s fluctuations by modulating the probability of regime transition. This effect was first observed and explained for a similar dynamical system in Kroujiline et al. (2019).

managers and analysts alike. Thus, economic feedback $\gamma\dot{y}$ vanishes as a decision factor in the supply-driven regime.

Following this argument, we account for regime-dependent variation in feedback strength by introducing a regime-specific factor $H(k_d, k_s)$ that regulates the impact of feedback in equation (7.19):

$$\tau_h \dot{h} = -h + \tanh(\gamma\dot{y}H(k_s, k_d) + \xi_t), \quad (7.25)$$

where

$$H(k_s, k_d) = \begin{cases} 1 & \text{if } k_d \leq k_s \\ 0 & \text{if } k_d > k_s \end{cases}. \quad (7.26)$$

The Dynamic Solow model is then represented in the general case by the following system of equations:

$$\tau_Y \dot{y} = e^{\rho k + \varepsilon t - y} - 1, \quad (7.27)$$

$$\dot{k}_s = \lambda e^{y - k_s} - \delta e^{k - k_s}, \quad (7.28)$$

$$\dot{k}_d = c_1 \dot{s} + c_2 s, \quad (7.29)$$

$$\tau_s \dot{s} = -s + \tanh(\beta_1 s + \beta_2 h), \quad (7.30)$$

$$\tau_h \dot{h} = -h + \tanh(\gamma\dot{y}H(k_s, k_d) + \xi_t), \quad (7.31)$$

$$k = \min(k_d, k_s), \quad (7.32)$$

$$k_s \sim k_d \quad \text{for } t \geq O(1/\varepsilon) \gg 1, \quad (7.33)$$

where (7.27) is the dynamic equation governing production; (7.28) describes the motion of capital supply; (7.29)-(7.31) govern the feedback-driven dynamics that link information h , sentiment s and capital demand k_d ; (7.32) defines the locally-inelastic market clearing condition; and (7.33) represents long-term market clearing that takes the form of an asymptotic boundary condition at large t .

Growth and Convergence in the Long Run

The Dynamic Solow model (7.27)-(7.33) covers two regimes with different dynamics: a demand-driven regime with endogenous fluctuations and a supply-driven regime without them. Both regimes are expected to participate in the model's general case, owing to the convergence of supply and demand in the long run under equation (7.33).

Equation (7.33) is central to our present analysis. Based on the regime definitions, this equation is satisfied when supply grows faster than demand in the supply-driven regime and, conversely, when demand grows faster than supply in the demand-driven regime. Under the demand-driven regime, the two possible mechanisms of fluctuations – limit cycle and coherence resonance – may entail different growth rates, validating the mechanism if demand grows fast enough to satisfy (7.33) and invalidating it otherwise.

7.3. BUSINESS CYCLES AND LONG-TERM GROWTH IN THE GENERAL CASE




Figure 7.7 – A long-term simulation that captures asymptotic growth in the enforced supply-driven regime ($k = k_s$). Output y and supply k_s grow steadily at the Solow rate R while demand k_d stagnates. As such, k_s will eventually reach and exceed k_d , at which point the supply-driven regime will be succeeded by the demand-driven regime. Parameters are from the base case in Appendix F.2.

This section aims to determine (i) the impact of fluctuations on growth; (ii) the mechanism of fluctuations compatible with equation (7.33); and (iii) the actual growth dynamics realized in the model. We first consider separately the supply- and demand-driven regimes (Sections 7.3 and 7.3) and then tackle the general case (Section 7.3). Appendix F.4 provides the derivations of the equations herein.

Asymptotic Growth in the Supply-Driven Case ($k_d > k_s$)

We show in Appendix F.4 that the economy's long-term growth in the supply-driven case is given by

$$y_0 = k_{s0} = \frac{\varepsilon}{1 - \rho} \equiv R, \quad (7.34)$$

$$k_{d0} = 0, \quad (7.35)$$

where y_0 , k_{s0} and k_{d0} represent, respectively, the log output, log supply and log demand growth rates; $\rho = 1/3$ is the capital share in production; and R denotes the classic Solow growth rate. As expected, the growth rate is not influenced by demand dynamics and matches R . These estimates are verified by numerical simulations (see Figure 7.7). Note that supply always catches up with demand as $k_{s0} > k_{d0}$ in this case.

Asymptotic Growth in the Demand-Driven Case ($k_d < k_s$)

We show in Appendix F.4 that the economy's long-term growth in the demand-driven case satisfies

$$y_0 = k_{s0} = R + \rho(k_{d0} - R). \quad (7.36)$$

According to this equation, output becomes dependent on demand, which itself undergoes endogenous fluctuations, and thus demand dynamics affect the long-term behavior of the economy. The magnitude of k_{d0} determines whether the economy expands faster

or slower than the classic Solow economy: $y_0 > R$ if $k_{d0} > R$ and $y_0 < R$ if $k_{d0} < R$ (the latter condition including an important case when $k_{d0} = 0$ that yields an especially slow growth rate $y_0 = k_{s0} = \varepsilon$). Next, we estimate k_{d0} numerically under the effect of limit cycle and coherence resonance.

Figure 7.8 depicts the growth dynamics driven by a periodic limit cycle ($\xi_t = 0$). We observe that k_{d0} stays close to zero and y_0 and k_{s0} match ε closely in accordance with (7.36), meaning the economy grows only through improvements in production efficiency. Figure 7.9 displays similar dynamics for the limit cycle perturbed by exogenous noise ξ_t . It follows that limit cycles, whether periodic or stochastic, lead to a growth rate of less than R .

The above result can be explained by noting that an economy on a limit cycle trajectory spends roughly an equal amount of time in expansion ($s > 0$) as in contraction ($s < 0$) and, consequently, s exhibits on this trajectory a long-term average value of zero. In Appendix F.4, we find that k_{d0} is proportional to the long-term average of s , implying k_{d0} tends to zero as well; therefore, demand can never catch up with supply due to the difference in their growth rates. In sum, the fluctuations generated by a limit cycle detract from long-term growth and fail to satisfy equation (7.33).

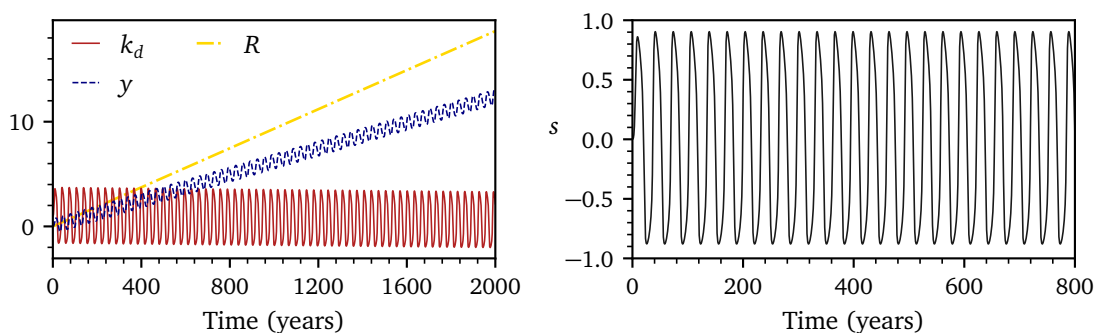


Figure 7.8 – A long-term simulation that captures asymptotic growth in the enforced demand-driven regime ($k = k_d$) powered by a periodic limit cycle ($\xi_t = 0$) with $\gamma = 1000$, $c_2 = 2 \times 10^{-5}$ and all other parameters from the base case in Appendix F.2. Left: Production y grows at a rate lower than the Solow rate R while demand k_d stagnates (and, in fact, appears to gradually decrease, which could be attributed to the slight asymmetry of the limit cycle with respect to s). Since k_s and y grow at the same rate (equation (7.36)), k_d cannot catch up with k_s . Right: Sentiment $s(t)$ demonstrates the limit cycle's periodicity.

Coherence resonance induces a drastically different long-term dynamic despite the visually similar fluctuations (see Figure 7.10). Demand grows asymptotically at $k_{d0} > R$, leading to accelerated economic growth of $y_0 = k_{s0} \equiv R^* > R$ in accordance with equation (7.36). This fast-paced growth, made possible by excess capital, $k_s > k_d$, available for investment in the demand-driven regime, is explained by technological progress ($\varepsilon > 0$), causing the economy to spend on average more time in expansion than contraction. We further observe that $k_{d0} > R^*$; that is, demand grows faster than both supply




Figure 7.9 – A long-term simulation that captures asymptotic growth in the enforced demand-driven regime ($k = k_d$) powered by a stochastic limit cycle ($\xi_t \neq 0$) with the same parameters as in Figure 7.8. Left: Production y grows at a rate lower than the Solow rate R while capital demand k_d stagnates (exogenous noise evidently erasing the limit cycle’s asymmetry visible in Figure 7.8). Since k_s and y grow at the same rate (equation (7.36)), k_d cannot catch up with k_s . Right: Sentiment $s(t)$ is no longer periodic due to the impact of ξ_t .

and output.²² Therefore, demand powered by coherence resonance always catches up with supply.




Figure 7.10 – A long-term simulation that captures asymptotic growth in the enforced demand-driven regime ($k = k_d$) powered by coherence resonance in the base case parameter regime (Appendix F.2). Left: Output y and supply k_s expand at the rate $R^* > R$ and demand k_d grows faster than both y and k_s . As such, k_d eventually reaches and exceeds k_s , at which point the demand-driven regime is succeeded by the supply-driven regime. Right: Sentiment $s(t)$ exhibits a bi-stable behavior typical of coherence resonance.

Asymptotic Growth in the General Case

We have shown that fluctuations affect growth in the demand-driven regime of the Dynamic Solow model. In particular, limit cycles generate fluctuations that contribute negatively to growth, thus failing to satisfy the asymptotic boundary condition (7.33).

²²This result is consistent with equation (7.36), which can be rewritten as $k_{d0} - R^* = ((1 - \rho)/\rho)(R^* - R)$, so that $k_{d0} > R^*$ since $R^* > R$ and $\rho < 1$.

Therefore, such fluctuations cannot be realized, which rules out limit cycles as the mechanism responsible for business cycles.




Figure 7.11 – A long-term simulation that captures asymptotic growth in the general case. Left: Output y , supply k_s and demand k_d grow at the Solow rate R , demonstrating the asymptotic convergence on the economy’s trajectory. Right: The interplay of supply- and demand-driven regimes on a subsection* of the same trajectory. Shaded segments correspond to periods when $k_d > k_s$. Parameters are from the base case in Appendix F.2.

By contrast, coherence resonance produces fluctuations that contribute positively to growth, so that demand always catches up with supply. As this occurs, the system transits into the supply-driven regime in which supply grows faster than demand. Once supply has exceeded demand, the system switches back into the demand-driven dynamics. The regime cycle has thus come full circle, ensuring (7.33) is satisfied in the long run. As such, the economy’s path realized in the general case is forged by a regime interplay where the supply-driven equilibrium dynamics and the demand-driven fluctuations, powered by coherence resonance, continuously succeed one another.

Our simulations show the economy grows asymptotically at the Solow rate R . This result is not entirely unexpected. As capital supply and demand converge over the long run, capital invested into production during the supply- and demand-driven regime segments of the economy’s trajectory must also match asymptotically, as follows from (7.32). Consequently, the economy’s average growth rate across supply-driven segments is equal to the average growth rate across demand-driven segments. As the economy expands at R in the supply-driven regime, the same growth rate is achieved, on average, across the demand-driven segments,²³ meaning R is also the overall rate of expansion. Figure 7.11 displays a simulation capturing the realized asymptotic growth path in the general case and highlights the interplay of the supply- and demand-driven dynamics.

To sum up, the asymptotic growth rates in the demand-driven regime depend on the mechanism underlying economic fluctuations. Fluctuations driven by a limit cycle cannot be realized since they do not satisfy the convergence between supply and demand in the long run. The economy’s trajectory realized in the general case of the Dynamic Solow model consists of a chain of supply-driven regimes in which the economy experiences

²³The demand-driven economy cannot reach the asymptotic growth rate $R^* > R$ as a result of the segments’ finite duration and an adverse growth bias due to the typical alignment of demand-driven segments with recessionary periods.

7.3. BUSINESS CYCLES AND LONG-TERM GROWTH IN THE GENERAL CASE

the equilibrium growth and demand-driven regimes in which fluctuations emerge via coherence resonance. Overall, the economy grows asymptotically at the classic Solow rate R . Although fluctuations can cause large excursions from this equilibrium growth path, the deviations disappear on the large timescales relevant for the convergence of supply and demand.

Business Cycle Dynamics

Our analysis of asymptotic growth in the preceding section has led us to conclude that coherence resonance is the relevant endogenous mechanism underlying economic dynamics as it enables the convergence of capital demand and supply over the long run. In this section, we focus on the intermediate timescale to examine endogenous fluctuations produced by the Dynamic Solow model (7.27)-(7.33) in the coherence resonance regime.




Figure 7.12 – A simulation of the economy’s trajectory over the medium term in the model’s general case, which highlights a sequence of supply- and demand-driven regimes. Shaded segments correspond to periods during which $k_d > k_s$. Parameters are from the base case in Appendix F.2

Figure 7.12 depicts a typical realization of the economy’s trajectory over the medium term. The economy undergoes a sequence of supply- and demand-driven dynamic behaviors, as indicated, respectively, by shaded and unshaded segments. In the demand-driven case, in which demand is below supply, sentiment (lower panel) exhibits distinctively bi-stable behavior, staying for long periods near the positive (expansion) and negative (con-

traction) equilibria and traversing quickly the distance between them during economic regime transitions. This sentiment behavior leads to fluctuations in demand (middle panel) that, in turn, induce business cycles around the long-term growth trend (upper panel). Conversely, during periods when supply is the limiting factor, sentiment follows a random walk due to the absence of economic feedback and the supply-driven economy exhibits the equilibrium growth dynamics.

The long-term simulations demonstrate that demand stays below supply on average $\sim 70\%$ of the time. This can be interpreted as the firms' decision to hold excess capital (as, for example, noted in Fair, 2020) as the entire capital supply is made available to firms, implying a capital utilization rate below 100% over extended periods.²⁴




Figure 7.13 – Histogram of the duration of simulated business cycles. The cycles are based on the detrended production, $y - Rt$, with respect to the best fit straight line which coincides with the equilibrium growth path, given by the Solow growth rate R (Section 7.3). Duration is calculated as the time interval between two successive zero crossings in the same direction by detrended production. The histogram, based on the 5 year bins, is truncated at 10 years to eliminate noise artifacts and at 150 years to highlight its peak. The parameters are from the base case in Appendix F.2.

Figure 7.13 is a histogram of business cycle periods simulated by the model. It displays a wide distribution with a peak in the 40-70 year interval (with over 50% of the periods falling into this range), indicating the presence of quasiperiodic fluctuations. To confirm the source of these fluctuations, we inspect the distribution of the lengths of sentiment cycles, defined as the roundtrip of sentiment between the positive and negative equilibria (such as those depicted in the lower panel in Figure 7.12). This distribution, shown in Figure 7.14, also peaks at 40-70 years. It follows that business cycles are, as expected, linked to sentiment transitions from one equilibrium to the other driven by coherence resonance. Therefore, we affirm coherence resonance is the relevant mechanism forming the quasiperiodic fluctuations in output captured in Figure 7.13.

²⁴The utilization rate measures current production as a proportion of potential production. In our context, the analogous measure is the ratio of capital in production in the demand-driven regime ($k = k_d$) to capital available (k_s). For empirical evidence of firms holding excess capital, see the FRED series “Capacity Utilization: Total Index” (CAPUTLB50001SQ) or the Census Bureau’s National Emergency Utilization Rate. See also Murphy (2017) for a study of the persistence of excess capital.




Figure 7.14 – Histogram of the duration of simulated sentiment cycles. The cycles are defined as the roundtrip of sentiment between the positive and negative equilibria. Duration is calculated as the time interval between two successive zero crossings in the same direction by sentiment s in the enforced demand-driven case. The histogram, based on the 5 year bins, is truncated at 10 years to eliminate noise artifacts and at 150 years to highlight its peak. The parameters are from the base case in Appendix F.2.

In Appendix F.1, we show that parameter c_2 , which defines the sensitivity of capital demand to sentiment, is key to the business cycle duration: the lower c_2 , the shorter the average duration of business cycles. We also show there that the model admits coherence resonance only if c_2 is above a certain critical value and tune the model to be in a regime with c_2 close to this value. It follows that coherence resonance – as a mechanism of business cycles driven by firms’ investment – imposes a natural minimum duration threshold, ruling out fluctuations with a characteristic timespan shorter than the Kondratieff-like 40-70 years.

In current literature, business cycles are typically estimated to last 8-12 years. However, a direct comparison of the duration would be misleading as our model, centered on capital demand dynamics, does not include links to the faster-paced processes, such as credit or equity market dynamics, that can accelerate business cycles through further interactions with the real economy. In other words, our model captures capital demand driven cycles, which are arguably just one of a number of fluctuation modes that reinforce or otherwise affect each other to produce the business cycles observed in the real world.

On that point, we take note of Kroujiline et al. (2019) that studies combined effects in a coupled macroeconomic system, attaching the interactions-based stock market model of Gusev et al. (2015) (capable of producing relatively short-term endogenous cycles) to the simple phenomenological model of the economy of Blanchard (1981) (within which output follows slow relaxation dynamics) to obtain quasiperiodic fluctuations with the same frequency as observed business cycles. A natural next step would be to investigate whether a more advanced coupled system, where both the financial sector and the real economy experience nonlinear endogenous dynamics at different frequencies can replicate and explain observed macroeconomic behaviors in greater detail.²⁵

²⁵Such as the stock market model of Gusev et al. (2015) and the present Dynamic Solow model of the

7.4 Conclusion

In this chapter we have developed the Dynamic Solow model, a tractable macroeconomic model that captures dynamic behaviors across multiple timescales, and applied it to study economic fluctuations and their impact on long-term growth.

Our model consists of a dynamic capital demand component, representing an interactions-based process whereby firms determine capital needs and make investment decisions, and a simple capital supply component in the form of a Solow growth economy. These components are interlinked via a capital market, which comprises a local inelastic market clearing condition reflecting short-term price rigidity and an asymptotic market clearing condition valid for timescales in which prices are sufficiently flexible to match supply and demand. Starting from the micro-level interactions among firms, we derived the macroscopic equations for capital demand that constitute the dynamic core of the model and attached them to a capital motion equation for a static representative household (providing a dynamic version of the Cobb-Douglas production equation to capture short-term processes) and the aforementioned market clearing conditions. As a result, we have obtained a closed-form nonlinear dynamical system that allows for the examination of a broad range of economic behaviors.

The Dynamic Solow model admits two characteristic regimes, depending on whether capital demand exceeds supply or not. When demand exceeds supply, supply drives output and the dynamic demand component decouples from the rest of the economy, placing the economy on the familiar equilibrium growth path. Otherwise, demand drives output and the model is shown, for economically realistic parameters, to possess two attracting equilibria, one where the economy contracts and the other where it expands. This bi-stable geometry gives rise to business cycles manifested as endogenous fluctuations, wherein the economy's long entrapment in recessions and expansions is punctuated by rapid alternations between them. We show that, in our model, the economy's realized trajectory is forged by an interplay of these regimes such that the supply-driven equilibrium dynamics and demand-driven fluctuations continuously succeed one another. We further show that the economy spends around 70% of its time in the demand-driven regime, indicating fluctuations represent a prevalent economic behavior.

We identify a coherence resonance phenomenon, whereby noise applied to a dynamical system leads to a quasiperiodic response, to be the mechanism behind demand-driven fluctuations. In our model, exogenous noise (representing news received by analysts) instigates the economy's transition from one equilibrium to the other, resulting in recurrent booms and busts. As such, news shocks act as a catalyst, which is compatible with the "small shocks, large cycle" effect observed in the real-world economy. In addition, under a different range of parameter values, we obtain a stochastic limit cycle (i.e. a limit cycle perturbed by exogenous noise) likewise capable of generating endogenous fluctuations. We show, however, that this type of fluctuations cannot be realized as

economy, which share a similar framework for micro-level interactions.

the growth dynamics induced by it do not allow supply and demand to converge in the long run. While both limit cycle and coherence resonance mechanisms are hardwired in our model, in the sense that the parameter ranges must be appropriately selected, we conjecture that in reality the economy self-regulates towards the coherence resonance parameter ranges via long-term price adjustment responsible for the convergence of supply and demand in the long run.

The distribution of the business cycle periods simulated by our model displays a peak in the Kondratieff range of 40-70 years, demonstrating the quasiperiodic character of demand-driven fluctuations. We further find coherence resonance imposes a minimum duration threshold that rules out fluctuations peaking at shorter lengths. This result seems sensible because our model, centered on capital demand dynamics, has no links to faster-paced processes (such as credit or equity market dynamics) that can accelerate fluctuations to be in line with the observed business cycles. A natural extension would be to develop a coupled system, within which both the financial sector representing such faster-paced processes and the real economy experience nonlinear endogenous dynamics at different characteristic frequencies.

Our simulations show that although demand-driven fluctuations occasionally cause large excursions from the equilibrium growth path, the deviations vanish in the long run as supply and demand converge. In our model, the equilibrium growth path is defined by the Solow growth rate in which technology growth appears, simplistically, as a fixed exogenous parameter. From this perspective, endogenizing technological progress in the model may help better understand the long-term impact of fluctuations on growth, presenting an intriguing topic for future research.

Key Messages

- This chapter has developed the Dynamic Solow Model: adding an interactions-based investment decision process into a simple Solow framework.
- The model has two regimes: a supply-driven dynamic recovering classic exponential growth, and a demand-driven regime where endogenous quasi-periodic crises occur.
- The mechanism of crises is coherence resonance: small noise in the decision-making process can push the investment outlook from converging to a positive to a negative sentiment.
- In the simple model presented, the distribution of business-cycle durations is fat-tailed with a peak of 40-70 years, as there are no additional accelerators such as credit constraints.

Economic Crises in a Model with Capital Scarcity and Self-Reflexive Confidence

Taken from Morelli et al. (2021) *Economic Crises in a Model with Capital Scarcity and Self-Reflexive Confidence* without modification. This work is a collaboration with F. Morelli, M. Tarzia and J-P. Bouchaud. My main contributions in this paper include the economic framing of sentiment in New Keynesian and DSGE models, the construction and implementation of the extended model, and its ensuing exploration. The drafting, analysis and review was done in collaboration with all of the authors.

The years following 2008 were marked by the great financial crisis, and with it a crisis for economic theory (Kirman, 2010). As for the great depression of the 1930s, there was a failure to predict the crisis among economic orthodoxy.¹ Despite its failures in predicting the recession (Christiano et al., 2018) or the sluggish recovery (Lindé et al., 2016), the mainstream Dynamic Stochastic General Equilibrium (DSGE) class of models have remained the core macroeconomic framework and workhorse tool of policy (Kaplan and Violante, 2018). While calls to reform these models have been made (Stiglitz, 2018; Vines and Wills, 2018, 2020), the basic framework with a single rational representative agent often remains a baseline assumption when studying business cycles, although heterogeneous agents new Keynesian models (HANK) have recently been considered as well.² The decisions of such a representative agent, which include capital investment decisions, determine the trajectory of the economy and are based on the optimisation of a utility function with static parameters, with no space for “animal spirits” or confidence

¹See the extensive discussions in (Buiter, 2009; Blanchard, 2018; Dosi and Roventini, 2019; Korinek, 2017; Romer, 2016; Stiglitz, 2018), with a recent review in Fair (2020).

²For a non-exhaustive list of prominent examples see (Kaplan et al., 2018; Kaplan and Violante, 2014; McKay and Reis, 2016).

effects (although see Barsky and Sims (2012) and the discussion below). In such models, “dark corners” are absent and crises can only be the result of major exogenous shocks.³

This must be contrasted with Keynes’ intuition, which, as rephrased by Minsky (1976), was that “the subjective evaluation of prospects over a time horizon is the major proximate basis for investment and portfolio decisions, and these subjective estimates are changeable”. Expectations can indeed be subject to rapid changes, disagreement and irrationality, as reflected in the high volatility of investment, and the abrupt nature of expansions and recessions (Stock and Watson, 1999). The investor behavior behind these swings are indeed often referred to as animal spirits or irrational exuberance (Shiller, 2005; Akerlof and Shiller, 2010). There is now a rich literature on irrational behavior across economics (see Hommes (2021) for a recent review). However, it has not been fully dovetailed into more traditional business cycle models. One can find some boundedly rational components in DSGE models,⁴ such as Cornea et al. (2019), Hommes and Lustenhouwer (2019), or Ozden (2021), who focus on learning in expectations formation in a single-actor model, as well as Jump and Levine (2019) and Gabaix (2020), who use various different utility specifications in DSGE. But apart from Barsky and Sims (2012), Deniz and Aslanoglu (2014), and Brenneisen (2020), there is surprisingly little work attempting to factor confidence or sentiment into the DSGE framework as an explicit variable. This is despite some empirical work suggesting that consumer confidence contains important information for forecasting personal spending and consumption (Blanchard, 1993; Carroll et al., 1994; Matsusaka and Sbordone, 1995).

The adapted New Keynesian models of Barsky and Sims (2012), and Brenneisen (2020) consider confidence as forecasting with imperfect signals, or private news about future technological states. In this chapter, we instead focus on Keynes’ animal spirits, and the non-rational self-reflexive facet of confidence, which can lead to abrupt shifts like in 1929 or 2008.⁵ As a first step to incorporate such effects and assess their impact on the economy, some of us recently proposed a generalization of a simple monetary model in which the household’s propensity to consume depends on the prior state of the economy, which generates either optimism or anxiety (Morelli et al., 2020). This feedback can amplify productivity shocks, and lead to the appearance of a second equilibrium characterized by low consumption and high unemployment, as well as crises resulting from self-induced confidence collapse. The existence of two very different macroeconomic equilibria has also been recently suggested in another context by Carlin and Soskice (2018).

In the present chapter, our aim is to significantly extend the work of Morelli et al. (2020) by including capital investment as a factor determining the trajectory of the

³“Dark corners” refers to a particularly insightful piece by O. Blanchard in 2014 entitled “Where Danger Lurks”, see Blanchard (2014).

⁴See Franke and Westerhoff (2017) for an early review of animal spirits in macroeconomic models.

⁵More recently Angeletos et al. (2018), and Angeletos and La’O (2013) consider sentiments as uncertainty about the beliefs of others. For another strand of the literature on sudden breakdown of confidence, see Bouchaud (2013) and Gama Batista et al. (2015) and references therein.




Figure 8.1 – Trajectories of the OECD confidence index, changes in consumption [%], and investment [%] in the United States over a period from the beginning of 2000 to the present. The data were taken from FRED and OECD. Taken from Morelli et al. (2021).

economy. We assume that capital and labor are essentially non-substitutable, and posit a behavioral rule for investment that accounts for both consumer confidence and for the quality of the returns generated by risky capital investment. This expanded framework allows us to investigate the joint dynamics of confidence, capital availability and output. In a nutshell, our model attempts to capture many of the ideas so clearly expressed by Keynes in the opening quote above, while keeping part of the scaffolding of standard business cycle models.

Our motivations for building such behavioral business cycle models stems from the Great Recession of 2008, sparked by Lehman’s bankruptcy, that led to a sudden collapse in the confidence of both households and investors. This was followed by an almost immediate downfall of both investment and consumption. These stylized facts can be observed in Figure 8.1. It took six years to recover to prior confidence levels, even as investment and consumption grew in the medium-term. It is difficult to believe that the Great Recession was the result of a major exogenous shock.⁶ Rather, Keynes’ story assigning the abruptness of the crisis to a shift in investment decisions may be more plausible. Indeed, anecdotal evidence reported by prominent actors at the time strongly suggests that confidence collapse played an essential role in the unfolding of the crisis – see the account of Bernanke et al. (2019).

The consumption and investment trajectories generated by our model can be grouped into four distinct categories differentiated by the prevalence of crises in consumption and the scarcity of capital for production. We recover the bi-stable behavior obtained in Morelli et al. (2020), alternating between enduring spells of high and low confidence. In

⁶Lo (2012) notes that there is no consensus narrative of the causes for the crisis. In addition, from a DSGE model perspective, Angeletos et al. (2020) recently ruled out many of the common exogenous DSGE shocks as explanatory candidates for business cycles.

addition, the household's investment behavior can lead to capital scarcity, i.e. periods where capital is the limiting factor to output. During these instances there is an increased risk of a confidence collapse and an ensuing low-consumption depression where the household consumes a small fraction of disposable income and invests cautiously. When compared with the results of our previous version of the model where capital is absent (or rather, assumed to be so abundant that keeping track of it is unnecessary), we find that low output periods can last orders of magnitude longer. This is because, in the absence of suitable policy measures, investment remains low and capital scarcity prevents the economy from recovering.

This multiple equilibria scenario is an attempt to move away from the over-simplified, but still dominant single equilibrium paradigm, following recent calls to that effect Vines and Wills (2020); Greene (2021). Note that the coexistence of different equilibria is also the hallmark of recent agent based models, such as Gualdi et al. (2015) or Sharma et al. (2020).

The qualitative results of our behavioral business cycle model suggests various policy measures, in terms of narratives that may change the perception of the future of the economy and the attractiveness of investment in productive capital (Shiller, 2019). In particular, we emphasize the crucial need to maintain capital investment at a sufficiently high level throughout crisis periods, in order to allow for a quick recovery when the economic conditions improve.

The manuscript is organized as follows. In Section 8.1 we build up our business cycle model based on Morelli et al. (2020), and outline our two novel additions. We then show the various dynamics the model can generate and reveal its phase diagram in Section 8.2. We discuss in particular how capital scarcity increases the probability of consumption crises, and lead to a multi-fold increase of the recovery time (see Section 8.2). Section 8.3 concludes by discussing the policy implications of our findings, and the avenues for possible extensions of the model.

8.1 A Behavioral Business Cycle Model

The framework presented here hybridizes some standard assumptions used in the New Keynesian Dynamic Stochastic General Equilibrium (DSGE) model (see (Galí, 2015)) with plausible behavioral assumptions about consumption propensity and investment strategies. The environment is based on two blocks: the representative consumer and the representative firm. At this point, we neglect inter-temporal effects and do not attempt to model inflation dynamics and monetary policy, although these features could be included at a later stage. Nonetheless, the phenomenology of our model is already quite rich and needs to be streamlined before exploring further the dynamics of prices. All variables and notations are reported in Table G.1 of Appendix G.




Figure 8.2 – A schematic representation illustrating the division of income, i.e. the budget constraint, by the household. Taken from Morelli et al. (2021).

The Household

The household sector derives utility from a composite consumption good and provides labor services to the firm. At each time t , the representative household maximizes its instantaneous utility,

$$U_t(c_t, n_t) := G_t \cdot \log c_t - \gamma \cdot n_t^2, \quad (8.1)$$

where c_t and n_t denote respectively the level of aggregate consumption and the aggregate amount of working hours the household provides to the firm, G_t is the (time dependent) propensity to consume out of income, and γ is the disutility of labor (which we fix to 1 for the numerical analysis).

Each period, the household faces a budget constraint given by its real income \mathcal{I}_t ,

$$\mathcal{I}_t := w_t \cdot n_t + \frac{b_{t-1}}{1 + \pi_t} + q_{t-1} \cdot \frac{k_{t-1}}{1 + \pi_t}, \quad (8.2)$$

which is funded by three sources: (i) the real wage rate w_t paid by the firm for a unit of labor n_t , (ii) the real value of the maturing single-period bonds b_{t-1} , purchased at time $t - 1$ at the price $(1 + r_{t-1})^{-1}$ and paying $(1 + \pi_t)^{-1}$ at time t , where r_t is the interest rate and π_t is the inflation rate, and (iii) the realized yield q_{t-1} per unit of real capital k_t that the firm pays to the household in return for investment. We henceforth assume a constant interest rate $r = 0.15\%$ and inflation $\pi = 0.1\%$, keeping in mind a unit time scale corresponding to a month or quarter.

Total spending, correspondingly, consists in good consumption (with the price of good set to unity), purchases of new bonds and topping up the firm's capital. Maximization

of the household's utility (Eq. (8.1)) leads to the familiar state equation

$$n_t \cdot c_t - \frac{G_t \cdot w_t}{2\gamma} = 0, \quad (8.3)$$

describing the trade-off between consumption and labor in the current period t .

Interestingly, Eq. (8.3) can also be interpreted in a way that lends itself to a natural generalization for investment decisions. Suppose one starts with a *time independent* utility function, Eq. (8.1) with $G_t \equiv 1$, which is now optimized under the constraint that the total budget devoted to consumption is a *fixed* fraction $G_t \in [0, 1]$ of the income \mathcal{I}_t , i.e.

$$c_t = G_t \cdot \mathcal{I}_t. \quad (8.4)$$

It is easy to show that the very same equation Eq. (8.3) immediately follows. We posit that the remaining fraction $1 - G_t$ of income is invested in bonds and capital, i.e.

$$i_t = (1 - G_t) \cdot \mathcal{I}_t, \quad (8.5)$$

where a fraction $F_t \cdot i_t$ (with $F_t \in [0, 1]$) is allocated to productive capital, and the remainder $(1 - F_t) \cdot i_t$ is invested in bonds – see Fig. 8.2 for a pie chart summarizing the household spending and investment decision.

The capital level available to the firm thus evolves as

$$k_t = (1 - \delta) \cdot k_{t-1} + F_t \cdot (1 - G_t) \cdot \mathcal{I}_t, \quad (8.6)$$

where δ is the capital depreciation rate. The remaining investment is allocated to bonds at price $(1 + r)^{-1}$, so

$$\frac{b_t}{1 + r} = (1 - F_t) \cdot (1 - G_t) \cdot \mathcal{I}_t \quad (8.7)$$

The quantities G_t and F_t aim to capture confidence effects and the attractiveness of risky capital investment, respectively, and are specified in section 8.1 below.

The Firm

The economy's productive sector is made up of a single representative firm, which transforms labor n_t and capital k_t into a composite good y_t consumed by the representative household. The firm's production technology is given by a Constant Elasticity of Substitution (CES) function with constant returns to scale,⁷

$$y_t = z_t \cdot \left(\alpha \cdot k_t^{-\rho} + (1 - \alpha) \cdot n_t^{-\rho} \right)^{-\frac{1}{\rho}}, \quad (8.8)$$

⁷In full generality, the CES function should be written as $(\alpha \cdot k_t^{-\rho} + (1 - \alpha) \cdot (\kappa n_t)^{-\rho})^{-\frac{1}{\rho}}$, where κ is another parameter. However, one can always set $\kappa = 1$ at the expense of re-scaling the disutility of labor parameter according to $\gamma \rightarrow \kappa\gamma$.

where $\alpha = 1/3$ is the capital share in production, $1/(1+\rho)$ is the elasticity of substitution between capital and labor with $\rho > 0$, and $z_t > 0$ is a stationary exogenous technological process. It is given by $z_t = z_0 e^{\beta t}$, where \mathfrak{z}_t follows an AR(1) process:

$$\mathfrak{z}_t = \eta \cdot \mathfrak{z}_{t-1} + \sqrt{1 - \eta^2} \cdot \mathcal{N}(0, \sigma^2) , \quad (8.9)$$

with first-order autocorrelation η , which affects the correlation time of the technology shocks. (In the following we will fix $\eta = 0.5$, corresponding to a correlation time of a few months). The base level z_0 corresponds to the most probable value of productivity. Note, importantly, that z_0 has units of $[\text{Time}]^{-1}$, i.e. the amount of goods that can be produced *per unit time* for a given level of capital and labor. As our focus is on economic fluctuations, we abstract from production growth in the present model, i.e. the secular dependence of z_0 on time.

The CES production function nests two important limits that affect economic dynamics. As $\rho \rightarrow 0^+$, the production function becomes perfectly elastic and recovers the Cobb-Douglas form ($y_t^{CD} = z_t n_t^{1-\alpha} k_t^\alpha$), whereas in the limit $\rho \rightarrow +\infty$ the firm produces via an inelastic Leontief function ($y_t^L = z_t \min(n_t, k_t)$).⁸ Throughout the following, we choose $\rho = 7$, corresponding to a near Leontief limit, i.e. a very small amount of substitutability between capital and labor. We will briefly comment in section 8.2 the impact of higher substitutability.

The firm maximizes its target profit \mathcal{P}_t

$$\mathcal{P}_t = p_t \cdot y_t - w_t \cdot n_t - q_t^* \cdot k_t , \quad (p_t \equiv 1), \quad (8.10)$$

with respect to the labor supply n_t and the capital level k_t , where p_t is set to unity and correspondingly w_t is the real wage and q_t^* is the real rent on capital. Under the assumption that the market clears, i.e.

$$y_t = c_t , \quad (8.11)$$

one finds

$$\tilde{w}_t = (1 - \alpha) \left(\frac{\tilde{c}_t}{n_t} \right)^{1+\rho} \quad (8.12)$$

$$\tilde{q}_t^* = \alpha \left(\frac{\tilde{c}_t}{k_t} \right)^{1+\rho} \quad (8.13)$$

where, generically, $\tilde{x} := x/z$. Note that, as it should be, wages, consumption and rent on capital are all in units of z , i.e. unit time scale (e.g. a month or a quarter).

Combining the household's state equation, Eq. (8.3) and the equation for the real wage, Eq. (8.12), the consumption level c_t must satisfy

$$\tilde{c}_t^2 = \frac{G_t}{2\gamma} (1 - \alpha)^{\frac{2}{\rho}} \left[1 - \alpha \left(\frac{\tilde{c}_t}{k_t} \right)^\rho \right]^{1 + \frac{2}{\rho}} , \quad (8.14)$$

⁸Keeping the parameter κ free (see previous footnote), the Leontief function would read $y_t^L = z_t \min(\kappa n_t, k_t)$, i.e. κ^{-1} measures the amount of labor equivalent to one unit of capital.




Figure 8.3 – The figure shows the behavior of re-scaled consumption \tilde{c}_t , wages \tilde{w}_t , and rent on capital \tilde{q}_t^* as a function of k_t , in the Leontief limit, i.e. $\rho \rightarrow +\infty$. Note the qualitative change of behavior between $k_t > \sqrt{G_t/2\gamma}$ and $k_t < \sqrt{G_t/2\gamma}$. Taken from Morelli et al. (2021).

As both sides of Eq. (8.14) are monotonous, this ensures a unique solution for any given level of capital k_t and consumption rate G_t . As expected, the consumption at time t increases if the capital k_t is increased and/or the consumption rate G_t is increased.

The Leontief Limit

In this section we discuss in detail the Leontief limit of the equations derived in the previous section. Such an analysis will greatly help understanding the dynamics of the model that will be described below.

Abundant Capital

Assume first that $\tilde{c}_t < k_t$ and $\rho \rightarrow \infty$. Then $(\tilde{c}_t/k_t)^\rho \rightarrow 0$ and one finds

$$\tilde{c}_t \approx \sqrt{\frac{G_t}{2\gamma}} \quad (8.15)$$

This is only consistent with our working hypothesis when

$$\frac{G_t}{2\gamma} < k_t^2. \quad (8.16)$$

In this regime one finds, using Equation (8.3):

$$n_t = \tilde{c}_t \tilde{w}_t, \quad (8.17)$$

which once plugged back in Equation (8.12) leads to

$$\tilde{w}_t = (1 - \alpha)^{1/(2+\rho)} \approx 1. \quad (8.18)$$

Since $\tilde{c}_t < k_t$, one concludes from Eq. (8.13) that the rent on capital \tilde{q}_t^* is exponentially small. Intuitively, as labor is the limiting factor, consumption is directly proportional

to how much the household chooses to work and the productivity at that time, while capital has no distinct effects on the economy.

Scarce Capital

Now let us look at the regime

$$\frac{G_t}{2\gamma} > k_t^2. \quad (8.19)$$

We introduce the notation $\beta_t := 2k_t^2\gamma/G_t$ for further use. We hypothesise that the solution for \tilde{c}_t in this regime is of the form

$$\tilde{c}_t = k_t e^{-x_t/\rho} \quad (8.20)$$

where $x_t = O(1)$ is to be determined. Plugging in Eq. (8.14), we find:

$$k_t^2 e^{-2x_t/\rho} = \frac{G_t}{2\gamma} (1 - \alpha)^{-2/\rho} (1 - \alpha e^{-x_t})^{1+2/\rho}, \quad (8.21)$$

or, for $\rho \rightarrow +\infty$,

$$e^{-x_t} = \frac{1 - \beta_t}{\alpha}. \quad (8.22)$$

The equation for \tilde{q}_t^* then leads to

$$\tilde{q}_t^* = \alpha e^{-x_t} = 1 - \beta_t, \quad (8.23)$$

which indeed vanishes when $\beta_t = 1$, correctly matching the regime where capital is plentiful, whereas \tilde{q}_t^* tends to unity when $\beta_t \rightarrow 0$, i.e. where $k_t \rightarrow 0$.

With $\tilde{c}_t = k_t e^{-x_t/\rho}$, one finds from Equation (8.3)

$$n_t = \frac{G_t w_t}{2\gamma k_t} e^{x_t/\rho}. \quad (8.24)$$

Finally, plugging into the equation for wages,

$$\tilde{w}_t^{2+\rho} = (1 - \alpha) e^{-x_t} \beta_t^{1+\rho}, \quad (8.25)$$

or in the limit $\rho \rightarrow \infty$,

$$\tilde{w}_t = \beta_t + O(\rho^{-1}), \quad (8.26)$$

and hence $n_t = k_t + O(\rho^{-1})$. Again, this solution matches with the result $\tilde{w}_t = 1$ obtained for $\beta_t \geq 1$.

Discussion

A summary sketch of the above results is provided in Fig. 8.3, as available capital k_t is varied. Part of the dynamical properties of our model can be inferred from this figure: when capital is lush, return on capital is small and investment decreases (i.e. F_t decreases). If investment falls below the level of capital depreciation δ , then k_t will

fall until the level $\sqrt{G_t/2\gamma}$ is reached. At this point, return on capital q_t^* increases, promoting investment. When consumption propensity G_t increases, k_t may fall behind, again leading to an increase of q_t^* . Hence we expect a regime where the economy stabilizes close to the point where $k_t \approx \sqrt{G_t/2\gamma}$, where capital and labor are tracking each other, and interest on capital and wages are neither very small, nor saturated to their maximum value $w_{\max} = q_{\max}^* = z$.

The Risk of Investment

Rigidity and costs to capital usage are typically introduced through adjustment costs to capital utilization (e.g. see Smets and Wouters (2007)). In this chapter we take a different route. Rather than empowering the household to choose the firm's utilization rate, we suppose the household invests in capital and gives operational control of capital to the firm. In exchange it is promised a return q^* per unit capital and per unit time scale (month or quarter). However, as the volatility of the stock-market attests to, such a return is not assured. Hence, we introduce an intrinsic state-dependent risk to the returns on capital⁹ $\xi \in [0, 1]$ as a modifier, such that the rate actually paid by the firm is:

$$q_t = q_t^* \cdot \xi_t \leq q_t^*, \quad (8.27)$$

where ξ is distributed as

$$p(\xi) = a \cdot \xi^{a-1}, \quad (8.28)$$

where parameter a controls the intensity of the risk. Note indeed that $\mathbb{E}[\xi] = a/(1+a)$ and $\mathbb{V}[\xi] = a/(2+a)(1+a)^2$. Hence the larger the value of a , the more $p(\xi)$ is concentrated around $\xi = 1$ (full payment). This formulation of risk implies that the representative firm pays out at most the marginal productivity of capital, but more likely only pays a fraction of this, corresponding to an effective description of financial distress and bankruptcy within a representative firm setup. In most simulations we set $a = 15$, such that the return is on average 93.75% of the promised return. In an extended version of the model, the parameter a could itself be a function of the state of the economy (in particular of the availability of capital), but we will not consider this possibility here.

Spending and Investing

The model laid out in this section contains two dynamic variables, the consumption rate G_t in Eq. (8.4) and the investment allocation rate F_t in Eq. (8.5), which have not been specified yet. These two variables are responsible for the feedback mechanisms which are at the core of the dynamical evolution of our model economy. Here we elaborate on these mechanisms and provide the economic intuition behind them.

The Consumption Rate

⁹More sophisticated distributions can be considered. We use this simple form to keep the number of parameters of the model as small as possible.

As in Morelli et al. (2020) and Deniz and Aslanoglu (2014), we postulate that the consumption rate G_t (or propensity, see section 8.1) is a function of the *consumer confidence index* \mathcal{C}_t , that we model as a real variable $\in [-1, 1]$ and, possibly, on the difference between the expected inflation rate $\hat{\pi}_t := \mathbb{E}_t[\pi]$ and the bond rate r_t :

$$G_t := G_t(\mathcal{C}_t, \hat{\pi}_t - r_t, \dots), \quad (8.29)$$

where the dependence on the second variable is a way to effectively encode the content of the standard Euler equation without explicitly introducing an inter-temporal optimization of utility, and where the \dots leaves room to possible additional variables. But since in the present chapter we assume both inflation and interest rates to be constant, the second variable will be dropped altogether. As far as the first variable is concerned, we follow our previous work in Morelli et al. (2020), where we postulated that confidence of a given household is impacted by the level of consumption of *other* households in the previous time step. In a mean-field limit, this self-reflexive mechanism writes

$$\mathcal{C}_t = \tanh(\theta_c \cdot (c_{t-1} - c_0)), \quad (8.30)$$

The parameter c_0 is a “confidence threshold” where the concavity of $\mathcal{C}(c)$ changes (if $c_{t-1} < c_0$, \mathcal{C} is closer to 1 while if $c_{t-1} > c_0$, $\mathcal{C}(c)$ is closer to -1). Parameter $\theta_c > 0$ sets the width of the consumption interval over which the transition from low confidence to high confidence takes place. One could introduce, as in e.g. Beaudry and Portier (2014), Naumann-Woleske et al. (2022) or Beaudry et al. (2020), the impact of macroeconomic news as an extra contribution to the argument of the tanh function. This would describe how the consumer confidence index is further modulated by some exogenous shocks, but we leave such an extension for future work.

Back to the consumption rate G_t , we write

$$G_t = \frac{1}{2}[G_{\min} + G_{\max} + (G_{\max} - G_{\min}) \cdot \mathcal{C}_t] \quad (8.31)$$

where $0 \leq G_{\min} < G_{\max} \leq 1$ are the minimum and maximum proportions of income the household will consume. We fix $G_{\min} = 0.05$ to ensure the household will consume whenever its income is positive (necessary consumption). Similarly, we set $G_{\max} = 0.95$ to account for a minimal form of precautionary savings in response to some uncertainty regarding the future.

The intuition behind Eqs. (8.30) and (8.31) is that when consumption is above the threshold, $c > c_0$, there is high confidence in the future of the economy, thereby a large fraction of income, $G_t \rightarrow G_{\max}$, is consumed. High confidence represents the belief that future income will be sufficient to maintain high consumption with a minimal amount of savings to sustain capital levels. Conversely, when consumption is below the threshold, $c < c_0$, the consumption rate collapses, $G_t \rightarrow G_{\min}$. Following a shock or deterioration in consumption to below the confidence threshold, there is uncertainty about the future economy and whether future consumption is assured. This induces the household to save more for the future, effectively reducing the current demand. Economically, c_0 is

thus similar to minimal consumption: it defines the threshold beyond which there is a panic where the household’s “survival” is in question.

The parameter θ_c modulates the households reaction to a breach of necessary consumption, and can be described as the household’s panic polarity. For high θ_c , the household requires only a relatively small shock below c_0 to reduce the consumption rate to its minimum. This leads to a bi-stable savings behavior with sharp transitions. Conversely, as $\theta_c \rightarrow 0$ the household becomes unresponsive to the state of the economy, consuming half its income regardless of high or low preceding consumption. The intermediate levels of θ_c describe the smoothness of the adjustment to consumption shocks.

According to Morelli et al. (2020) there are four distinct “phases”, i.e. regions of qualitatively comparable dynamics, that are distinguished by the bi-stability of G_t . We can observe in particular a phase of high persistent consumption with no crises, high consumption with short downward spikes, or a phase with alternating periods of high consumption (booms) and low consumption (busts).

The Investment Allocation

In each period, the household must allocate its savings between one-period bonds and capital. It does so through an allocation decision F_t based on the household’s observation of the economy, and its beliefs about future risk and return. The novelty of our model lies in the behavioral foundation that determines the proportion of new investment dedicated to bonds, F_t .

Investment beliefs are shaped by two factors: (i) an estimate of the expected risk-adjusted excess returns to capital investment, given by a Sharpe ratio \mathcal{S}_t (Sharpe, 1966), and (ii) the current confidence level \mathcal{C}_t about the future state of the economy.

The Sharpe ratio \mathcal{S}_t is an estimate of the risk-adjusted real return, $q_t - \delta$, of investing capital in the firm versus holding risk-free bonds (b_t) paying r . It increases as the returns to capital increase or become less volatile. We assume that estimates of the future Sharpe ratio are only based on exponential moving averages of past (observable) realized returns, which is a form of extrapolative beliefs.¹⁰ In other words, i.e. the mean μ^q and standard deviation σ^q of the return stream are computed as

$$\mu_t^q = \lambda \cdot \mu_{t-1}^q + (1 - \lambda) \cdot q_t \quad (8.32)$$

$$(\sigma_t^q)^2 = \lambda \cdot (\sigma_{t-1}^q)^2 + (1 - \lambda) \cdot (q_t - \mu_t^q)^2 \quad (8.33)$$

$$\mathcal{S}_t := \mathcal{N} \cdot \frac{\mu_t^q - r_t - \delta}{\sigma_t^q} \quad (8.34)$$

with an exponential moving average defined by a gain parameter $\lambda \in (0, 1)$, corresponding to a memory time scale equal to $\mathcal{T}_\lambda := 1/|\log \lambda|$: a larger λ implies that a higher weight is given to recent observations. The factor $\mathcal{N} \approx 1/4$ is quite arbitrary, but chosen such that, when compared to the confidence in Eq. (8.35) below, the two terms are

¹⁰See Da et al. (2021), and Kuchler and Zafar (2019) for recent empirical work on extrapolative beliefs.

of similar magnitude. (Note that this choice is in fact immaterial, since changing \mathcal{N} is equivalent to re-scaling the parameter ν defined in Eq. (8.35) below.)

The interpretation of the Sharpe ratio is as follows: a positive signal $\mathcal{S}_t > 0$ suggests that the expected real return to capital investment exceeds the returns to risk-free bonds. The magnitude of \mathcal{S}_t is inversely proportional to the risk of capital investment, as measured by the estimated volatility σ_t^q . Thus in a high-volatility environment the signal might be positive but weak.

The second indicator potentially influencing the household investment decision is the confidence index, \mathcal{C}_t , as previously defined. In periods where the household has low confidence, there is a reduced impetus to invest in risky assets because households wish to guarantee next-period income. These are often periods of crisis with a higher volatility in returns. Since bonds are risk-free, this leads to a higher allocation of funds to bonds, *ceteris paribus*. Conversely, higher confidence about the future means more appetite for risk, and hence a higher fraction of the savings invested in the capital of firms and a lower fraction invested in bonds.

We postulate that the propensity, F_t , to make risky bets is a function of the overall *sentiment* Σ_t , computed as a linear combination of the Sharpe ratio and of the confidence:

$$\Sigma_t = \nu \cdot \mathcal{S}_t + (1 - \nu) \cdot \mathcal{C}_t, \quad (8.35)$$

where $\nu \in [0, 1]$ is the weight the household gives to its estimates of risk-adjusted return \mathcal{S}_t and its confidence level \mathcal{C}_t . When $\nu = 1$, the household's confidence plays no role in the investment rule. For positive Sharpe ratio and confidence indicator, the sentiment is positive, $\Sigma_t > 0$, indicating a willingness to invest in risky capital. But if $\nu < 1$ sentiment can turn negative even when the Sharpe ratio is high, because of a high level of anxiety about the future state of the economy, encoded as a negative value of \mathcal{C}_t .

Finally, the unbounded sentiment Σ_t is transformed into a portfolio allocation to capital $F_t \in [0, 1]$ via,

$$F_t = \frac{1}{2} [F_{\max} + F_{\min} + (F_{\max} - F_{\min}) \cdot \tanh(\theta_k \cdot \Sigma_t)], \quad (8.36)$$

where F_{\max} and F_{\min} represent the maximum and the minimum proportion of total investment i_t invested into capital. In the following, the allocation decision is bounded between $F_{\min} = 0$ and $F_{\max} = 1$, which precludes any divestment (or short-selling) of capital or bonds.¹¹ The parameter θ_k represents the sensitivity of the portfolio allocation to the agent's sentiment and sets the width of the sentiment interval over which the capital allocation goes from F_{\min} to F_{\max} , that is how *polar* the investment decision is. For $\theta_k \rightarrow \infty$, the allocation becomes binary, leading to either $F_t = F_{\min}$ when sentiment is negative or $F_t = F_{\max}$ when sentiment is positive.

¹¹One could allow for divestment by $F_{\min} < 0$, however, this would require a more elaborate form for Eq. (8.36).

In the following part, we fix the sensitivities to a rather high value $\theta_k = \theta_c = 15$, such that the transitions between different regimes are sharp.

Summary & Orders of Magnitude

In this section we have set up a business cycle model incorporating two behavioural mechanisms: a self-reflexive consumption rate decision, already advocated by Morelli et al. (2020), and an investment allocation decision. The novelty of this chapter lies in the behavioural foundation that determines the proportion of new risky investment F_t , which depends directly on three key parameters, λ, ν, θ_k , describing the “sentiment” of the household, i.e. its risk aversion. F_t also indirectly depends on the risk intensity parameter a and the capital depreciation rate δ . The consumption decision depends on two parameters c_0 and θ_c that define the household confidence about its future welfare.

In the following we discuss how the parameters of these two feedback mechanisms strongly affect the model’s dynamics. Note that a very important parameter of the model is the baseline productivity z_0 , which fixes the scale of the consumption, wages and rent on capital (all per unit time scale). In the following, we choose $z_0 = 0.05$, corresponding to an annual productivity of capital of 20% if the unit time step is a quarter and 60% if it is a month.¹²

Among all the parameters of the model, three have an interpretation in terms of time scales:

- η , which appears in the dynamics of the productivity shocks, that we have fixed to 0.5 throughout this study, corresponding to a time scale $\mathcal{T}_\eta = 1/|\log \eta|$ of a few months ;
- λ , which is the gain parameter used by investors to estimate the Sharpe ratio of risky investments, corresponding to a time scale $\mathcal{T}_\lambda = 1/|\log \lambda|$. Our default value will be $\lambda = 0.95$, corresponding to $\mathcal{T}_\lambda \approx 20$ or 5 years if the unit time is a quarter or a month respectively;
- δ , the capital depreciation rate, which we choose in the range 0.001 – 0.02, corresponding to a typical replacement time of capital $\mathcal{T}_\delta = 1/|\log(1 - \delta)| \approx 12 - 250$ years when the unit time is a quarter, and three times less if it is a month. Hence $\delta = 0.001$ means essentially no depreciation of capital.

The role and the effect of varying these timescales is studied in detail in Section 8.2. An important remark, at this stage is that, while our choice of one quarter as the unit time step is quite arbitrary, a combination of parameters that is crucial for the properties of the model is the dimensionless product $z_0 \cdot \mathcal{T}_\delta \approx z_0/\delta$, i.e. how many goods can be produced (per unit capital) over the life-cycle of capital.

¹²To estimate an appropriate order of magnitude for z_0 , we considered the gross value added by non-financial corporations in the U.S. divided by the current cost net stock of fixed assets together with total wages (as a proxy for labor), which shows a downward trend to approximately 28% p.a.




Figure 8.4 – This figure shows different sections of the phase space. The three rightmost upper panels display the \log_{10} probability of capital shortage Ξ_k as a function of the Sharpe ratio weight ν and of the confidence threshold c_0 . The bottom row display the \log_{10} probability of consumption crises Ξ_c as a function of the same parameters. The dark red zones correspond to the regions where $\Xi_{k/c} > 0.99$, where crises are permanent. Each column corresponds to a different choice of the capital depreciation, from $\delta = 0.001$ to $\delta = 0.02$. We set $a = 15$ and $\sigma_z = 0.15$. As δ increases, the $(\mathcal{H}_k, \mathcal{H}_c)$ regimes becomes widespread. In each panel we have marked the points chosen to illustrate the different phases of the model, together with their label (see text). The dynamics trajectories and the corresponding histograms are reported in Fig. 8.5. The two leftmost panels show another section of the phase space, varying $\log a$ and σ_z , with $c_0 = 0.017$, $\nu = 1$ and $\delta = 0.005$ fixed (the two solid dots there correspond to the same solid dots of the middle panels, in the $\mathcal{H}_k \mathcal{H}_c$ phase).

Taken from Morelli et al. (2021).

8.2 Crises & Phase Diagrams

In this section we first investigate numerically the phase diagram of our self-reflexive business cycle model and highlight the different dynamical features that the model can generate. We choose as control parameters those which govern the behavior of our two feedback mechanisms: the consumption propensity G_t and the risky investment decision F_t . In order to navigate through the following paragraphs, let us explain in a nutshell what is expected to happen in the model.

If a productivity shock causes confidence to drop, consumption propensity G_t and consumption both drop as well, whereas the saving rate $1 - G_t$ increases. Because consumption drops, unemployment rises and capital becomes superfluous, leading to a decrease of the rent on capital q^* . Because the fraction of savings invested in capital F_t depends both on q^* (through the Sharpe ratio) and on the level of confidence (with a weight $1 - \nu$), the amount invested in risky capital, given by $(1 - G_t) \cdot F_t \cdot \mathcal{I}_t$ can either increase (if the factor $1 - G_t$ dominates) or decrease (if the factor F_t dominates), depending on parameters and conditions. In the second situation, and if capital depreciation is fast, one may face a situation where consumption is impaired and capital becomes scarce at the same time, making recovery more difficult and leading to long periods where the economy is trapped in a low output state.

Crises Indicators

We focus on two distinct phenomena exhibited by our model: consumption crises and capital scarcity.

- Consumption crises occur in periods where the household's consumption, c_t , falls below its threshold, c_0 . In other words, we have a low demand for consumption which leads the economy into a stagnating low-output state. The severity of such consumption crises is measured as

$$\Xi_c = \frac{1}{T} \sum_{t=0}^T \left(1 - \frac{c_t}{c_0}\right) \Theta(c_0 - c_t), \quad (8.37)$$

where $\Theta(x \geq 0) = 1$ and $\Theta(x < 0) = 0$ and T is the total simulation time. This indicator counts the fraction of time consumption c_t is low, weighted by the relative distance between c_t and c_0 .

- Since we are considering an economy defined by low substitutability between capital and labor (i.e. $\rho \gg 1$ in the CES production function), we define capital scarcity as the periods where production is determined by capital levels, i.e. $k_t \leq n_t$. The severity of capital crises is similarly measured as

$$\Xi_k = \frac{1}{T} \sum_{t=0}^T \left(1 - \frac{k_t}{n_t}\right) \Theta(n_t - k_t). \quad (8.38)$$

In a sense, one can consider these two phenomena as demand and supply crises.

- In the consumption crisis state the household does not wish to spend on consumption, hence we see a low aggregate demand. Provided capital depreciation is low, this is also a state of excess capital ($k > n$) and low returns on capital.
- In the capital scarcity state, the firm is bound in its production by the supply of capital, hence it can be viewed as a form of supply crisis.

Both phenomena can be more or less frequent, and at first glance unrelated, but closer scrutiny reveals that in some regions of parameters, these two types of crises interact with one another. To differentiate between characteristic behaviors we distinguish between four different phases in the space of the parameters defined by the values that the indicators Ξ_k and Ξ_c take: $(\mathcal{L}_k, \mathcal{L}_c)$, $(\mathcal{L}_k, \mathcal{H}_c)$, $(\mathcal{H}_k, \mathcal{L}_c)$, $(\mathcal{H}_k, \mathcal{H}_c)$, where \mathcal{L} and \mathcal{H} represent the “low prevalence” and “high prevalence” of each phenomena c or k , respectively. There is however no strict definition of the boundary between high and low prevalence regimes. As a convention, we consider that the crisis prevalence is high when $\Xi \gtrsim 10^{-2}$.

Given this setup, we first focus on the effects of three key parameters: the depreciation rate δ , the weight ν of the Sharpe ratio in the investment decision, and the consumption threshold c_0 . Other parameters are fixed to $z_0 = 0.05$, $\lambda = 0.95$ (i.e. $\mathcal{T}_\lambda = 20$), $a = 15$ and $\sigma_z = 0.15$.

Figure 8.4 presents heat-maps of the severity of capital crises $\log_{10} \Xi_k$ (top) and consumption crises $\log_{10} \Xi_c$ (bottom) across parameter combinations, where red indicates high prevalence. We show two representative sections of the parameter space: the planes (c_0, ν) (left) and (σ_z, a) (right). Since we present the logarithms of $\Xi_{k,c}$, the crossovers between high and low prevalence of different phases are rapid but smooth, i.e. there are no sharp phase transitions that characterize the system’s behavior. From each phase we study a point of the line $\nu = 1$ and different values of c_0 (marked by points in Figure 8.4), with all other parameters fixed and plot the dynamics of consumption c_t , labor n_t , capital k_t , and the measured Sharpe ratio \mathcal{S}_t in Figure 8.5. Note that changing the value of parameters (including ν) while staying in the same phase leads to qualitatively similar trajectories.

Prosperous Stability

As shown in Figure 8.5, leftmost column, the $\mathcal{L}_k \mathcal{L}_c$ phase is characterized by a stable capital surplus, low interest on capital and rare consumption crises. The depreciation of capital δ is so small that even with a puny level of investment, capital is always in excess and labor is the limiting factor. The stable capital surplus, in combination with a low confidence threshold c_0 , means that productivity shocks z_t hardly ever reach the required magnitude to trigger a consumption crisis, and if it does, recovery is almost immediate.

The corresponding bottom panel of Figure 8.5, shows that the consumption level has normal fluctuations, entirely due to exogenous productivity shocks z_t , around a single




Figure 8.5 – The upper panels show sample dynamics trajectories for each phase presented in Fig. 8.4 (marked by circles). In the topmost panels the solid black line corresponds to the consumption c_t , while the dashed red and blue horizontal lines show, respectively, the confidence threshold c_0 and the average consumption in the $\delta = 0$ scenario. Grey (resp. pink) background indicate capital scarcity (resp. consumption crisis). The middle row presents the dynamics of capital k_t (solid orange), labor n_t (solid green) and Sharpe ratio S_t (solid grey, with levels shown on the right y-axis). The lower panels show the histograms of consumption (black dots) and labor (green crosses) in a log – log scale. The green (resp. grey) dashed curve corresponds to the $\delta = 0$ baseline value for labor (resp. consumption), with c_0 indicated as a vertical red line. In the \mathcal{H}_c phase, the histograms of consumption and labor become bi-modal, corresponding to high output and low output regimes. For all simulations $\nu = 1$, $a = 15$, $\sigma_z = 0.15$. Specific parameters are $\mathcal{L}_k \mathcal{L}_c$: $\delta = 0.001$, $c_0 = 0.001$, $\mathcal{L}_k \mathcal{H}_c$: $\delta = 0.001$, $c_0 = 0.019$, $\mathcal{H}_k \mathcal{L}_c$: $\delta = 0.02$, $c_0 = 0.001$, $\mathcal{H}_k \mathcal{H}_c$: $\delta = 0.005$, $c_0 = 0.017$. Taken from Morelli et al. (2021).

high-consumption equilibrium. A corollary of the large capital excess is that the labor supply is nearly constant (extremely narrow-distribution in the $\mathcal{L}_k \mathcal{L}_c$ panel of Figure 8.5).

As the depreciation rate δ increases, the average excess of capital supply over labor shrinks, increasing the prevalence of capital scarcity. Accordingly, the phase \mathcal{L}_k quickly disappears upon increasing δ , leading to a pervasive $\mathcal{H}_k \mathcal{L}_c$ phase (see e.g. Fig. 8.4, third column, which shows that capital is always scarce when $\delta = 0.02$). As δ is further increased, the \mathcal{L}_c phase is more and more confined to small values of c_0 , i.e. when confidence is intrinsically robust.

Prevalent Capital Scarcity

The $\mathcal{H}_k\mathcal{L}_c$ phase is characterized by persistent capital scarcity with rare consumption crises, and is confined within a low c_0 “band” in the (c_0, ν) plane when δ is large enough. Since c_0 is low, confidence is generally high and therefore the household systematically consumes a large proportion G_t of its income, leaving only a small share for investment. Because of capital depreciation, the economy settles in a regime where $k_t < \sqrt{G_t/2\gamma}$, meaning that production is limited by capital, wages are low and rent on capital is high (i.e. the left region in Fig. 8.3). Hence, the average consumption level is lower than the maximal consumption level reached in the $\mathcal{L}_k\mathcal{L}_c$ phase – see Figure 8.5, second column.

But since c_t is now closer to c_0 , consumption crises are lurking around and the economy can flip into the $\mathcal{H}_k\mathcal{H}_c$ if c_0 increases and/or if productivity shocks are stronger (higher σ_z). This is clearly confirmed by the phase diagram of Fig. 8.4. In fact, comparing the phase diagrams for $\delta = 0.005$ and $\delta = 0.02$, we see that faster depreciation of capital converts large swaths of $\mathcal{H}_k\mathcal{L}_c$ phase into $\mathcal{H}_k\mathcal{H}_c$. Hence, in this case, investment crises (i.e. the supply side) do trigger consumption crises (i.e. the demand side) by reducing the difference between c_t and c_0 – see also the discussion in section 8.2.

Prevalent Consumption Crisis

When the depreciation rate is sufficiently small but the confidence threshold increases, capital remains abundant but self-reflexive confidence crises can hurl the system into a low consumption, low employment regime as a result of random productivity shocks. This is the $\mathcal{L}_k\mathcal{H}_c$ phase. Since capital is high, its level does not impact the level of production, and interest on capital is small. Hence the model becomes completely equivalent, in this regime, to the one studied in Morelli et al. (2020), where the dynamics is characterized entirely by the consumption propensity G_t and is dominated by frequent consumption crises, induced by breakdown of collective confidence.

As argued in Morelli et al. (2020) and shown in Figure 8.5, third column, consumption then displays bi-stable dynamics, where high and low consumption regimes alternate. Correspondingly, the distributions of consumption and labor reveal a secondary peak centered around the low-consumption equilibrium.

Note that during consumption crises (i.e. $G_t \searrow$) capital becomes even more abundant relative to labor (recall that one needs to compare k_t with $\sqrt{G_t/2\gamma}$) and therefore return on capital and Sharpe ratio both fall, as can be seen in the pink shaded region of Fig. 8.5, third column. If we are in a region where consumption crises are short enough compared to both the time \mathcal{T}_λ over which the Sharpe ratio is estimated and the capital depreciation time \mathcal{T}_δ , then one can avoid a capital crisis when confidence comes back. Otherwise, the economy enters a turbulent $\mathcal{H}_k\mathcal{H}_c$ phase with both capital and consumption crises.

Capital and Consumption Crises

This final $\mathcal{H}_k\mathcal{H}_c$ phase has both persistent capital scarcity and consumption crises. As anticipated above, capital crises can trigger consumption crises, because capital scarcity drives consumption closer to the confidence threshold c_0 , below which consumption drops and precautionary savings increase. One can then enter a doom loop (similar to Keynes' famous paradox of thrift) where now capital is too high and leads to a reduction of incentive to invest away from bonds. Hence, as shown in the fourth column of Fig. 8.5, capital and labor fluctuate around low levels, with intertwined periods of capital scarcity (gray regions) and high unemployment (pink regions). The Sharpe ratio gyrates rather strongly between negative values and values close to unity, with a significant negative skewness. The economy is unstable and always far from its optimal state.

Recall that we have fixed the interest rate on bonds to a constant value. But with a massive demand for bonds, as expected in the $\mathcal{H}_k\mathcal{H}_c$ phase, one should expect the government to borrow at low rates and prop up the economy with public investment, a feature not modeled in the current framework, but certainly worth accounting for in a later version of the model.

Summary

To summarize, we have identified four qualitatively different phases of the dynamics. Possibly the most interesting (and novel) one is $\mathcal{H}_k\mathcal{H}_c$, where capital scarcity is persistent, thereby triggering consumption crises. In this phase the economy is unstable, as capital becomes scarce the likelihood of a consumption crisis increases, and vice versa, low consumption drives rent on capital down and increases the risk aversion of investors.

We have underlined the role of the capital depreciation rate δ in determining the fate of our model economy. In fact, when capital and infrastructure are sufficiently long-lived such that $z_0 \cdot \mathcal{T}_\delta$ is large, the economy reaches a stable and prosperous state $\mathcal{L}_k\mathcal{L}_c$, provided self-induced confidence crises are rare enough (i.e. c_0 small). Conversely, when $z_0 \cdot \mathcal{T}_\delta$ is low, capital depreciates too quickly and this dents the rents that can be expected by investors. The economy quickly becomes under-capitalized and inefficient, especially because the dearth of capital makes confidence crises more probable, paving the way for the existence of a dysfunctional $\mathcal{H}_k\mathcal{H}_c$ region in the phase diagram.

Investment and Crisis Recovery

In the previous section, we have explained how capital depreciation can cause instabilities, and the appearance of a $\mathcal{H}_k\mathcal{H}_c$ phase where both capital and consumption undergo regular crises. In this section, we want to explore the influence of the memory timescale \mathcal{T}_λ , which is the history span over which investors assess the Sharpe ratio of capital investment, and of the sentiment parameter ν on the time needed for recovery when in a crisis period. We focus on this turbulent phase of the economy. We will fix the other two timescales $\mathcal{T}_\delta, \mathcal{T}_\eta$ defined in section 8.1 to, respectively, 200 and 2.




Figure 8.6 – The upper (resp. lower) left heat-map shows $\log_{10} \Xi_c$ (resp. average Sharpe ratio) as a function of the memory time scale \mathcal{T}_λ and the weight parameter ν . The dark red zones correspond to the regions where $\Xi_c > 0.99$, where crises are permanent. The right set of panels shows snapshots of the dynamics of c_t , k_t and n_t , corresponding to recovery (left) and crisis formation (right), both for $\mathcal{T}_\lambda = 20$ and $\nu = 0.75$. The two bottom rows show the histograms of Sharpe ratio \mathcal{S} and crisis duration $\mathcal{T}_<$ for four values of \mathcal{T}_λ : 2, 10, 20 and 50, all for the same value of ν (shown as symbols in the two heat-maps on the left). The faded grey lines show the same histograms in the benchmark case $\nu = 1$ that correspond to the $\mathcal{H}_k \mathcal{H}_c$ point in Fig. 8.4. Other parameters used are: $\delta = 0.005$, $c_0 = 0.017$, $\sigma_z = 0.15$ and $a = 15$. Taken from Morelli et al. (2021).

Our benchmark will thus be the $\mathcal{H}_k \mathcal{H}_c$ point in Fig. 8.4, corresponding to $\lambda = 0.95$, $\delta = 0.005$, $\nu = 1$ and $c_0 = 0.017$ (with a and σ_z also fixed at their baseline values). Looking at the statistics of the high consumption periods and of the low consumption periods, we conclude that the prosperous periods last a time $\mathcal{T}_>$ of the order of $\mathcal{T}_\delta = 200$ (data not shown), whereas crises are rather short, of the order of $\mathcal{T}_< \approx 10$, see Fig. 8.6, plain vertical lines in the third graph of the bottom row, which corresponds to $\mathcal{T}_\lambda = 20$, i.e. $\lambda = 0.95$. The full distribution of $\mathcal{T}_<$ and of the Sharpe ratio \mathcal{S} for $\nu = 1$ are shown in light grey, and reveals that whereas its time averaged value of \mathcal{S} is clearly positive and equal to $\langle \mathcal{S} \rangle_t \approx 0.71$, its full distribution is uni-modal but quite broad and negatively skewed.

Fig. 8.6, left graphs shows the consumption crisis prevalence Ξ_c and the average Sharpe ratio \mathcal{S} as $1 - \nu$ (weighing confidence in the investment allocation decision) and \mathcal{T}_λ are varied. One sees that decreasing ν or increasing \mathcal{T}_λ leads to an increase of Ξ_c , at least in the range shown, $\mathcal{T}_\lambda \lesssim 300$. The evolution of the average Sharpe ratio is more complex, reflecting the non-trivial shape of its distribution function (bi-modal and skewed, see

below). But certainly as agents pay less attention to the actual return on capital and are more affected by the level of confidence, the average Sharpe becomes strongly negative (red region of the diagram) and the economy gets trapped forever in a low consumption, low investment regime where $\theta_k \Sigma$ is negative (see Eq. (8.36)).

Now, let us look at a cut along the direction $\nu = 0.75$, corresponding to a 25% weight given to confidence in the allocation decision, as \mathcal{T}_λ is varied. For this particular value of ν , the average Sharpe ratio is close to zero and only weakly depends on \mathcal{T}_λ (Fig. 8.6, bottom left graph). But from the bottom row of Fig. 8.6, we see that when $\mathcal{T}_\lambda \gtrsim 10$, the distribution of Sharpe ratios becomes *bi-modal* and with a skewness that decreases as \mathcal{T}_λ increases. This can be rationalized as follows:

- The peak corresponding to positive Sharpe ratios comes from prosperous periods, where consumption is high and capital relatively scarce, leading to a positive return on capital q^* : see Fig. 8.6, top right graphs: in the high consumption phase, the orange line (capital) is below the green line (labor).
- The peak corresponding to zero Sharpe comes from crises periods, where capital is in slight excess of labor, leading to a small return on capital (see again Fig. 8.6, top right graphs, and section 8.1).
- The fat left tail corresponding to negative Sharpes comes from the transitory periods between high confidence and low confidence, when consumption and labor collapse but capital depreciates much more slowly. In this case, return on capital plummets and the Sharpe ratio becomes negative.
- As \mathcal{T}_λ increases, the weight of these transitory regimes in the estimate of the Sharpe ratio becomes small, and the fat left tail disappears, as crises become less frequent and much longer.¹³

Whereas the length of the prosperous periods $\mathcal{T}_>$ is unchanged compared to the benchmark $\nu = 1$ case for all values of $\mathcal{T}_\lambda \gtrsim 10$, the length of the crisis periods $\mathcal{T}_<$ increases by more than a 100 times as ν is decreased from 1 to 0.75 (compare the grey line and the coloured points in bottom row in Fig. 8.6). The first observation is due to the fact that the Sharpe ratio estimated when in a high output period is clearly in positive territory and quite insensitive to \mathcal{T}_λ (see the histograms in Fig. 8.6). This means that capital supply is also independent of \mathcal{T}_λ and that the confidence collapse mechanism must be identical to the one described in our previous work and not triggered by a lack of investment.¹⁴

On the contrary, the mechanism by which confidence is *restored* is strongly impacted by the value of the memory time \mathcal{T}_λ . When the economy is in a consumption crisis,

¹³For very large \mathcal{T}_λ , the situation changes again, see below.

¹⁴This is not to say that the crisis frequency is not related to capital abundance. As already noted in section 8.2, as capital depreciation δ increases, available capital decreases, which leads to a lowering of output c_t . Since the distance between c_t and the threshold c_0 is a crucial determinant of the probability of a confidence crisis, the region $\mathcal{H}_k \mathcal{H}_c$ becomes pervasive as δ increases, see again Fig. 8.4.

the returns on capital are very small. Thus, averaged over a sufficiently long time period, the Sharpe ratio is well defined and also small (see the narrow peaks in the Sharpe ratio distribution in Figure 8.6). Combined with the low confidence dampener on sentiment $(1 - \nu) \cdot \mathcal{C}_t$, this leads to a negligible investment flow. So whereas positive productivity shocks should put the economy back on an even keel, the level of capital is lagging, which creates a ceiling that prevents consumption (and hence confidence) from increasing substantially and returning to the high consumption case.

Interestingly, the dependence of $\mathcal{T}_<$ on \mathcal{T}_λ is in fact non-monotonic. For very large $\mathcal{T}_\lambda \gtrsim 1000$, the memory of prosperous periods persists even during the crises, so that the Sharpe ratio and investment always remain high. In such cases, $\mathcal{T}_<$ abruptly drops back to small values $\lesssim 5$ (data not shown). With extremely small probability, however, the system remains trapped in a crisis forever.

In the opposite case of a small enough \mathcal{T}_λ , the short periods where consumption increases due to productivity shocks allow sufficiently rapid increases in capital rent to encourage immediate investment. This is enough to prop up capital and allows confidence to be fully restored as labor and consumption will grow with the limiting factor k_t . For an example of these positive spikes of consumption, see top center panel in Fig. 8.6. The same effect takes place if ν is increased back to 1, where only realized Sharpe affects investment. In this case, the drag on capital due to low confidence levels is absent, and the system is able to pull itself out of the rut much more efficiently, leading to shorter crisis periods. But for lower values of ν (higher impact of household confidence on the investment propensity), the dearth of capital in crisis periods is such that the economy is unable to ever recover, i.e. $\mathcal{T}_< = +\infty$ for all purposes.

From a policy point of view, reducing interest rates has the direct effect of increasing the Sharpe ratio and reducing the return to bonds, thus promoting investment and making the transition back to the high consumption state easier. However, this may require the central bank to set interest rates r to negative values, as r which might already be close to zero due to prior crises. Besides monetary policy, other measures that improve confidence (e.g. central bank messaging) and/or promote investment into productive capital would have a similar impact (for instance if the government decides on strong fiscal measures that include investment into productive capital, such as through mission-oriented policies or infrastructure spending).

Finally, let us mention that while the existence of consumption crisis is independent of the substitutability parameter ρ , the duration of the low investment, low consumption periods is also highly sensitive to substitutability effects. We have indeed found that when ρ is sufficiently small, i.e. for production functions closer to Cobb-Douglas than to Leontief, recovery is much faster (data not shown). This could have been expected: lack of capital can now be compensated by labor, expediting the transition back to a prosperous state of affairs.

8.3 Discussion and Conclusions

We have constructed a behavioral real business cycle model in which labor and capital are nearly unsubstitutable. In the model, consumption and investment are controlled by (a) the confidence of households, which is self-reflexive (i.e. agents take cues from the consumption of other agents to determine their consumption budget) and (b) the quality of the excess real return to capital, as measured by the Sharpe ratio. As we have shown in Morelli et al. (2020), the self-referential nature of confidence amplifies the effect of productivity shocks on output, and can lead to crises where consumption abruptly jumps from a high equilibrium level to a low equilibrium level. Depending on the parameters of the model, these crises can be more or less frequent, and the low consumption periods can be of various duration: short spikes (“V-shaped crises”) or long drawn-out phases (“L-shaped crises”).

In the present study we investigate how the introduction of capital affects these dynamical patterns. In our model, capital can either be abundant (in which case labor is the limiting factor to production) or scarce. The main factors determining the quantity of working capital are the depreciation rate and the propensity of the households to save and invest, which itself depends on the return on capital. The resulting phenomenology of the model is quite rich. Our analysis reveals the following main takeaways:

1. Higher capital depreciation rates, *ceteris paribus*, lead to capital scarcity and limit production. This makes the economy more prone to confidence crises, increasing their prevalence;
2. Increasing the influence of the level of confidence in capital allocation decisions creates a feedback loop similar to Keynes’ paradox of thrift, destabilizing and trapping the economy into a non-optimal low consumption state;
3. The time during which the economy remains in a low output state is highly sensitive to the time span over which investors compute the Sharpe ratio. Increasing this memory timescale leads to sluggish adjustments of investment.¹⁵ Consequently, instantaneous increase of capital returns due to productivity upticks are not sufficient to boost the investment propensity. This leads to a persistence of capital scarcity, and prevents the economy from escaping the low output trap.

Our findings have different policy implications. As already emphasized in Morelli et al. (2020), if self-reflexive feedback loops exist, then governments and monetary authorities should not only manage inflation expectations but more broadly confidence in the future prospects of the economy. Although confidence indices are routinely measured by polling institutes (see e.g. Fig. 8.1), the inclusion of such indices in macroeconomic DSGE models and the importance of narratives (Shiller, 2019) have never really been considered seriously beyond the impact of news shocks on productivity.¹⁶

¹⁵Note however, as reported in the previous section, that extremely long memory timescales allow the Sharpe ratio to stick to high values.

¹⁶e.g. the work of Beaudry et al. (2020) reflecting interactions and complementarity. Also Angeletos

Beyond communication and narratives, our model suggests that monetary authorities should also directly promote investment, in particular during recessions. This is needed to prevent the economy being trapped in a low output, low confidence environment. Although this conclusion looks perfectly intuitive, our model reveals that a lack of capital can prolong crisis periods by orders of magnitude, and convert V-shaped crises into L-shaped crises. Boosting investment in working capital can be done through traditional channels, by lowering the risk-free interest rate (possibly making it negative) or by direct Keynesian investments in infrastructure and in innovation, which have the double effect of increasing the productivity of capital and propping up household confidence.

There are of course many directions in which our model should be extended and improved. The first obvious one is to allow interest rates and inflation to be dynamical variables, and to introduce an explicit monetary policy with the central bank monitoring inflation and confidence. A fully developed DSGE model building upon the framework proposed here would be welcome. Other relevant extensions could be to include a feedback mechanism between confidence and the time scale \mathcal{T}_λ or the sentiment parameter ν . This would allow potentially relevant panic effects to set in the model, and capture what happened in 2008, for example.

Another possible extension is to allow the parameter a which describes the default risk on capital to depend on the state of the economy, since bankruptcies are more frequent when the economy is in a low output, low investment regime.

Last, but not least, we have assumed that confidence is only a function of past realized output, but other factors should obviously be taken into account to model the dynamics of the confidence index, in particular financial news (like in 2008) or geopolitical news. Our framework would lead to scenarios where a shock like Lehman's bankruptcy simultaneously affects both consumption and investment, leading to a deep and prolonged recession, even in the absence of any "true" productivity shock. Conversely, good news about the future (e.g. technology shocks) could help to recover faster from the low output trap.

It would also be interesting to look for a "grand unification" between the type of behavioral business cycle/DSGE models considered in this chapter and heterogeneous agent based models studied in the recent literature, which generically give rise to similar crises and bi-stable dynamics between high output and low output regimes of the economy (see e.g. Gualdi et al. (2015) and Sharma et al. (2020)).

et al. (2020) showed that single news shocks (confidence shocks) are sufficient to fit empirically the effects of business cycles. See also Deniz and Aslanoglu (2014).

Key Messages

- This chapter extends the behavioral real business cycle model of Morelli et al. (2021) by introducing capital as a factor to production, which, together with consumption, is driven by household confidence.
- Depending on the parameterisation, multiple phases are possible: a stable high or low equilibrium, or crises with short or long drops.
- These phases depend on model-timescales: the degree of depreciation leading to capital scarcity, the memory timescale of investors, and the degree of influence confidence in the general state of the economy has on capital allocation.
- The policy implications are that beyond managing the exact expectation of inflation, policymakers must also address general confidence in the state of the economy. A task made difficult when, as in the United States, consumer sentiment is sharply divided by political affiliation.

Post-COVID Inflation & the Monetary Policy Dilemma

This chapter has been taken from Knicker et al. (2023) *Post-COVID Inflation & the Monetary Policy Dilemma* with only minor modifications to fit the thesis. This was a joint work with M. Knicker, F. Zamponi, and J-P Bouchaud. My main contributions in this paper were the economic framing of the debate on the causes of inflation, the calibration of the shocks applied to Mark-0, and contributing to the analysis of the policy outcomes. Overall, the drafting and analysis was done in discussion and collaboration with all other co-authors.

Inflation has captured global attention since the onset of the COVID-19 pandemic (from now on, simply “COVID”) in 2020. In the United States annual inflation reached 4.8% in 2021 and peaked at 9.1% in June 2022,¹ while Europe experienced highs of 11.5% in October 2022.² Competing narratives have emerged to explain the mechanisms driving this inflationary surge, which has persisted longer than expected. Policymakers have been blindsided by inadequate models, with the Bank of England admitting it had “big lessons to learn” from failure to forecast inflation using existing models.³ This chapter aims to explore the influence of fiscal and monetary policies on prevailing inflationary dynamics within a complex macroeconomic environment, modeled using the Mark-0 Agent-based Model that is based on alternative modeling foundations.

¹see e.g. U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; May 24, 2023.

²see e.g. Eurostat EuroIndicators, report n. 31/2023.

³see for example the Financial Times, <https://www.ft.com/content/b972f5e3-4f03-4986-890d-5443878424ac>

During the period of high inflation in 2021 and 2022, several theories emerged to explain the underlying causes of persisting inflationary trends. These interpretations led to differing views on appropriate policy responses, ranging from monetary policy interventions such as interest rate hikes to more targeted fiscal policies and price controls. Our analysis contributes to the debate on the appropriate policy responses to post-COVID inflation by providing a flexible framework to assess different policy options in the context of various inflation drivers, including demand-pull, cost-push, and profit-driven inflation. Our framework can accommodate varying behavioral foundations within a complex economy, including agents' trust in the Central Bank's clout and the anchoring of inflation expectations. Our main conclusions are that (i) the economic recovery after the shocks, especially in absence of mitigating fiscal policies, can be much more sluggish than expected – or even fall into deep recessions beyond dangerous tipping points or “dark corners” (Blanchard, 2014); (ii) the policy response (both from the government, the Central Bank, and other public authorities) has to navigate a narrow path, facing the trade-off between high inflation (with the risk of a runaway scenario) and high unemployment (with the risk of an economic collapse). In particular too weak a fiscal stimulus is ineffective, and too large a stimulus fuels high inflation; (iii) the success of monetary policy in curbing inflation is primarily due to expectation anchoring, rather than to direct impact of interest rate hikes; (iv) the two most sensitive model parameters (in terms of the inflation outcome) are those describing wage and price indexation, or in other words the bargaining power of workers and the market power of firms.

An initially dominant view of the post-COVID inflation was based on the “too much money chasing too few goods” theory of inflation (see e.g. de Soyres et al., 2022b,a), also known as demand-pull inflation. This narrative focuses on the large amount of fiscal stimulus, such as the Coronavirus Aid, Relief and Economy Security Act (the “CARES Act”) or the American Rescue Plan in the U.S., that have led to excess demand due to increases in disposable income while supply has not adjusted, thus pushing up prices. Some, such as Ferguson and Storm (2023) agree with the demand-pull inflation analysis but contend that “the final cause of the inflationary surge in the U.S., therefore, was in large measure the unequal (wealth) effects of ultra-loose monetary policy during 2020-2021”.

In contrast to the excess demand interpretation, scholars such as Stiglitz and Regmi (2022) conclude that “today’s inflation is largely driven by supply shocks and sectoral demand shifts, not by excess aggregate demand”. This claim is supported by the study of Cavallo and Kryvtsov (2023), who present empirical evidence spanning the years 2020 to 2022, illustrating the significant inflationary consequences of unexpected disruptions in product availability and stockouts across various sectors. In particular, the energy and food price shocks following COVID and the Russian invasion of Ukraine in 2022 were major causes of inflation, but this time from a cost-push inflation perspective wherein firms pass on increases in costs to consumers through prices. The strength of the inflationary surge was in large part due to the systemically important nature of the sectors where inflation occurred: energy and food (Weber et al., 2022). Combined with

supply chain bottlenecks that have also made headlines after COVID, these observations suggest an alternative cost-push scenario that is outside the realm of Central Bank's policy toolkit. This is also the conclusion of Bernanke and Blanchard (2023), where a detailed decomposition of inflation over different factors is proposed.⁴

Finally, a recent debate has emerged around profit-driven inflation, wherein sellers have increased prices beyond the increase in costs they face, thus expanding their profit margins. This is based on the observation that profits have increased sharply in the current inflationary climate, as compared to the observed amount of cost increases, even when including concerns of a wage-price spiral, and the excess demand cited by a monetarist perspective (Glover et al., 2023; Stiglitz and Regmi, 2022). Weber and Wasner (2023) analyzed firms' earnings calls and posit an interpretation based on imperfect competition and market power. Specifically, in an environment of cost increases and supply bottlenecks, firms' market power temporarily increases. This may be enhanced by the fact that when inflation is high, uncertainty about prices increases, such that consumers may be prone to accepting unreasonably high prices (an effect sometimes called "consumer discombobulation"), thus allowing companies to enlarge their margins.

All of these different interpretations lead to different guidelines for how policy should respond, and whether it should at all. From a standard monetarist perspective, the Central Bank should raise interest rates in order to push down demand, thus solving the excess demand scenario. On the other hand, in a cost-push scenario, the Central Bank's rate has no effect on the external cost increases and may actually harm the situation if firms choose to pass on the increased costs of debt to consumers. On this, Stiglitz and Regmi (2022) suggest "monetary policy, then, is too blunt an instrument because it will greatly reduce inflation only at the cost of unnecessarily high unemployment, with severe adverse distributive consequences", as in their view inflation is not due to excess demand. In reality, the U.S. Federal Reserve has hiked rates at the fastest rate since Paul Volcker was chairman, in line with a monetarist view. Simultaneously, in 2023 inflation has begun easing, which begs the question of whether the reduction in inflation throughout 2022 is due to monetary policy or to external factors such as the easing of energy prices as with the oil crises of the 1970s (Blinder, 1982; Blinder and Rudd, 2013), or else to a spontaneous tendency of inflation to self-heal due to economic forces. But since the Central Bank's mandate to keep inflation low and stable, they are expected to act, generally through interest rate mechanisms and communication.⁵ The question remains, to what effect? Are dips in inflation due to exogenous or endogenous factors or to Central Bank policy? What are the consequences of raising rates in this environment?

⁴We became aware of this paper, dated May 23, 2023, in the last week before finalizing our own work. Similarly to our own approach, Bernanke and Blanchard (2023) include wage bargaining, labor tightness and trust anchoring in their framework. Among the most important differences, however, are the absence of (i) supply-demand imbalances in the price setting mechanism and (ii) possible price gouging effects. Furthermore, the multiple equilibria and corresponding "tipping points" found within our model do not exist in their simplified specification.

⁵There is an ongoing impression that monetary policy is the "only game in town" when it comes to inflation response.

Are the consequences of not acting on inflation with monetary policy greater than those when one does? After all, Bruno and Easterly (1998) found that “countries can manage to live with relatively high – around 15-30 percent – inflation for long periods”. It is these questions upon which we aim to shed some light by considering various scenarios generated through an Agent-based Model, that allows us to run various counterfactual scenarios.

To conduct our analysis, we use the Mark-0 Agent-based Model (ABM) originally proposed by Gualdi et al. (2015), and extended in Gualdi et al. (2017); Bouchaud et al. (2018) to study monetary policy and inflation, with an early application to the effects of the COVID pandemic by Sharma et al. (2020) and a model-exploration study by Naumann-Woleske et al. (2023, which I have extended in Part I of this thesis). The Mark-0 ABM is a simplified model of a closed macroeconomy that nonetheless generates a wide variety of phenomena, from stable low unemployment and inflation, to endogenous crises that may oscillate regularly or punctuate long periods of recovery, or even to runaway inflation (see Chapters 5 and 6 for an introduction to the model’s dynamics). The philosophy behind the model is to generate qualitatively plausible scenarios. In this context, Gualdi et al. (2017) studied the efficacy of monetary policy in maintaining low unemployment and inflation, and “find that provided the economy is far from phase boundaries (or ‘dark corners’ (Blanchard, 2014)) such policies can be successful, whereas too aggressive policies may in fact, unwillingly, drive the economy to an unstable state, where large swings of inflation and unemployment occur.” An analysis using the Mark-0 model in fact predicted, as early as June 2020, that the post-COVID recovery could be more sluggish than expected and lead to a period of sustained inflation (Sharma et al., 2020).

In this chapter, we present a more detailed study of the post-COVID recovery by considering three distinct shocks occurring in the Mark-0 model: (1) a COVID-shock that negatively impacts firms productivity and consumer demand for the period of lockdowns, (2) a supply-chain shock on firms’ productivity that mimics the after-effects of COVID on global value chains, and (3) an energy price shock that is exacerbated by the Russian invasion of Ukraine. Each of these shocks is calibrated to macroeconomic time-series for the United States, such that the magnitude and duration match the observed data. We then study the effects of these shocks under different assumptions about the activity and efficacy of monetary policy, including the strength of monetary policy reactions to inflation, the ability of the Central Bank to influence firms’ expectations, and the structure of firm decision-making on prices and wages.

Our first results confirm that in the presence of properly calibrated shocks, the economic recovery in absence of any mitigating policy is extremely sluggish, taking least several years (if not much more) for the economy to return to the pre-crisis equilibrium. We note that this happens with shocks that are calibrated on the observed post-COVID macroeconomic time series, which already incorporate the effect of actually implemented policies; we expect that in absence of any policies the impact would have been even greater. This result thus points to the necessity of some sort of fiscal policy to prevent

the economy from spiraling into a crises or sluggish recovery.

A second set of results concerns the impact of such policies. We find that the model properly accounts for the inflation-unemployment trade-off, which leads to a narrow window in which policy can be efficient. For example, a disproportionate response of the Central Bank to inflation leads to unnecessary high unemployment, and an oversized injection of Helicopter Money leads to an unnecessary high inflation. Finally, we find that the anchoring (or de-anchoring) of economic agents' trust in the Central Bank strongly influences the dynamics of inflation due to the divergence in expected and realized inflation rates, which confirms the necessity for the Central Bank to manage inflation expectations. However, keeping expectations anchored to the Central Bank's target may come at a cost: if households believe that inflation is under control while rates are going up, consumption might be hobbled. Similarly, if firms believe that inflation is under control when it is not, wages will lag behind, again leading to a drop in consumption. Both effects, perhaps paradoxically, lead to increased unemployment compared to the case where expectations are floating. We conclude that each policy must be carefully tuned to achieve the desired result, and we believe that our modeling tool can guide policy makers in this respect.

Our primary objective is to provide possible scenarios and counterfactuals. As discussed at length in our previous papers (Gualdi et al., 2015; Bouchaud et al., 2018; Sharma et al., 2020), our ambition is not to provide precise predictions based on a fully calibrated model, but rather a tool for decision makers to help them apprehend different possible outcomes and anticipate unintended consequences and potential counter-intuitive impacts of their policies. We hope that Mark-0 can be usefully added to the policymakers toolbox and help them navigate a radically complex world (see e.g. King and Kay (2020); Bouchaud (2021)).

The remainder of the chapter is structured as follows: Section 9.1 gives an overview of the Mark-0 model and the adaptations made for this chapter. Pursuing this, Section 9.2 then outlines the policy channels that the model contains: interest rate, expectation management, access to credit, Helicopter Money and Windfall Tax. Section 9.3 discusses the model's dynamics in absence of exogenous shocks. In Section 9.4 we calibrate our three shocks, and show how they affect the macroeconomic dynamics with or without an easy credit policy, but without any monetary policy interventions by the Central Bank. In Section 9.5 we introduce and discuss monetary policy through interest rates and expectations management, and highlight its effects. In Section 9.6 we add fiscal policy by examining two distinct kind of stimuli, Helicopter Money and a Windfall Tax. Section 9.7 presents a discussion of the robustness of our model to variations of its parameters, and highlights the risk of hyperinflation. Finally, Section 9.8 summarizes our results and lays out some perspectives and ideas for future work.

9.1 The Mark-0 Model

The Mark-0 model was already briefly introduced in Chapter 5, with an outline of its main equations. However, in this Chapter we extend the model to incorporate three different shocks and study the agents' confidence in the Central Bank. Thus we here outline again the relevant equations for the model and explain the economic intuition behind all of them.

Model Overview

The model comprises a large set of firms producing a homogeneous consumption good that is purchased by a representative household sector. The household sector goes to the market with a consumption budget C_B computed as a fraction of savings, $S(t)$, wages $W(t)$, payouts from the energy sector $\delta_e \mathcal{E}_e(t)$, and interest on deposits $\rho_d(t)S(t)$,

$$C_B(t) = c(t) [S(t) + W(t) + \delta_e \mathcal{E}_e(t) + \rho_d(t)S(t)], \quad (9.1)$$

where $c(t) \in (0, 1)$ is the consumption propensity out of wealth and income, defined as

$$c(t) = c_0 [1 + \alpha_c (\hat{\pi}(t) - \rho_d(t))], \quad \alpha_c \geq 0, \quad (9.2)$$

where $\hat{\pi}(t)$, specified by Eq. (9.23) below, is the expected future inflation. Eq. (9.2) mimics the behavior of the classical Euler equation, as the household consumes more when real interest rates, $\rho_d(t) - \hat{\pi}(t)$, are lower. The savings of households over time can then be written as:

$$S(t+1) = S(t) + W(t) + \rho_d(t)S(t) - C(t) + \Delta(t) + \delta_e \mathcal{E}_e(t) \quad (9.3)$$

with dividends received from firms profits $\Delta(t)$ and actual consumption $C(t) \leq C_B(t)$ that is given by the matching of demand and production.

Households choose to split their consumption budget between N_F different firms (labeled by an index $i = 1, 2, \dots, N_F$), based on an intensity of choice model

$$D_i(t) = \frac{C_B(t)}{p_i(t)} \frac{\exp(-\beta p_i(t))}{\sum_j \exp(-\beta p_j(t))}, \quad \text{s.t.} \quad \sum_i p_i(t) D_i(t) \equiv C_B(t), \quad (9.4)$$

with a price sensitivity β .⁶

Because the model considers out-of-equilibrium situations, demand D_i for good i and production Y_i may not match, leading to a realized consumption c_i^R given by $\min(D_i, Y_i)$. Meanwhile, firms produce consumption goods based on a linear production function,

$$Y_i(t) = \zeta(t) N_i(t), \quad (9.5)$$

⁶Good differentiation could easily be included at this stage by replacing p_i by p_i/ψ_i , where ψ_i is the preference for good i .

dependent only on the firm's employed labor force N_i and time-dependent labor-productivity ζ . The unemployment rate $u(t)$ is then given by $u(t) = 1 - \sum_i N_i(t)/N$, where N is the number of workers. At each time step, corresponding to one month throughout this chapter, firms can adapt to the current economic environment by choosing three firm-specific variables: the target production $Y_i(t)$, price $p_i(t)$, and wage offered to their employees, $w_i(t)$. We describe these in turn, but note that the time step of one month and the different update parameters γ, g are assumed to be small such that production, prices and wages evolve slowly between t and $t + 1$, barring the role of external shocks of the COVID type, as discussed in section 9.4 below.

1. *Production update.* Production adapts to the observed gap between supply and demand at the previous time step as follows:

$$\begin{aligned} \text{If } Y_i(t) < D_i(t) &\Rightarrow Y_i(t+1) = Y_i(t) + \min\{\eta_i^+(D_i(t) - Y_i(t)), \zeta u_i^*(t)\} \\ \text{If } Y_i(t) > D_i(t) &\Rightarrow Y_i(t+1) = Y_i(t) - \eta_i^-[Y_i(t) - D_i(t)] \end{aligned} \quad (9.6)$$

where $u_i^*(t)$ is the maximum number of unemployed workers available to the firm i at time t , which depends on the wage the firm pays relative to the production-weighted average wage $W_i(t)/\bar{w}(t)$ ⁷

$$u_i^*(t) = \frac{e^{\beta W_i(t)/\bar{w}(t)}}{\sum_i e^{\beta W_i(t)/\bar{w}(t)}}. \quad (9.7)$$

The speed at which firms hire and fire workers depends on their level of financial fragility Φ_i , defined as the debt-to-sales ratio, where debt is $\mathcal{D}_i(t)$:⁸

$$\Phi_i(t) = \frac{1}{\Theta} \frac{\mathcal{D}_i(t)}{\min(p_i(t)D_i(t), p_i(t)Y_i(t))}. \quad (9.8)$$

Non-indebted firms have zero fragility and Θ is the maximum debt-to-sales ratio allowed by firms creditors, beyond which firms are declared bankrupt.⁹ Firms that are close to bankruptcy (i.e. $\Phi \approx 1$) are arguably faster to fire and slower to hire, and vice-versa for healthy firms. The coefficients of the hiring and firing rates η_i^\pm for firm i (belonging to $[0, 1]$) are given by:

$$\eta_i^\pm = [[\eta_0^\pm (1 \mp \Gamma(t)\Phi_i(t))]], \quad (9.9)$$

where η_0^\pm are fixed coefficients, identical for all firms, and $[[x]] = x$ when $x \in (0, 1)$, $[[x]] = 1$ for $x \geq 1$ and $[[x]] = 0$ when $x \leq 0$. The factor $\Gamma > 0$ measures how the financial fragility of firms influences their hiring/firing policy, since a larger value of Φ_i then leads to a faster downward adjustment of the workforce when the firm is over-producing, and a slower (more cautious) upward adjustment when the firm is

⁷The production weighted average wage is defined as $\bar{w}(t) = \frac{\sum_i W_i(t)Y_i(t)}{\sum_i Y_i(t)}$.

⁸Note that this definition of fragility Φ slightly differs from that used in our previous publication, in particular the normalization by the bankruptcy threshold Θ .

⁹For the detailed bankruptcy settlement, see Gualdi et al. (2017); Bouchaud et al. (2018).

under-producing. Γ itself depends on the inflation-adjusted interest rate and takes the following form:

$$\Gamma(t) = \max \left[\alpha_{\Gamma} (\rho_{\ell}(t) - \hat{\pi}(t)), \Gamma_0 \right], \quad \Gamma_0, \alpha_{\Gamma} \geq 0, \quad (9.10)$$

where $\rho_{\ell}(t)$ is the rate at which firms can borrow, Γ_0 and α_{Γ} are free, non negative parameters, the latter being similar to α_c that captures the influence of the real interest rate on the hiring/firing policy of firms.

2. *Price update.* Compared to the previous versions of the Mark-0 model, in this chapter the firms' price update, $\Delta p_i(t+1) = (p_i(t+1) - p_i(t))/p_i(t)$, explicitly takes into account the demand-output gap, and a contribution for the change in exogenous (e.g. energy) price $\Delta p_{e,ema}$ (in %), to with:

$$\Delta p_i(t+1) = \gamma \xi_i(t) \cdot \frac{D_i(t)}{Y_i(t)} + g_p \hat{\pi}(t) + g_e \Delta p_{e,ema}(t) \quad \text{if} \quad \begin{cases} D_i(t) > Y_i(t) \\ p_i(t) < \bar{p}(t) \end{cases} \quad (9.11)$$

$$\Delta p_i(t+1) = -\gamma \xi_i(t) \cdot \frac{Y_i(t)}{D_i(t)} + g_p \hat{\pi}(t) + g_e \Delta p_{e,ema}(t) \quad \text{if} \quad \begin{cases} D_i(t) < Y_i(t) \\ p_i(t) > \bar{p}(t) \end{cases} \quad (9.12)$$

$$\Delta p_i(t+1) = g_p \hat{\pi}(t) + g_e \Delta p_{e,ema}(t) \quad \text{otherwise,} \quad (9.13)$$

where

- $\xi_i(t)$ are independent uniform $U[0, 1]$ random variables scaled up by the actual demand-output ratio $D_i(t)/Y_i(t)$ (or its inverse), in order to mimic an increased pressure on prices in case of supply (or demand) gluts. When $D_i(t) \approx Y_i(t)$, this reduces to the rule used in previous papers, see Gualdi et al. (2015);
- $\hat{\pi}(t)$ is the expected next-period inflation in % per month, that firms take into account – or may even amplify – in their price setting mechanism;
- $\Delta p_{e,ema}(t)$ is the exponentially weighted moving average (with parameter ω , as in Eq. (9.22) below) of the exogenous price variations that are partially or fully passed on to final customers. The weighted average reflects that changes in energy price are not instantaneously transmitted to customers but rather distributed gradually over several months.

The magnitudes of the corresponding price adjustments are determined by parameters: γ for supply-demand imbalance, g_p for inflation expectations, and g_e for changes in the exogenous price. In this regard, g_e can be thought of as the effective energy-share of production that firms want to pass on to customers, whereas $g_p > 1$ would correspond to the much discussed concept of “greedflation”.

Note that \bar{p} is the consumer price index, defined as

$$\bar{p} = \frac{\sum_i c_i^R p_i}{\sum_i c_i^R}, \quad c_i^R = \min(D_i, Y_i), \quad (9.14)$$

with c_i^R the realized consumption of product i .¹⁰

3. *Wage update.* In similar fashion, firms update wages as

$$\Delta w_i(t+1) = \gamma(1-u(t))(1-\Gamma(t)\Phi_i(t))\xi'_i(t) + g_w \hat{\pi}(t) \quad \text{if} \quad \begin{cases} D_i(t) > Y_i(t) \\ \mathcal{P}_i(t) > 0 \end{cases} \quad (9.15)$$

$$\Delta w_i(t+1) = -\gamma u(t)(1+\Gamma(t)\Phi_i(t))\xi'_i(t) + g_w \hat{\pi}(t) \quad \text{if} \quad \begin{cases} D_i(t) < Y_i(t) \\ \mathcal{P}_i(t) < 0 \end{cases} \quad (9.16)$$

$$\Delta w_i(t+1) = g_w \hat{\pi}(t) \quad \text{otherwise} \quad (9.17)$$

with $\Delta w_i(t+1) = (w_i(t+1) - w_i(t))/w_i(t)$, and where $u(t)$ represents the unemployment rate, $\xi'_i(t)$ is a $U[0, 1]$ random realisation, and g_w is the wage sensitivity to expected inflation, which could be seen as worker's bargaining power. Firm profits \mathcal{P}_i include the cost of debt \mathcal{D}_i , the revenue on cash \mathcal{E}_i and the cost of energy $g_e p_e Y_i$:

$$\mathcal{P}_i = p_i(t) \min\{Y_i(t), D_i(t)\} - w_i(t)Y_i(t) + \rho_d \mathcal{E}_i(t) - \rho_\ell \mathcal{D}_i(t) - g_e p_e(t)Y_i(t). \quad (9.18)$$

The conditions of wage change therefore depend on the demand-supply imbalance, current firm profit $\mathcal{P}_i(t)$ and firm financial health, and also on the tension on the labor market, since lower unemployment leads to higher wage increase. This allows the model to reproduce Phillips curve effects. When firms experience positive profits and maintain a positive cash balance, they distribute dividends to households as a fraction δ of their cash balance:

$$\Delta(t) = \delta \sum_i \mathcal{E}_i(t) \mathbb{I}(\mathcal{P}_i(t) > 0) \mathbb{I}(\mathcal{E}_i(t) > 0), \quad (9.19)$$

where $\mathbb{I}(E)$ is the indicator function of event E .

4. *Energy Sector.* In this chapter, we consider an exogenously varying energy price, $p_e(t)$. Before the explicit shock introduced in Section 9.4, we consider the real energy price to be constant and equal to the average price of goods, i.e. $p_e(t) =$

¹⁰Here again, our present definition of \bar{p} slightly differs from that used in our previous papers, where we used a production weighted index instead of a consumption based index. We conform in this chapter with the more standard definition of inflation based on the Consumer Price Index (CPI).

$\bar{p}(t)$. In each period, firms pay a total amount $g_e p_e(t) \sum_i Y_i(t)$ to the energy sector. Subsequently, the energy sector pays a fraction of its accumulated profits to households at a rate δ_e . Therefore the cash balance \mathcal{E}_e of the energy sector writes as

$$\mathcal{E}_e(t+1) = \mathcal{E}_e(t) + g_e p_e(t) \sum_i Y_i(t) - \delta_e \mathcal{E}_e(t) \quad (9.20)$$

with the fraction δ_e of the energy profits that is paid out to households as income. Here, δ_e can be understood as dividends, share sales, or other channels through which energy sector profits circulate back into the economy. The way we introduce the energy sector in this chapter as an accounting identity with exogenous prices is arguably simplistic, and can certainly be improved in a future version of the model.¹¹

Inflation Expectations and Monetary Policy

The equations above depend on inflation rate expectations $\hat{\pi}(t)$, that we now define. We consider that the measure of realized inflation is given by the change in consumption-weighted average price,

$$\pi(t) = \frac{\bar{p}(t) - \bar{p}(t-1)}{\bar{p}(t-1)}, \quad (9.21)$$

where \bar{p} is defined in Eq. (9.14).

In the model, agents form expectations of future inflation partly on the basis of past realizations and partly on the basis of their trust in the Central Bank ability to enforce its inflation target. More precisely, they use an exponentially weighted moving average over realized inflation,

$$\pi_{\text{ema}}(t) = (1 - \omega)\pi_{\text{ema}}(t-1) + \omega\pi(t), \quad (9.22)$$

where ω sets a memory time over which agents perceive realized inflation, equal to $-\log(1 - \omega)^{-1} \approx \omega^{-1}$ for small ω . Together with the Central Bank's communicated inflation target, π^* , all agents form the same inflation expectation, given by a weighted average of $\pi_{\text{ema}}(t)$ and π^* :

$$\hat{\pi}(t) = (1 - \tau^T(t))\pi_{\text{ema}}(t) + \tau^T(t)\pi^*, \quad (9.23)$$

where $\tau^T(t)$ is the degree to which expectations are anchored around the Central Bank's target.¹² Consistently with the definition of π_{ema} and the long-term nature of the inflation target, we interpret $\hat{\pi}(t)$ as the expectation of *long-term* inflation.¹³

¹¹Note that in the model there is a resource constraint such that the total money M (that is created by the Central Bank) is kept fixed during the simulation. The balance sheet of the banking sector can then be written as $M = S(t) + \mathcal{E}^+ - \mathcal{E}^- + \mathcal{E}_e$, with the household savings S and the cash balance of firms $\mathcal{E}^{+/-}$.

¹²For empirical work on the question of trust, see e.g. Christelis et al. (2020).

¹³This is in contrast to the common formulations in DSGE models that operate with one-period-ahead expectations ($t \rightarrow t+1$).

As far as monetary policy is concerned, we assume that the Central Bank sets the baseline interest rate, ρ_0 , via a classical Taylor rule based on observed (realized) inflation π_{ema} :

$$\rho_0 = \rho^* + \phi_\pi (\pi_{\text{ema}}(t) - \pi^*), \quad (9.24)$$

with reaction strength ϕ_π . The baseline rate ρ_0 is then translated into a time-dependent rate on loans, $\rho_\ell(t)$, and deposits, $\rho_d(t)$, adjusting for the cost of bankruptcies (see Gualdi et al., 2017; Sharma et al., 2020, for more details). These interest rates are defined as

$$\rho_\ell(t) = \rho_0(t) + f \frac{\mathcal{D}(t)}{\mathcal{E}^-(t)}, \quad (9.25)$$

$$\rho_d(t) = \frac{\rho_0(t)\mathcal{E}^-(t) - (1-f)\mathcal{D}(t)}{S + \mathcal{E}^+(t)} \quad (9.26)$$

with the positive cash balance $\mathcal{E}^+(t) = \sum_i \max(\mathcal{E}_i, 0)$ and firms total debt $\mathcal{E}^-(t) = -\sum_i \min(\mathcal{E}_i, 0)$. The parameter f determines how the consequences of defaults are allocated to lenders and depositors, interpolating between the costs borne entirely by borrowers ($f = 1$) and those shouldered entirely by depositors ($f = 0$). In this chapter, we do not consider a double mandate for the Central Bank as in Gualdi et al. (2017), instead focusing only on inflation. We will comment about this below, as strict inflation control may turn out to be highly detrimental to unemployment.

Finally, we allow inflation expectation anchoring to evolve dynamically via¹⁴

$$\tau^T(t+1) = (1-\omega)\tau^T(t) + \omega \exp \left[-\alpha_I \frac{|\pi(t) - \pi^*|}{\pi^* \phi_\pi} \right]. \quad (9.27)$$

This equation aims at capturing the fact that the degree of expectation anchoring depends on how closely the realized inflation actually matches the Central Bank target. This is factored into an exponentially weighted moving average with memory time $\approx \omega^{-1}$: realized and target inflation must differ significantly, and for sufficiently long times, for agents to lose trust in the Central Bank. A larger α_I means economic agents lose trust in the Central Bank more abruptly as the gap between realized inflation and inflation target becomes significant. We have included the factor ϕ_π in the denominator to emphasize the fact that stronger commitment of the Central Bank should decrease the sensitivity of anchoring on realized inflation, but of course this extra factor can be reabsorbed into α_I .

The calibration of the parameters we choose for the model is based on the results and studies done in the previous work around Mark0 (Gualdi et al. (2015, 2017); Bouchaud et al. (2018); Sharma et al. (2020)). Here, we discuss only the parameters that are relevant to the current work, in particular those identified by the sensitivity analysis in Section 9.7. The parameters used in this study can be found in the Table H.1.

¹⁴Note that in our previous paper (Bouchaud et al. (2018)), the anchoring parameter τ^T was assumed to be time independent, a case we will call ‘‘Anchored Trust’’ henceforth.

9.2 Policy Channels in Mark-0

In the augmented Mark-0 model, the central authorities can influence macroeconomic dynamics through (i) the manipulation of interest rates, (ii) anchoring inflation expectations, and (iii) regulating the amount of debt firms can accumulate. Additionally, there is the possibility of fiscal policy in the form of direct injection of cash in the economy (“Helicopter Money”) and Windfall Tax on the energy sector. These manipulations take effect through firms’ financial fragility, which directly determines the probability of bankruptcies and indirectly changes the propensity of firms to hire and fire workers, as well as the wage setting behavior. Interest rates and cash injection also have effects on the willingness of consumers to spend. In addition, there is an expectations channel whereby firms and households may or may not trust the Central Bank’s communicated inflation target depending on how close actual inflation is to the target. We now consider these mechanics in turn, as these will turn out to be important in response to the COVID related shocks that we will introduce in the next section.

The Interest Rate Channel

The Central Bank’s baseline interest rate ρ_0 additively influences both deposit (ρ_d) and lending (ρ_ℓ) rates. Through this channel, the interest rate affects firms’ wages and hiring strategies, as well as households’ spending.

The impact on households is straightforward and parallels that of the standard Euler equation of inter-temporal substitution: higher interest rates, all else being equal, reduce the propensity to consume out of income and wealth through parameter α_c (Eq. (9.2)), effectively decreasing current demand in favor of later consumption. This decrease leads firms to reduce or maintain prices due to excess supply (see Eq. (9.12)), thereby lowering inflation.

The impact of interest rates on firm behavior occurs through the coefficient α_Γ , defined in Eq. (9.10), which influences firms’ production target and wage adjustment. In all cases α_Γ affects firm behavior in combination with the firm’s fragility Φ (where $\Phi = 1$ implies bankruptcy, see Eq. (9.8)). An increase in the baseline rate will increase Γ , which implies that firms have stronger reactions to excess demand or excess supply. Unhealthy firms ($\Phi > 0$) will therefore more cautiously expand production when needed (η^+), but more abruptly fire staff in the case of excess supply (η^-). Hence, an increase of the baseline rate ρ_0 will tend to decrease production and increase unemployment as fragile firms feel the brunt of the cost of debt. For similar reasons, an increase of ρ_0 implies a stronger downward pressure on wages.

Therefore, in Mark-0 an increase of the baseline rate *ceteris paribus* induces a reduction of consumption and an increase of unemployment. The success of a rate hike in curbing inflation depends on the strength of the response of households and of firms to interest rates, and hence on the value of the parameters α_c and α_Γ . If, as some authors argue (Reis, 2022), the reaction of households and of firms to interest rates are subdued, large

hikes would be necessary for monetary policy to have a significant effect on inflation, in part by bringing about a recession (Stiglitz and Regmi, 2022).

The Expectation Channel

With rising global inflation, the debate about the importance of inflation expectations for macroeconomic dynamics has been rekindled. Some, such as Rudd (2022), argue that short-term inflation expectations play no role in macroeconomic dynamics, while others, like Reis (2021), maintain their importance but question whether they are “de-anchored” from Central Bank’s targets. In the present version of the Mark-0, long-term inflation expectations $\hat{\pi}$ play a role in price and wage settings: formally, the factor g_p in Eqs. (9.11) and (9.12) determines the fraction of expected inflation reflected in firms’ next-period prices. In a similar style, the factor g_w in Eqs. (9.15) and (9.16) controls the rate of wage adaptation to inflation in response to firms’ expected inflation, and is thus adjacent to the idea of labor-bargaining power.

The long term inflation expectation, $\hat{\pi}$, is determined as a mixture of past realized inflation and the Central Bank target, with the Central Bank’s effectiveness in reaching its target dictating the level of “anchoring” τ^T (see Eq. (9.23) and (9.27)). With this formulation, economic agents lose confidence in Central Bank actions when inflation significantly and persistently deviates from the Central Bank’s inflation target. The parameter α_I represents the sensitivity of trust to the Central Bank’s ability to control inflation.

De-anchored beliefs imply firms’ inflation expectation is based solely on past observation, running a risk of a wage-price spiral or a hyper-inflation runaway, depending on the imbalance between the parameters g_p and g_w . The simplest case is equality in bargaining and market power, $g_w = g_p$, which implies that on average both prices and wages are raised in the same proportion to expected inflation, thus having no real effects on the economy. One then faces either stable inflation when $g_w = g_p < 1$ or possible hyper-inflation when $g_w = g_p > 1$, see section 9.7.

A more interesting case is $g_p \neq g_w$, implying a differentiation in bargaining vs. pricing power. For $g_p > g_w$, labor bargaining power is lower than firms pricing power. In this case, firms raise prices by more than wages, eroding the purchasing power of workers. This results in lower demand that makes price hikes less likely until firms’ excess demand becomes excess supply, reversing the cycle. For $g_w > g_p$, on the other hand, an unstable wage spiral may set in, with higher wages driving demand up, leading to further increases of prices as firms face excess demand. The situation $g_p > 1$, at least for firms with a healthy balance sheet, would correspond to “greedflation”, i.e. firms trying to use inflation to hide increases in their profits.

Credit Regulatory Policy

The last channel for the monetary authority to affect the economy within the framework of Mark-0 is by means of the default threshold, Θ . Earlier work on the Mark-0 model

found that the default threshold is a key parameter causing a phase transition between a regime of high unemployment and/or endogenous crises (low Θ), and of stable full employment (higher Θ conditional on low interest rates) (see Gualdi et al., 2015). This is because higher thresholds give firms more time to adjust their production and price strategy to return to a profitable state, assuming interest payments are low. Empirically, both in the CARES act and the CAA there were relaxations of the bankruptcy laws in the United States.¹⁵ Direct state aid to firms may also prevent bankruptcies from occurring. In the realm of our model these policies are mimicked by increasing the bankruptcy threshold Θ of firms during the shock and reducing it again after the shock. The “Easy-Credit policy” proposed in Sharma et al. (2020) amounts to set the value of Θ such that it remains commensurate to the current average firm fragility, to wit

$$\Theta(t) = \max(\mu\langle\phi\rangle(t), \Theta_0), \quad (9.28)$$

where Θ_0 is the bankruptcy threshold before the shock, μ a multiplier, and $\langle\phi\rangle(t)$ the firm-wide average debt-to-sales ratio at time t , meaning that only the firms whose debt-to-sales ratio exceeds $\mu\langle\phi\rangle$ must file for bankruptcy. Throughout the following we will use $\mu = 1.3$ and $\Theta_0 = 3.2$, which represents a compromise between “too robust” economies when Θ_0 is large, and “too fragile” economies when Θ_0 is small.

This implies that with easier credit access, firms can accumulate debt during the shock without the fear of bankruptcy. This, in turn, affects the firing and hiring policies of firms as they can continue with their business as usual, leading to a hiring and firing as if they were not fragile (see Eq. (9.8)). Sharma et al. (2020) have shown that such policy is effective in speeding up the recovery to pre-shock levels after severe COVID-like shocks, without such policy long-lasting recessions would otherwise ensue.

Helicopter Money

When interest rates are low and access to credit is loose, but shocks nonetheless persist, central authorities may turn to less conventional policies such as “Helicopter Money” to address economic downturn. Introduced by Friedman (1969), this implies an injection of money directly into the hands of economic agents to increase spending and thus economic output.¹⁶ While some view Helicopter Money as a potential cause of hyperinflation and currency devaluation, it can be an effective response to financial crises and pandemics by stimulating real economic activity (Reis, 2022). In this chapter we consider Helicopter Money as introduced by Friedman: a distribution of newly printed money by the Central Bank directly to households, resulting in a decrease in the Central Bank’s net worth and an increase in the net worth of households. This injection of money is modeled as a one-time increase of savings of households by a factor $\kappa_H > 1$, such that $S \rightarrow \kappa_H S$

¹⁵CARES refers to the Coronavirus Aid, Relief, and Economic Security Act, and the CAA refers to the Consolidated Appropriations Act of 2021.

¹⁶By contrast, quantitative easing implies that money created by the Central Bank is used to purchase government bonds and distributed through the government, whereas Helicopter Money is directly distributed in the economy (Ugai et al., 2007).

instantaneously after all shocks. The increased savings of households boost demand through an increased consumption budget C_B (see Eq. (9.4)) and finally end up on the balance sheet of firms where it is then redistributed to households as wages and dividends.

Windfall Tax

Whereas Helicopter Money stimulates consumption by injecting external funds into households savings, a tax on “windfall” profits, which refer to large unforeseen profits (Chennells, 1997), achieves a similar outcome by redistributing money through taxation. Following the Russian invasion of Ukraine, there has been a significant surge in fossil fuel prices, leading to unexpectedly high returns for utilities and fossil fuel producing companies (Weber, 2022), and profit-driven price markups for other firms (Weber and Wasner, 2023), leading to strong debate among policymakers about implementing a Windfall Tax. The specific design of the Windfall Tax as a rent-sharing fiscal policy may vary depending on the particular circumstances (Baunsgaard and Vernon, 2022).

In this chapter, we implement a “Windfall Tax” as a temporary increase of the fraction of energy sector cash balance that is re-injected to the savings of households, i.e. $\delta_e \rightarrow \delta_e + \Delta\delta_e$. This tax has a duration of two years, commencing one year prior to the end of the price shock. The primary objective of increasing δ_e is to boost consumption by enhancing households’ savings through the redistribution of the increased cash balance in the energy sector that follow the energy price increases.

9.3 Stationary Dynamics: the Role of the Central Bank

The Economic Equilibrium Without Monetary Policy

To begin our analysis, we consider the economy without shocks. Our first choice of parameters corresponds to a Central Bank that does not react to inflation (i.e. $\pi^* = \phi_\pi = 0$) while maintaining low baseline interest rates ($\rho^* = 1.2\%$ p.a.). By consequence, agents form inflation expectations based only on their own observations ($\tau^T = 0$, see Eq. (9.23)). In this scenario, called “Inactive CB”, and for the choice of parameters detailed in Table H.1, the economy spontaneously evolves into a steady state of full capacity, as evidenced by the low, oscillating unemployment rate around 0.4% (see Figure 9.1). There are small endogenous output and unemployment fluctuations of roughly two years in duration. Simultaneously, there are stronger oscillations in the demand and supply for goods, and in real wages. These fluctuations cause firms to adjust both their production and prices, thus generating a small “business cycle” with periodic imbalance between demand and output due to price adjustments. But in the absence of an active Central Bank target, the amplitude of the resulting inflation oscillations is found to be substantial, ranging between 1.5% and 8% p.a., suggesting that monetary policy indeed iron out business cycles. As already found by Gualdi et al. (2015), these cycles emerge endogenously from the oscillating feedback loop of prices, demand and savings. In this scenario, called “Inactive CB”, and for the choice of parameters detailed in Table H.1,

9.3. STATIONARY DYNAMICS: THE ROLE OF THE CENTRAL BANK




Figure 9.1 – Economic dashboard for the three Central Bank scenarios described in sections 9.3, 9.3 in absence of any policy and shock. Blue lines: Inactive Central Bank scenario ($\pi^* = \phi_\pi = 0$, $\rho^* = 1.2\%$ p.a.). Orange lines: Reactive Central Bank with Anchored Trust scenario ($\pi^* = 2.4\%$ p.a., $\phi_\pi = 1$, $\tau^T = 0.95$). Green lines: Reactive Central Bank with Floating Trust scenario. The complete parameter set can be seen in table H.1. The two insets correspond to the demand/output ratio and the bankruptcy rate.

the economy spontaneously evolves into a steady state of full capacity, as evidenced by the low, oscillating unemployment rate around 0.4% (see Figure 9.1). There are small endogenous output and unemployment fluctuations of roughly two years in duration. Simultaneously, there are stronger oscillations in the demand and supply for goods, and in real wages. These fluctuations cause firms to adjust both their production and prices, thus generating a small “business cycle” with periodic imbalance between demand and output due to price adjustments. But in the absence of an active Central Bank target, the amplitude of the resulting inflation oscillations is found to be substantial, ranging between 1.5% and 8% p.a., suggesting that monetary policy may iron out business cycles, as indeed found below. As already discussed in Gualdi et al. (2015), these cycles emerge endogenously from the oscillating feedback loop of prices, demand and savings.

Note however that depending on the choice of parameters, the economy may settle into a much less favorable state. In particular, when the hire-to-fire ratio is too small or when the baseline interest rate is too high, the economy collapses, see Gualdi et al. (2015); Bouchaud et al. (2018). Similarly, when the bankruptcy threshold Θ_0 is too low, unemployment remains at a relatively high level.

Stabilizing Inflation through Monetary Policy

At this point, we introduce an active Central Bank and study its impact on inflation and unemployment. Specifically, we set the Central Bank’s inflation target to $\pi^* = 2.4\%$ p.a. and its reaction to inflation variation to $\phi_\pi = 1$. In this configuration, we assume that the Central Bank works with a well-anchored inflation expectation fixed at $\tau^T = 0.95$. In the pre-COVID period core inflation was quite stable as a result of an arguably successful monetary policy (Miles et al., 2017). Simultaneously, the Central Banks had been acting in a transparent and measurable way despite the historically low interest rates (Reis, 2021), suggesting a high degree of trust. We call this constant τ^T scenario “Reactive CB, Anchored Trust” in the following. In this case, Figure 9.1 shows that inflation is reduced to an average of $\langle \pi_{AT} \rangle \approx 2.7\%$, close to the 2.4% target, while nearly completely suppressing the business cycle with a very low stable unemployment. On the other hand, there is a slight reduction in real wages and higher interest rates.

Finally, we introduce a third scenario, referred to as “Reactive CB, Floating Trust”, where we allow agents’ trust in the Central Bank to vary depending on the perceived success of the CB to bring inflation to roost, see $\tau^T(t)$ defined in Eq. (9.27). We choose as parameter values $\omega = 0.2$ (corresponding to a memory time of 5 months) and a sensitivity to off-target inflation of $\alpha_I = 0.4$. Again, all other parameters are taken from Table H.1. In the absence of external shocks, this configuration leads to a similarly high degree of trust in the Central Bank, and is otherwise identical to the case where trust is anchored – simply because inflation is on target.

9.4 The COVID Shock and its Aftermath in the Absence of Monetary Policy

Modeling the Three COVID Shocks

In order to assess the effects of recovery and monetary policies in response to the events of 2020-2022, we introduce in Mark-0 three shocks calibrated on U.S. data: (1) a COVID shock, (2) a supply-chain shock, and (3) an energy price shock. The data was retrieved from the Federal Reserve Bank of St. Louis, and are shown in Figure 9.2 where we have indexed values to January 2020. We now discuss these shocks in turn:

1. **COVID:** The first major COVID outbreak emerged in the US in February 2020, prompting the US government to declare a public health emergency, followed by the implementation of stay-at-home orders in March of the same year. The COVID outbreak led to a significant reduction in personal consumption expenditure by households, falling by around 15% in February 2020 (Figure 9.2, dashed orange line), which was of the same order of magnitude as the decline in the Industrial Production Index during the same month (dashed blue line). It took about 5 months for aggregate Personal Consumption Expenditure to recover to its pre-pandemic level. In accordance with this, we consider a shock to the consumption propensity $c(t)$ of households by 15%, which appears in Eq. (9.2), recovering linearly over 5




Figure 9.2 – Empirical Shocks: Fitting of shock scenarios to macro data from the US. The empirical series (dashed lines) are indexed to January 2020. (Left Panel) Empirical series for the Personal Consumption Expenditure (orange, dashed line) and Industrial Production Index (blue, dashed line). Solid lines represent our model for the shocks to consumption propensity c (orange) and productivity ζ (blue). Dark gray area corresponds to the COVID Shock and the gray area to the Supply Chain Shock in the shock scenarios that we apply to Mark-0. (Right Panel) WTI Crude oil spot price during the periods of shock (green, dashed line). Modeling of the artificial price series (green, solid) in the Mark-0 model for a period of two years (light gray area). Data retrieved from FRED, specifically series PCE, WTISPLC, CPALTT01USM657N, INDPRO.

periods (solid orange line in Figure 9.2).

2. **Supply Chains:** During the initial COVID outbreak, firms in the US laid off a large number of workers, reducing their production to a similar degree as the personal consumption expenditure. However, while personal consumption recovered within 5 months, it took an additional 10-15 months for industrial production to return to its pre-pandemic state, as can be seen in the industrial production index (dashed blue line, Figure 9.2). This was due to a plethora of idiosyncratic supply chain disruptions, such as logistics and transportation difficulties, semiconductor shortages, pandemic-related restrictions on economic activity, and labor shortages that led to the slower recovery of production in the industrial sector (Attinasi et al., 2022). In the context of this chapter, we model this by a shock to firm productivity $\zeta(t)$ (defined in Eq. (9.5)) of an initial magnitude of 15%, with a recovery of 15 months, see solid blue line in Figure 9.2.
3. **Energy Prices:** Finally, we consider the energy price shock. The reduction of demand and production throughout the pandemic led to a supply glut in energy markets, which led to a steep decrease in energy prices, such as oil, by up to 70% for immediate delivery of West Texas Intermediate crude oil (Figure 9.2, dashed green line). As the recovery period began, external factors such as extreme weather conditions in various parts of the world and maintenance work that had been postponed during the pandemic caused a surge in demand, and the energy prices thus rebounded quickly (Zakeri et al., 2022). Unfortunately, with the Russian invasion of Ukraine in February 2022 this rebound was further exacerbated to a global energy crisis, due to Russia’s position as a major global exporter of natural

gas and oil. By June 2022, the WTI crude oil spot price had peaked, rising nearly 100% compared to pre-pandemic levels. Following this, the global recovery and adjustment of energy markets has led to a sharp easing of energy prices. For Mark-0, these processes are external to the model, such that we introduce here an exogenous price shock to firms' price update Eqs. (9.11)-(9.13). Specifically, firms' prices change by an additional exogenous factor $g_e \Delta p_{e,ema}(t)$, where g_e is a constant factor akin to the energy-share in production, and $\Delta p_{e,ema}(t)$ is an exponentially weighted moving average of the time-dependent monthly percentage change in energy prices. Thus, the transmission of the change in energy prices to firms' product prices is smoothed, as expected to be the case in reality.

The form of the moving average is the same as for trust in Central Bank Eq. (9.27) and inflation expectations Eq. (9.22), with the very same memory time parameter ω . Our artificial energy price series is shown in Figure 9.2 as the solid green line, and is based on the WTI Spot Price.¹⁷ Alternative calibrations with the energy component of the U.S. Consumer Price Index or the Henry Hub Natural Gas price lead to similar shapes and magnitudes. In the following we will choose $g_e = 3.25\%$, i.e. half of the energy share of the GDP in the US, where the factor 2 accounts for inventories and partial substitutability. We also posit that the profits of the energy sector are transferred back to households at an effective rate of $\delta_e = 4\%$ per time step, unless stated otherwise.

¹⁷Prior to the shock, we set $p_e(t) = \bar{p}(t)$.

9.4. THE COVID SHOCK AND ITS AFTERMATH IN THE ABSENCE OF MONETARY POLICY

The Effect of Shocks Without Stabilization Policies




Figure 9.3 – Economic dashboard for the three shocks in the Inactive Central Bank scenario and in the absence of any policies. The dynamics for the three shock scenarios, COVID only (blue), COVID and Supply Chain shock (orange) and all shocks (green) when the Central Bank is inactive and no policy is applied. The areas shaded in gray indicate the duration of the three shocks: the COVID shock lasting until the end of the dark gray area, the supply chain shock until the end of the gray area, and the energy price shock until the end of the light gray area. The dashed black line represents pre-shock averages. In the first two cases, the economy is able to recover on its own (although with significant fluctuations that last for a few years), but significantly gets worse when the energy shock is turned on.

To develop some intuition, we now build three counterfactual cases where (i) only the COVID shock, (ii) COVID + supply chain shocks and (iii) COVID + supply chain + energy price shocks hit the Mark-0 economy, without any stabilization policy put in place. It should be noted that this exercise is somewhat optimistic, as our shocks have been calibrated on an economy where emergency policies were actually implemented. This is discussed further below. Moreover, in the Inactive Central Bank scenario discussed in this section, economic agents form their inflation expectations solely on the basis of realized inflation, which means that there is no anchoring of expectations ($\tau^T = 0$, see Eq. (9.23)). All other parameters of this scenario can be found in Table H.1.

- Case (i) After the impact of only the COVID shock, the economy remains operational with full capacity but experiences a persistent and high inflation that peaks at 15.5% and only starts to recede by the end of 2023, with significant fluctuations that persist for a few years (Figure 9.3, blue). All other observables return to their pre-shock levels by the end of 2023, with small oscillations around the steady state in accordance with the Inactive Central Bank scenario without shocks.

- Case (i) When adding the supply chain shock (Figure 9.3, orange), the dynamics remain very similar: the economy returns to its steady state without fiscal or monetary policy. Remarkably, inflationary dynamics are almost identical, albeit with a slightly higher peak value (16.0%) reached slightly earlier in 2023. The biggest difference is the prolongation of the recovery period, which leads to a steeper real wage dip.
- Case (i) Taking now the energy price shock into consideration, the situation abruptly changes. Initially, we observe a short deflation period due to the steep drop in energy prices, increasing real wages, supply-demand imbalance and the fragility of firms due to the decreased prices.¹⁸ Immediately following this, we observe an explosion in inflation peaking at 23.2% mid-2023. Firms fragility increases due to the price shock which causes bankruptcies (peaking at 3 % rate), unemployment and a decrease of wages. In the long run, this disruption is so strong that wages recover only very slow. This starts a feedback cycle of low demand that leads to a decrease of output, which causes a reduction in savings and therefore demand drops. This feedback cycle consistently increases unemployment, and decreases savings and wages, which takes the economy significantly longer than 10 years to recover fully with a high and persistent unemployment between 6%-8%.

In the realm of the Mark-0 model, the COVID and supply chain shocks would thus have had a minor long term impact on output, but would have led to substantial medium term inflation due to excess savings, as predicted using the Mark-0 model as early as June 2020 in Sharma et al. (2020). The energy price shock, however, is the last straw on the camel's back and, in the absence of monetary and fiscal policies, has a strongly detrimental long-term impact on the economy. This is a consequence of the existence of discontinuous transition boundaries (a.k.a. *tipping points*) in parameter space, as emphasized in Gualdi et al. (2015) and, within the specific context of the COVID shock, in Sharma et al. (2020). As a scenario-generating tool for policymakers, our model demonstrates that in the absence of mitigating policies, the full sequence of COVID, supply chain and energy shocks can trigger a negative feedback cycle, manifested as a downward wage spiral, that results in a collapse of demand and a full blown crisis. However, we see that the first two shocks, as modeled above, are of sufficiently mild amplitude not to trigger a complete collapse of the economy. In the following section, we study counterfactuals with larger "bare" (unmitigated) shock amplitudes, where mitigating policy measures are indeed needed to prevent a catastrophic collapse of output beyond a certain tipping amplitude of the initial COVID shock alone.

9.4. THE COVID SHOCK AND ITS AFTERMATH IN THE ABSENCE OF MONETARY POLICY




Figure 9.4 – Counterfactual COVID shock in the Inactive Central Bank scenario and in the absence of any policies. Unemployment (left) and inflation (right) in the Inactive Central Bank scenario without any policy, with an amplified COVID shock by a factor κ and Supply Chain shock. $\kappa = 1$ (blue line) corresponds to the COVID shock considered in the previous section, with a very low level of unemployment. With $\kappa = 1.1$ (red line), the unemployment and inflation dynamics change only very slightly, as the shock is mild. As soon as $\kappa \geq 1.2$, the economy is unable to recover spontaneously, and output collapses as demonstrated in the $\kappa = 1.2$ case (yellow line).

Sensitivity to Shock Magnitude and the Role of Easy-Credit Policy

Our model is based on data that already incorporates policy measures, which were indeed in place in 2020-21 to mitigate shocks. As such, it is possible that the unmitigated shocks were actually more severe than what we observe retrospectively in data. Here, we investigate counterfactual (ii)-scenarios with different shock magnitudes, where we shut down the final energy shock but scale the amplitude of the observed COVID shock by a factor κ (see Figure 9.2, left panel dark gray area). This stronger shock could have taken place if for example the CARES act and the CAA were not authorized.¹⁹ As κ increases, firms go progressively bankrupt due to their high financial fragility Φ , leading to a higher unemployment rate during the shock. Households save money during the shock, resulting in higher demand, leading to upward pressure on prices and somewhat higher inflation. Therefore in the long run, the scenarios with intensified shocks $\kappa \lesssim 1.2$ very slightly increases inflation and long-term unemployment. However, for larger κ , a tipping point is reached, as in Sharma et al. (2020): more initial bankruptcies and firms' hesitation to increase wages after the shock leads to a collapse of output and deflation, see Figure 9.4 (for a full dashboard see Figure H.1 in Appendix H.2). Eventually, the shock is too strong for the economy to recover, and the negative feedback loop of a downward wage spiral causes the economy to collapse.

The default threshold of firms Θ , holds a pivotal role in the economic downturn; if chosen

¹⁸We consider here a direct price-shock transmission. In practice, energy is purchased in various forms and often with inventories and financial contracts to insure against price volatility. In this respect, it might make more sense to define the effective price of energy $p_e(t)$ in the production process as an exponential moving average of the WTI spot price. This would smooth out the initial dip and lead to more realistic inflation time series. We leave this for a later study.

¹⁹The supply chain shock is left unchanged in the present exercise, but we observe almost the same phenomenon without such additional shock.

CHAPTER 9. POST-COVID INFLATION & THE MONETARY POLICY DILEMMA




Figure 9.5 – Economic dashboard for a stronger COVID shock in the Inactive Central Bank scenario, with and without Easy-Credit policy. The dynamics of the Inactive CB scenario in case (ii) (COVID and supply chain shock but no energy price shock), for a COVID shock strength of $\kappa = 1.2$ without Easy-Credit policy (blue) and $\kappa = 1.2$ with Easy-Credit policy, $\mu = 1.3$ (orange). We see that policy is effective in preventing the economic collapse in the case of the stronger COVID shock.

too low an excessive number of firm bankruptcies ensues, which can be detrimental to the overall economy. In this case, allowing firms to accumulate more debt by easing bankruptcy rules is highly effective to keep the economy on an even keel. In fact, as shown in Figure 9.5, the “Easy-Credit” policy defined by Eq. (9.28) with $\mu = 1.3$ for the bankruptcy threshold Θ manages to substantially reduce the impact of strong COVID-related shocks, even with an increased initial COVID shock ($\kappa = 1.2$) that would collapse the economy on its own. With Easy-Credit policy, firms can maintain wages and do not need to fire employees to remain solvent, allowing the economy to recover. However, unemployment still peaks above 8% (with a second hump mid 2024) and inflation reaches 16.3% at the end of 2022. As we will discuss in section 9.7 below, the main issue with such a period of high inflation is the risk of de-anchoring inflation expectations, opening the path to possible hyper-inflation.

In order to control inflation, one needs to consider the effect of monetary policy. We thus now turn to the study of the same sequence of three non-amplified shocks, with Easy-Credit policy and with a fully active Central Bank, distinguishing between the case of Anchored Trust and the case of Floating Trust in its ability to curb inflation.

9.5 Monetary Policy Response to Inflationary Shocks

In the previous section, we showed that the loosening of regulatory bankruptcy policy during severe shocks prevents the Mark-0 economy from collapsing, but does not address the issue of high inflation rates (Section 9.4). To study inflation mitigation policies, we now introduce a Taylor-rule based monetary policy into the mix, i.e. we combine the situation explored in Section 9.4 (all shocks and Easy-Credit policy without Central Bank) with the different monetary policy scenarios discussed in Section 9.3. All numerical experiments henceforth are run with the Easy-Credit policy described in Eq. (9.28) with $\mu = 1.3$, as in Section 9.4. As the economy returns to its steady state, the Easy-Credit policy becomes equivalent to the fixed default threshold, such that its duration depends on the recovery period.²⁰ Note that in what follows, we consider the three shocks together, keeping the amplitude of the COVID shock to the one calibrated on the observed data (i.e. $\kappa = 1$ in the language of Section 9.4). As already discussed, such calibration may underestimate the severity of the unmitigated shock, since the US government immediately implemented Easy-Credit and Helicopter Money measures to alleviate the COVID shock. Note also that unemployment is not an explicit target of our Central Bank, in the sense that the Taylor-rule is only responsive to realized inflation and not to unemployment.

An Inactive Central Bank

We begin our study by considering again the situation with an Inactive Central Bank as in Section 9.3, but now with all shocks together and with the Easy-Credit policy being implemented (Figure 9.6, full blue line; see Figure H.5 in Appendix H.3 for all macroeconomic time series). In this case, over the first few months, the combination of shocks causes a spike in unemployment that peaks around 8% during a deflationary phase, but the economy quickly recovers to full employment, yet at the price of a sustained inflationary period that peaks above 20% before slowly reverting to equilibrium, consistently with the discussion of Section 9.4: the Easy-Credit policy can avoid full collapse, but only at the price of high inflation.

Monetary Policy with Unanchored Trust

Before fully introducing trust, we consider the case where agents form their inflation expectations without considering the Central Bank's target inflation, but the Central Bank is nonetheless actively pursuing monetary policy. We thus set $\tau^T = 0$ such that inflation expectations $\hat{\pi}(t)$ are simply an exponentially weighted moving average over realized inflation. In this case, we find that the Central Bank initially decreases rates before rapidly increases interest rates as inflation increases (see Appendix H.4 for the relevant dashboard). However, there is no resulting effect on inflation, which displays the same dynamics as the case where the Central Bank is inactive. However, the de-anchored

²⁰Eq. (9.28) indeed implies $\Theta(t) = \Theta_0$ in the steady state when the average fragility is low, which is the case for the choice of parameters made in this chapter.

expectations lead to a doubling in magnitude of the initial spike in the unemployment rate, together with a higher (and oscillating) peak following the end of the shocks, as compared to the case where monetary policy is inactive. We can thus draw a first conclusion that in the absence of any sort of expectation anchoring the effects of monetary policy are concentrated on the unemployment rate, which is not actually the Central Bank’s target in our model setup.

Monetary Policy with Anchored Trust

We next consider a monetary policy experiment with a responsive Central Bank in the “Anchored Trust” scenario. In the Mark-0 model, as in real life, the Central Bank’s aim is to keep inflation on target. Here, following Section 9.3, the inflation target is chosen to be $\pi^* = 2.4\%$ p.a. and the Taylor rule strength is $\phi_\pi = 1$. Moreover, we optimistically assume that the Central Bank has successfully communicated to all economic actors that it has inflation under control, thereby convincing them to believe that long-term inflation will be close to the Central Bank’s inflation target, with a fixed anchor weight $\tau^T = 0.95$.

In the case of Anchored Trust and variable interest rates with all three shocks (Figure 9.6, dashed orange line), the Central Bank is successful in taming inflation: peak inflation is lower than in the case without an active central bank (peak at 8.8% vs 23.2% in the Inactive Central Bank scenario) and quickly tapers off. Although the inflation rate is not in line with the Central Bank’s target throughout the crisis, it appears that monetary policy is able to reign in inflation. The reduction in inflation is not primarily due to the impact of interest rate policy but rather to the strong anchoring of expectations, which significantly dampens expected (and thus realized) inflation, as upward pressure on both prices and wages are reduced. Such anchoring moderates the price increases after the initial COVID shock, but does not dampen the deflationary dynamics towards the end of the energy price shock, as this is driven by a strongly increasing unemployment that puts downward pressure on wages (see Figure H.5 in Appendix H.3 for all macroeconomic time series).

The initial price shock in the first few months causes bankruptcies that are less numerous with a reactive Central Bank, hence also reducing the initial unemployment spike slightly (peak at 8.1% in the Inactive Central Bank scenario and 7.4% with Active Central Bank with Anchored Trust). However, at the end of the shock the Central Bank effort to control inflation results in a peak unemployment of 12.8 % in Feb 2024, 5% higher than in the Inactive Central Bank scenario. Despite realized inflation surging during shocks, firms inflation expectations are anchored to the central bank’s target. This results in a substantial decline in real wages during the shock and a sluggish recovery thereafter. Similarly, households expect a controlled level of inflation and hence reduce consumption in response to higher interest rates that favor savings. Both effects lead to higher unemployment, impeding the overall post-shock recovery. We see that in the Mark-0 model, the Central Bank is stuck between a rock and a hard place with its inflation-unemployment trade-off, meaning that in the context of these shocks, monetary




Figure 9.6 – Unemployment and Inflation with all shocks, Easy-Credit policy, and three distinct monetary policy scenarios. (Blue lines) Inactive Central Bank scenario. (Orange lines) Reactive Central Bank with Anchored Trust, with anchor parameter $\tau^T = 0.95$. A Taylor rule policy successfully decreases peak inflation but increases peak unemployment. (Green lines) Reactive Central Bank with Floating Trust. Here, trust is eroded during the high inflation period. Monetary policy then fails at reducing inflation, with the risk of hyper-inflation lurking (see section 9.7); unemployment remains lower because of the higher inflation.

policy is not a panacea: some form of Keynesian stimulus or “Helicopter Money” is needed in conjunction with monetary policy to restore the economy to its pre-crisis steady-state – see section 9.6 below. In this regard, a Central Bank with an explicit dual-mandate would potentially navigate the trade-off better.

Monetary Policy with Floating Trust

The scenario of high and Anchored Trust in monetary policy for long-term inflation expectations might be a good approximation during long periods of stability, but in periods of shocks, the anchoring τ^T may decrease when inflation deviates strongly from the Central Bank’s target (Reis, 2021). We model such a Floating Trust effect through Eq. (9.27), where the anchoring parameter decreases when the spread between observed and target inflation increases. We set $\alpha_I = 0.4$ and a memory time of ~ 5 months ($\omega = 0.2$). These values imply that when inflation reaches 4 times the Central Bank’s target π^* , trust in the Central Bank falls to approximately a third of its initial value after approximately one year. Larger values of α_I would lead to an even steeper de-anchoring of inflation expectations. Such a loss of faith in monetary authorities further increases realized inflation as economic agents expect higher inflation in the future and take this into account when setting prices and wages. This may lead to a self-fulfilling feedback loop between expected inflation and real inflation through a wage-price spiral (see section 9.7).

Compared to the Anchored Trust case, the surge of inflation after all COVID shocks leads to a loss in trust that almost vanishes (i.e. $\tau^T \rightarrow 0$) during the energy price shock. The consequence of this is a higher realized inflation rate, that reaches a peak value only slightly below the Inactive Central Bank scenario (Figure 9.6, right).²¹ Un-

²¹For the full economic dashboard in the case of dynamic expectations, see Appendix H.3, Fig. H.4,

employment, on the other hand, turns out to be *higher* than with an Inactive Central Bank (see Figure 9.6, left), but *lower* than when trust is anchored. As noted above, low expected inflation is detrimental to wages and to consumption when interest rate are high. Therefore, perhaps paradoxically, loss of trust improves the situation in terms of unemployment, but is of course detrimental to inflation, with the possibility of runaway situations. This again illustrates the quandary faced by Central Banks and the importance of a dual-mandate.

A strong factor of the loss in trust is the exogenous nature of the energy price movements. Despite the strong hike up to a real interest rate of 3.4% per year (16.8% nominal target rate), inflation remains far above target. Once these shocks have passed, the Central Bank is arguably quite successful in keeping inflation closer to its steady state value. However, similarly to the Anchored Trust case, the economy now contends with a persistent, relatively high unemployment rate. The unemployment here is driven by the energy price shock that causes a strong surge in inflation, which erodes household savings and reduces real wages by 4% compared to its steady state value, thus resulting in a drop in demand (see Figure H.4 in Appendix H.3). In contrast to the Anchored Trust scenario, unemployment in the Floating Trust scenario recovers to full employment faster. This paradoxical result is driven by the comparatively higher real wages, which allows for a faster recovery of the consumption budget and consequently output. This implies that a loss of trust may be beneficial in terms of unemployment by promoting a faster adjustment to exogenous price levels as they occur – in a sense, a favorable aspect of a tight wage indexation, provided of course it does not spiral out of control.

As with both the Inactive Central Bank and Anchored Trust scenario, strong exogenous inflation outside of the Central Bank's sphere of influence leads to a severe contraction of the economy, which recovers only slowly. The convergence to pre-crisis equilibrium takes more than 18 years in the Anchored Trust scenario and 11 years in the Floating Trust scenario. This suggests that in the face of exogenous price shocks, fiscal policy shoring up consumer's budget or reducing energy price effects on firms and households can be a more effective tool than monetary policy.

Sensitivity to the Strength of the Monetary Policy

In circumstances of crises, as with our shocks, Central Banks should possibly react more strongly to inflation and keep expectations anchored. We thus consider a larger Central Bank Taylor reactivity parameter $\phi_\pi = 2.0$ (see Appendix H.5, Figure H.7 and H.8 for details).

In the case of Anchored Trust, an increase in ϕ_π hardly changes inflation, but stronger interest rate hikes result in significantly higher unemployment rates. Hence, stronger monetary policy is clearly detrimental in this case. By contrast, in the Floating Trust scenario with dynamic trust, an increase in ϕ_π signals a stronger commitment of the

and compare with Fig. H.3.

Central Bank to control inflation. Therefore, all else being the same, inflation expectations are lower which, in a self-fulfilling manner, cuts realized inflation from 16.8% to 11.2% and keeps trust better anchored. However, as in the Anchored Trust case, the sharper interest rate hikes by the Central Bank lead to a significant increase in the unemployment rate (from 12.8% to 20.5%) and a much longer recovery time for the economy as a whole. Again, our model confirms and quantifies the observation made by Stiglitz and Regmi (2022), quoted in the introduction: monetary policy interest tools tend to be too blunt, curbing inflation at the cost of unnecessarily high unemployment.

Sensitivity to Transmission Channels

Within our modeling framework, the impact of monetary policy relies on the efficiency of three transmission channels: (a) expectation anchoring, discussed above; (b) sensitivity of consumption on real interest rate, through parameter α_c ; (c) sensitivity of firms' hiring and wage policies on real interest rate, through parameter α_Γ . We have run some simulations to check the dependence of the state of the economy on these last two parameters in the Anchored Trust case, see Appendix H.6. The conclusion of this exercise is that increasing α_c and α_Γ has a minor direct effect on inflation but a significant effect on unemployment, depending on the sign of the real interest rate $\rho - \hat{\pi}$. As expected, a raise of interest rates degrades economic activity, all the more so when the sensitivity of firms and households to the real rate is higher. As already stated, the main transmission channel through which monetary policy impacts inflation is expectation anchoring, rather than directly through cost of loans or income on savings.

9.6 Fiscal Stimulus

Section 9.5 showed that in the presence of exogenous inflation drivers, the Central Bank can reduce inflation only at the cost of high unemployment. The mechanisms responsible for unemployment are the increased burden of debt that weighs on wages and increases the speed at which firms lay off workers. This in turn lowers demand, an effect amplified by the erosion of household's purchasing power by the lingering inflation. One way to tackle this issue is to ensure that demand does not drop as severely by increasing it through fiscal stimulus (see Section 9.2). We test this policy device here by considering different kinds of fiscal stimuli. Unless otherwise indicated, the fiscal stimuli are applied on the scenario of Section 9.5, i.e. with all shocks, Easy-Credit policy, and an Active Central Bank with Floating Trust, which we believe to be the scenario closest to reality.

The Effects of Helicopter Money

Following Sharma et al. (2020), we begin by considering a one-time increase of the households' budget by a factor $\kappa_H \in [0\%, 60\%]$ one month after the end of the energy price shock (July 2023). The goal is to find an optimal multiplier κ_H that reduces unemployment while keeping excess inflation to a minimum in amplitude and duration. Here we find that this minimum is reached for $\kappa_H \approx 20\%$ (see Figure 9.7 for the Reactive

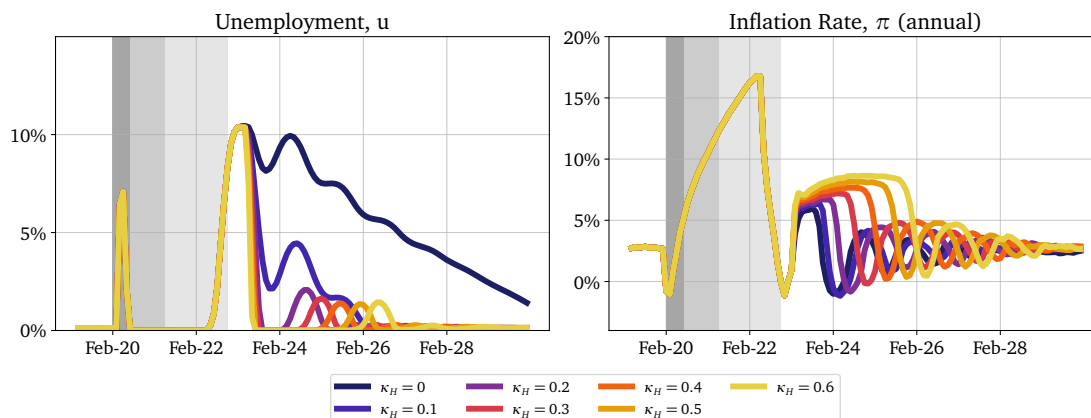


Figure 9.7 – Helicopter Money in the Reactive Central Bank with Floating Trust scenario, Easy-Credit policy, and all shocks. Unemployment (left) and inflation (right) for a helicopter drop of size κ_H times savings one month after the price shock. Already with $\kappa_H \geq 0.2$ unemployment is reduced almost to zero. A further increase of Helicopter Money only increases the duration of the high-inflation period, without further reducing unemployment. Note the small unemployment and inflation “ripples” persisting several years after the initial shock.

Central Bank with Floating Trust case).²² This choice of κ_H matches the magnitude of the “Emergency Money for the People Act”, which entailed the US government providing direct payments of 2000 USD per month for a maximum of 12 months to support individuals during the COVID pandemic.

In all considered scenarios, the fiscal stimulus package leads to a quick recovery of production to its pre-shock levels and unemployment at near 0%, thus successfully eliminating the steep recession and long recovery following the energy price shock, as shown in Figure 9.7. However, policymakers face an inflation-unemployment trade-off, as the stimulus generates a spurt of inflation. In the absence of monetary authority, inflation due to the injection of money rises to 15.8% (see Appendix H.7, Figure H.11). However, this is an endogenous inflation, which means that in the Floating Trust case, the Central Bank can raise rates to curb consumption propensity and keep inflation below 10% p.a. (Figure 9.7). Increasing the stimulus leads to a higher inflation peak and a longer duration of high inflation.

The inflationary period after the stimulus is in part due to a short-term de-anchoring of expectations as the demand stimulus leads to higher prices in the context of a tight labour market preventing production increases. Yet, here one can take the position that the economy is in good shape from a macroeconomic perspective, with near-zero unemployment despite an above-target inflation.²³ In this case, the Anchored Trust case

²²See Appendix 9.6, Figures H.13, H.12 for the other two monetary policy scenarios, and Figures H.13, H.14 and H.15 for the complete economic dashboards of all scenarios.

²³Because Mark-0 treats households at the aggregate level, we do not analyze the distributional consequences of inflation here, though inflation is always and everywhere differential in its effects.

would be much more favorable as the high unemployment problem is resolved with only a minimal rise in inflation, around 7% annualized over 8 months, as soon as $\kappa_H \gtrsim 0.2$.

The Effects of a Windfall Tax




Figure 9.8 – Windfall Tax in the Reactive Central Bank with Floating Trust scenario, Easy-Credit policy, and all shocks. Unemployment (left) and inflation (right) for a Windfall Tax of $\delta_e + \Delta\delta_e$ one year before the end of the price shock with a duration of two years. With $\Delta\delta_e \approx 4\%$ unemployment is reduced strongly. A further increase of tax does only increase unemployment again. For even larger $\Delta\delta_e$, inflation increases because of increased demand due to increased savings.

Another stimulus that central authorities can use to alleviate the impact of the price shock is a Windfall Tax. This policy redistributes the excess profits generated by the energy sector as a result of rising energy prices to household savings with the aim to stimulate consumption and consequently reduce unemployment. We test the effectiveness of the policy, which in our modeling consists in an increase of the energy sector payout rate $\delta_e \rightarrow \delta_e + \Delta\delta_e$ one year before the end of the price shock with a duration of two years, with $\Delta\delta_e \in [0\%, 12\%]$.²⁴

Our aim is to design the fiscal stimulus in a manner that minimizes both unemployment and inflation. We find that for $\Delta\delta_e = 2\% - 4\%$ the excess profits of the energy sector are redistributed such that unemployment is significantly reduced at all times while inflation remains under control. Further increasing $\Delta\delta_e$ reduces the unemployment in 2023, but at the cost of increasing unemployment in 2024 after the Windfall Tax has ended (see Figure 9.8).²⁵ As long as the tax is active, the increase in δ_e results in higher savings, leading to an increased consumption budget and therefore higher demand. This has the effect of reducing unemployment as firms expand production, as well as reducing firm fragility due to increased profitability. However, as the Windfall Tax increases in the

²⁴We remind here that δ_e is the fraction of the energy sector cash balance redistributed to households at each time step, see Eq. (9.20).

²⁵For more detail, see Appendix H.8, Figures H.16, H.17.

Floating Trust scenario, inflation also increases and expectations remain de-anchored for longer, compelling the Central Bank to increase interest rates. As the tax ends, the household’s consumption budget has been reduced to below the case without a Windfall Tax, and the real interest rate remains high. This leads to a negative spiral of lower demand and output, higher unemployment, leading to lower income, until the system equilibrates again.

The inflationary effects of the increase in savings due to the Windfall Tax remain minor for small values of $\Delta\delta_e$. However, once the Windfall Tax exceeds $\Delta\delta_e \sim 10\%$, a longer period of high inflation after the shock is observed (see Figure 9.8 and Appendix H.8, Figure H.20) as inflation expectations remain de-anchored for a longer period.

The Anchored Trust scenario presents a different challenge. Again, up to a threshold, the Windfall Tax reduces the initial and long-term unemployment spike, with $\Delta\delta_e = 4\%$ still remaining a robust choice (see Appendix 9.6, Figures H.17, H.19). However, this time at the cost of inflation (or less deflation) on the way back to the pre-crisis steady state. In this regard, policymakers might opt instead to minimize the volatility of inflation to let it smoothly return to the Central Bank’s target rate, which actually suggests a higher Windfall Tax of $\Delta\delta_e = 8\%$ for instance.

We conclude that when implementing a Windfall Tax as a fiscal stimulus, excessive tax levels can have the unintended consequence of exacerbating long-term unemployment, while a measured approach can both reduce unemployment and smooth inflation dynamics.

9.7 Model Sensitivity: The Dangers of a Wage-Price Spiral

To develop a deeper understanding of the reaction to shocks of our model economy, we explore the parameter sensitivity of this configuration of the Mark-0 model using the “sloppy model” methodology put forth in Part I of this Thesis.

With respect to both inflation and unemployment, the balance of wage indexation (bargaining power) g_w to price indexation (market power) g_p , has by far the strongest influence on the outcome dynamics for the inactive CB and Floating Trust scenarios (see detailed results in Appendix H.9 and H.10). Beyond inflation itself, eigenvectors are combinations of various pricing parameters including γ , g_p and g_e for both unemployment and inflation. Central Bank parameters are present but not significantly. This aligns with our findings that the Mark-0 economy always maintains a tight connection between unemployment and inflation, indicating that it may not be possible to minimize one without affecting the other (see Appendix H.9, Figures H.23, H.24).

Gualdi et al. (2015) have demonstrated the significance of the default threshold of firms Θ in determining the economic phase of the Mark-0 model. This parameter exhibits a non-linear behavior, where minimal alterations have negligible effects on the overall dynamics until a critical “tipping point” is reached. Beyond this threshold, the dynamics undergo substantial transformations. It is important to note that with the “sloppiness”




Figure 9.9 – Hyperinflation tipping points. Left: Stable inflation vs. Hyperinflation in the plane ($g_p = g_w, \tau^T$) in the Anchored Trust scenario and in absence of exogenous shocks. Note that inflation can be stable even when $g_p = g_w > 1$ for strong enough anchoring. Middle: Stable inflation vs. Hyperinflation in the plane ($g_p = g_w, \alpha_I$) in the Floating Trust case and in absence of exogenous shocks. Right: inflation rate (top) and unemployment (bottom) as a function of time, in the Floating Trust case, $\alpha_I = 0.3$ and different values of $g_p = g_w$, after the three COVID shocks. When indexation is too strong, an hyperinflation regime sets in at the end of the energy price shock.

analysis, we can only identify local sensitivity in parameter combinations, where slight deviations from these combinations result in noticeable changes in observed dynamics. As already found in Naumann-Woleske et al. (2023), multiple phase transitions occur along the parameter axis of Θ , although minor perturbations do not alter the curvature of the loss function, therefore we do not find Θ to be crucial in our results. However it is crucial to distinguish between local sensitivity, which can be determined by the sloppiness approach, and global sensitivity to parameters, which can trigger tipping points in the model Gualdi et al. (2015).

The Risk of a Hyperinflation Spiral

We can interpret the sensitivity to g_w and g_p in light of the dangers of a hyperinflation episode resulting from a wage-price spiral. Alvarez et al. (2022) cite a concern that hyperinflation may occur if firms increase wages in response to higher inflation, leading to an increase in purchasing power and ultimately feeding into a wage-price spiral in the current macroeconomic environment. This feedback loop is influenced by the indexation of prices and wages to firms' inflation expectations (Holland, 1988), i.e. by the value of parameters g_p and g_w in our model. In the simplest case where bargaining power and market power are equal ($g_p = g_w$), the economy reaches a stable inflationary state, marked by cyclical fluctuations due to mismatches in supply and demand (see Figure 9.1 in Section 9.1, for which $g_w = g_p = 0.8$). When indexation is weak ($g_p = g_w < 1$), neither wages nor prices fully incorporate inflation expectations, and fluctuations in the inflation rate due to mismatches in demand and supply are dampened. Conversely, for

strong indexation ($g_p = g_w > 1$), the economy may enter a state of hyperinflation, though as wages and prices increase equally fast, there are no real effects (omitting, of course, the cost of inflation itself due to “menu costs”). Only when $g_p \neq g_w$ does the economy collapse.

Introducing monetary policy implies that hyperinflation can be staved off due to the anchoring of inflation expectations. In the case of Anchored Trust, inflation remains stable until a critical point $g_p = g_w < g^*$ where $g^* > 1$ depends on the commitment of the Central Bank as well as the strength of expectation anchoring to the its target (see Figure 9.9, left panel). The same holds true with Floating Trust (middle panel). However, in this case an interesting range of parameters where inflation remains stable for $g^* > g_w = g_p > 1$, until a strong enough shock occurs, which triggers a loss of trust in the Central Bank, tipping the economy into a hyperinflation phase (as illustrated in Figure 9.9, right panel). This is precisely the scenario that Central Bankers want to avoid by doing “whatever it takes”.

If bargaining and market power differ ($g_p \neq g_w$), the effect on unemployment is significant (see Appendix H, Figure H.26). In the Anchored Trust case, when market power is greater than bargaining power – a.k.a. “greedflation” – ($g_p > g_w$), unemployment increases after the shock and only slowly decreases as wages adjust at a slower pace to pre-crisis levels, leading to sluggish demand recovery. Consequently, firms do not see the need to increase their production faster. In contrast, if bargaining power exceeds market power ($g_w > g_p$), wages recover more quickly, resulting in an earlier recovery of pre-crisis purchasing power, which helps to reduce unemployment. This conclusion is in line with the result of section 9.6 on the impact of Helicopter Money or, more generally, of Keynesian stimulus: increasing the consumption budget of households is quite efficient at reviving an ailing economy, but only if hyperinflation can be avoided.

9.8 Summary & Conclusions

In this chapter, we have expanded the Mark-0 Agent-based Model to assess the inflationary dynamics following the COVID pandemic and energy crisis of 2020-2023. Our results highlight the narrow path of monetary policy to balance unemployment and inflation, as well as the benefits of joint monetary and fiscal policy packages.

We extended the Mark-0 model used in Sharma et al. (2020) in two directions (Section 9.1): First, a dynamically evolving trust in the Central Bank’s ability to control inflation, such that expected inflation remains anchored to the Central Bank’s target if trust is high, but persistent off-target inflation realisations lead to a loss of trust in the Central Bank. Second, a simple “exogenous” energy sector from which firms buy energy that allows us to introduce an energy price, and slowly re-inject the energy profits into the economy as dividend payouts. This slightly expanded Mark-0 model allows us to investigate several policy channels (Section 9.2): (a) monetary policy via the management of interest rates and expectations, (b) an Easy-Credit regulatory policy, and fiscal policies such as (c) Helicopter Money and (d) a Windfall Tax.

The Mark-0 steady state economy is then perturbed by three calibrated shocks (Section 9.4): a consumption propensity drop (representing lockdowns), a productivity drop (representing supply chains disruptions), and an energy price shock (first a drop, representing reduced demand during lockdowns, followed by a rise due to demand recovery exacerbated by the Russian invasion of Ukraine). We show in Section 9.4 that the Easy-Credit policy, which was one of the first emergency responses to the COVID pandemic, is able to alleviate the impact of the shocks, but at the price of high and sustained inflation, as observed in the data and predicted as early as June 2020 by Sharma et al. (2020) within the Mark-0 framework.

We investigated whether monetary policy alone can control a surge of inflation, and at what cost to unemployment (Section 9.5). Our results show that if agents' place no trust in the Central Bank, there is no benefit to monetary policy to tame inflation. On the contrary, if agents' expectations remain fully anchored, inflation remains closer to the Central Bank target, but at the price of a strong recession leading to a wave of unemployment. Meanwhile, if agents' expectations evolve dynamically, inflation rises well above target during the shock leading to de-anchored expectations. Monetary policy is then essentially ineffective in controlling inflation while causing unemployment to increase in comparison to a baseline case with no monetary policy at all. This is robust to variations in parameters related to the channels of monetary policy. Our framework thus supports the view that the efficacy of monetary policy comes at the cost of unnecessarily high unemployment (Stiglitz and Regmi, 2022), in particular, if inflation expectations are strongly anchored (see Section 9.5).

Within Mark-0, trust anchoring is thus a crucial determinant of the success of the Central Bank inflation mitigation policy, far more important than the direct economic impact of higher interest rates (noting that we do not model in detail a financial sector). This resonates with Bernanke's statement: "Expectations matter so much that a Central Bank may be able to help make policy more effective by working to shape those expectations",²⁶ and points to the importance of *narratives* in shaping expectations (Shiller, 2019). In turn, fine-tuned fiscal policy combined with monetary policy can be successful in controlling inflation while keeping unemployment around acceptable levels (Section 9.6). However, this requires a high degree of precision to be effective: too weak a fiscal stimulus is ineffective, and too large a stimulus leads to further high inflation. Finally, in Section 9.7 we show that the pricing power of firms and the bargaining power of workers play a crucial role. Depending on their relation, the economy can experience a runaway hyperinflation wage-price spiral. This points to an additional difficulty in properly calibrating monetary and fiscal policy due to a variety of possible tipping points or "dark corners" lurking around.

Overall, this study of the Mark-0 economy shows that (i) the economic recovery can be very sluggish due to self-fulfilling expectations and other non-linear feedback loops,

²⁶see Bernanke, B. (2013), "Communication and Monetary Policy", speech at "National Economists Club Annual Dinner".

(ii) there is a tension between inflation and unemployment that is robustly captured by the model, (iii) in the Mark-0 economy, monetary policy is effective at controlling inflation provided trust is anchored (but not because of real economic effects of higher interest rates), and (iv) fiscal policy can alleviate some of the negative unemployment effects of inflation-focused monetary policy. While our study did not exhaust all possible cases and parameters, we believe that it illustrates both the realism and the richness of the Mark-0 economy. It offers a versatile tool with which policy makers can easily play in order to forge their intuition about what may happen if they turn this knob or add that policy measure. Indeed, almost all possible narratives that emerged during the recent debate about post-COVID inflation can be captured and reproduced by our model with proper choices of parameters and shock specifications. Furthermore, the procedure proposed in Section 9.7 can be used to rigorously establish which predictions of the model are robust to the choice of parameters, most of which are difficult to properly calibrate on data.

Our model is obviously incomplete and improvable on many counts. In particular, it could be extended to include (i) considering a dual mandate central bank as in Bouchaud et al. (2018); Gualdi et al. (2017) to address negative unemployment side-effects, (ii) a disaggregated household sector (e.g. distinguishing wage and rent earners, and by accumulated wealth) to assess distributional implications of policy, (iii) production networks and energy supply chains with differentiated products and commodities (as in Dessertaine et al., 2022, for instance), which introduces firm distribution effects and systemic risk, (iv) a financial sector to more closely model interest rate pass-through, lending choices and financial systemic risk, and (v) more detailed labour market structures to account for the change in occupational structure post-COVID and demographic trends. However, we do believe, as we already argued in Gualdi et al. (2015), that the Mark-0 model should already be part of the toolkit of Central Banks, if only as an inspiring *scenario generator*, or “telescope for the mind”, especially in times of great modelling uncertainty during which it is crucial to be at least “roughly right” and avoid being blindsided by spurious Black Swans.²⁷

²⁷See for example Mark Buchanan, “This Economy Does Not Compute”, New York Times, October 2008, Gualdi et al. (2015); King and Kay (2020); Bouchaud (2021), or the recent Financial Times piece cited in the introduction, <https://www.ft.com/content/b972f5e3-4f03-4986-890d-5443878424ac>.

Key Messages

- This chapter extends the Mark-0 MABM to incorporate dynamic trust in the Central Bank and an energy sector before perturbing it with shocks to the consumption propensity, productivity and energy prices.
- With only an easy-credit policy, the model suggests high and sustained inflation following the COVID and energy-price shocks.
- Depending on expectation anchoring, monetary policy may be ineffective (de-anchoring) at controlling inflation, and may exacerbate unemployment effects.
- By contrast, adding fiscal policy such as direct transfers, can ameliorate both inflation and unemployment if the amount is calibrated well.

Conclusion

Chapter 10

Summary of Results

This chapter serves as a brief summary of what has been presented in this thesis across the topics of parameter space exploration and the effects of confidence and collective behavior on macroeconomic dynamics.

Exploration of the Parameter Space

Part I of this thesis developed a preliminary approach to addressing the questions of *What is the set of phenomena that an Agent-based Macroeconomic Model can generate? And what determines their dynamics?* To address this question, I have taken inspiration from work in biophysics. Specifically, the finding that models with high-dimensional parameter spaces can nonetheless be well-fitted with little data and make good predictions. This finding, dubbed *sloppiness*, is based on the observation that the observations of these models are very sensitive to only on a handful of well-constrained (stiff) combinations of parameters. The mechanics of this method are developed in Chapter 3 and demonstrated on the simple example of Kirman's Ants (Chapter 4). In Chapter 5, this method is applied to the Mark-0 model, showing that this macroeconomic agent-based model has a sloppy parameter hierarchy, specifically, that the observed dynamics of the unemployment rate depend on only one or two well-constrained (stiff) parameter combinations, while the remaining combinations have little effect on the outcomes. This alone is a central result of this thesis, answering the question of what drives a given dynamic. Moreover, it has critical implications for the calibration of ABMs. If a given dynamic generally depends only on a handful of parameter combinations, then as long as these combinations are well-constrained by data, the models should be able to make good predictions (at least within sample). Moreover, these combinations often involve many parameters implying that individual parameters can have extremely large confidence intervals as long as their combination is constrained. Given the criticism of MABMs as being hard to calibrate, this result in particular should be useful to researchers working on calibration techniques, suggesting a potential move away from attempts to constrain

all individual parameters to narrow confidence intervals to a focus on fitting the collective behavior and stiff parameter directions.¹ This approach to the construction of models should make future exploration and calibration of ABMs significantly easier, though perhaps not necessarily faster.

Taking the results of the sloppy analysis one step further, in Chapter 6 I develop a simple parameter space exploration heuristic and apply it to the Mark-0 model. The principle is simple, in a gradient-ascent style approach, one perturbs the model's parameters in the well-constrained sensitive directions to generate a set of parameters with maximally different dynamics. In the case of the Mark-0 model, this heuristic can recover all of the model's phases as defined by Gualdi et al. (2015), and more, therewith providing an approach to answering the first research question on determining the set of dynamics. Introducing a self-consistent metric of algorithmic efficiency, one can see that the heuristic can recover a variety of different dynamics outside of the a priori phase definitions. Studying the set of directions in parameter space followed by all exploration attempts suggests that the heuristic traverses a much lower-dimensional set of parameters to vary the unemployment rate in the Mark-0 model, responding to the second research question on the drivers of dynamics but also pointing to possible model reductions to elicit key mechanisms in Mark-0. This simple heuristic is not free of limitations and requires further research to improve its robustness. First, there are limitations inherent to the choice of a mean-squared loss function, such as the sensitivity to small changes in frequency of an oscillation. Second, it is unclear what the most effective strategy is for determining the sign of the eigenvector or the exact distance to travel in the direction of the eigenvector.

Confidence and Collective Behavior

Part II of this thesis, presents three different macroeconomic models incorporating confidence effects and collective behavior. The first of these, the Dynamic Solow Model (Chapter 7), together with the extended Real Business Cycle model (Chapter 8), focus on the transitions of the economy between two attractors, a positive and negative one, due to changes in the level of confidence in the state of the economy. For the Dynamic Solow Model, we build an interactions-based process from the ground up to derive a set of non-linear differential equations describing the aggregate confidence, information and economic dynamics. Tuning this system leads to business cycle-like fluctuations due to a coherence resonance phenomenon where noise, here in the form of exogenous news, applied to a dynamical system can lead to a quasi-periodic response. Admittedly, the economic closure in the form of a Solow model is a very simplified representation that omits many of the additional drivers of business cycles, notably the financial system. Turning to the extended Real Business Cycle model of Chapter 8, here too we incorporate an explicit term for confidence into economic decision-making on consumption and investment. Here the confidence depends on the state of agents' neighbors, leading to

¹Data requirements for each direction scale with $1/\sqrt{\lambda_i(\Phi)}$, implying extreme data requirements for sloppy directions.

a “keeping up with the Joneses” effect, and as with the Dynamic Solow Model, it can be written as a non-linear aggregate agent. With the addition of capital as a factor in production, this model generates a variety of phases, with more or less frequent crises of various duration (short V-shaped spikes or long L-shaped recessions). As with the Dynamic Solow Model, the occurrence of these crises is endogenous and not dependent on any large exogenous shock pushing the economy away from equilibrium. Quoting Bernanke et al. (2019), we see a “small shocks, large cycle” effect arise naturally in both of these models.

Finally, Chapter 9 takes a slightly different approach to confidence, moving from the theoretical study of simple economies to assessing the post-COVID inflationary surge. Taking the Mark-0 model, we extend it to include three shocks resembling events of recent years: COVID-19 related lockdowns, supply chain disruptions, and energy price volatility. Additionally, we introduce confidence into the model in the form of the degree of anchoring of inflation expectations to the target of the central bank, together with a stylized energy sector. Studies done with the model indicate that in the absence of any policy, the recovery to the full-employment steady state following the shocks is sluggish and can last many years. The application of monetary policy by itself is akin to walking a tight-rope due to the inflation-unemployment tradeoff of raising interest rates. Critically, interest rate policy takes time to be effective, with an intermediate spike of untamed inflation, partially due to the exogenous nature of the energy prices driving inflation. In such a scenario, agents’ trust in the Central Bank diminishes temporarily, thus exacerbating the downturn. Not only is inflation eroding agents’ purchasing power, but their unanchored expectations lead to larger upward price and wage adjustments, that double down on the inflationary spike. In this case, we find a positive role to be played by fiscal interventions in the form of direct transfers to agents, as well as in the form of profit-taxes levied on the energy sector.

A key implication across all three models presented in Part II is the importance of confidence in the state of the economy in order to perpetuate this state, whether that is maintaining high confidence during expansions or remaining stuck in a low-confidence regime during recessions. From a policy perspective, this implies that messaging, such as that of the Central Bank with respect to its actions and inflation target, is of crucial importance to maintain a steady ship.

Conclusions and Perspectives

There are several ways in which the work presented in this thesis have already contributed to ongoing research into macroeconomic agent-based models and the macroeconomic implications of confidence and collective behavior (see Chapter 10 for a summary). However, in many ways these contributions are also the start of many future research directions. In this chapter, I will outline some of these new research directions. Some of these have already been mentioned throughout the preceding parts, but here I wish to expand on some of the areas that are of particular interest to me and which I believe to be perhaps the most impactful. Specifically, Section 11.1 outlines new avenues to investigate with respect to the estimation of Hessians, and the exploration and reduction of the parameter space in macroeconomic agent-based models. Beyond the methodological innovations, I am also interested in the wider applications of these exploration methods, such that phase diagrams of the various MABMs can be developed. In this regard, agent-based approaches to modeling the economy-environment-energy nexus are of particular interest to me. Section 11.2 presents a review of four Agent-Based Integrated Assessment Models (ABIAMs) that could serve as a starting point for this (the section is taken from Naumann-Woleske, 2023). Finally, Section 11.3 outlines some of the continuing research into the effects of different inflationary mechanisms in the Mark-0 model.

11.1 Diving Deeper into the Estimation and Exploration of MABM

Probabilistic Approaches: Working with the Fisher Information Matrix

The most natural extension to the work done in this thesis is to expand on the probabilistic interpretation of MABMs; such an approach should be more robust and also yield a much smoother model manifold as it may be less prone to noise. While the MSE is a good first approach, it has some drawbacks in its sensitivity to changes in oscillation. By contrast, working with the probability density function, a small change in frequency should yield a small change in the probability density. Chapter 3 intro-

duced the symmetrized Kullback-Leibler as the measure of choice, but one could also select the Hellinger or Bhattacharyya distance as a metric instead, as all of these metrics are f -divergences, implying that their Hessian is some transformation of the Fisher Information Matrix

$$g_{i,j}(\Phi) = \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{\partial \log P(\mathbf{x}|\Phi)}{\partial \Phi_j} \frac{\partial \log P(\mathbf{x}|\Phi)}{\partial \Phi_i},$$

whose estimation can also be troublesome. The principal reason for this being the possibility that $P(\mathbf{x}|\Phi) = 0$ implying that $\log P(\mathbf{x}|\Phi) \rightarrow -\infty$, which can be troublesome when approaching this estimation numerically. For example, applying a finite difference approach near a phase boundary might lead to points with non-zero density in one phase but zero density in another, meaning that one would obtain an ill-defined Hessian matrix as $\mathcal{L}^{sKL}(\Phi, \Phi + \delta\Phi) = \infty$, or one would have to discard this information. In an MABM context this occurrence may not be uncommon as there can be events taking place in one regime but not the other. A similar problem might occur within a phase if there are rare events happening. Secondly, estimating $P(\mathbf{x}|\Phi)$ itself introduces additional noise due to the Kernel Density smoothing, which is highly sensitive to the choice of bandwidth. An additional issue is that the tails of the distribution may not be well approximated (especially if they tend to infinite density as in Chapter 4).

One method to address this process might be to view the estimation of the FIM as a hopping diffusion problem. When computing the derivative, instead of trying to infer $P(\mathbf{x}|\Phi)$ and both $P(\mathbf{x}|\Phi \pm \varepsilon)$ by means of a Kernel Density approach, one could focus only on estimating $P(\mathbf{x}|\Phi)$ and then use the datapoint-specific dynamics $J_{s,t}(\Phi)$, which one can obtain by automatic differentiation, to estimate $H^{SKL}(\Phi)$ in a diffusion approximation.

Parameter Space Reduction: Isolating a Phase

One avenue of the sloppy models research pursued in the physics domain has been model reduction. In particular, Transtrum (2014) developed a methodology that, unlike the explorations in this thesis, follow the most ill-constrained directions. Rather than attempting to define the set of possible phenomena, this approach takes a given phenomena as fixed and tries to reduce the underlying model to the core parameters that have an effect on the phenomena. This is done by sequentially perturbing the most ill-defined direction, which, by definition should have a very small effect on the model's prediction. Doing this for multiple steps, Transtrum (2014) find that the most ill-defined direction often reduces to a parameter combination involving only a handful of parameters. These can then be analyzed and often replaced by a simpler single parameter, thus reducing the model's dimensionality. The simplest example of this would be a chemical reaction: there is a forward and a backward rate. Sequentially following the ill-constrained directions yields extremely large values for these rates, but keeps their ratio fixed. This allows the modeler to define a simpler parameter, the effective rate, thus reducing the parameter space by one dimension. Doing so sequentially can yield much simpler models

than one started with. Based on the results of Section 6.5, one might hypothesize that the Mark-0 model could be reduced to only the labor-market parameters for a given unemployment rate dynamic.

The Model Phase Space: A Visual Journey

One of the aims of the MABM community is to be useful for policy-making in a complex system by acting as a hypothetical world in which to test the effects of different policies and understand the drivers of their outcomes. While all of the exploration and reduction methods are valuable and interesting tools to ascertain these drivers and their potential policy levers, their specificity may not be easy to communicate. It would be ideal to have a visual representation of the phase space of a given model. One could develop this manually, with a series of two-dimensional phase diagrams such as those in the original work of Gualdi et al. (2015, 2017) and Bouchaud et al. (2018). However, this restricts the representation of the well-constrained directions to two dimensions, which may not be appropriate. For instance, Section 6.5 shows that in almost all cases, the first eigenvector contains at least three significant components. Quinn et al. (2019) developed an “Intensive PCA” method. Inspired by replica theory, the method generates an intensive embedding that preserves local distance but allows one to visualize the global structure of the model manifold. Applying this method to the Mark-0 model with respect to the unemployment rate dynamics, or in a multidimensional setting, has the potential to visualize the space of different phases such that one can understand its boundaries (i.e. reduced form models) and understand better how the model behaves. Such an approach might render MABMs more approachable, quite literally visualizing the boundaries to the “ABM can generate anything” critique.

11.2 Applications: Agent-based Integrated Assessment Models

Taken from Naumann-Woleske (2023) *Agent-based Integrated Assessment Models: Alternative Foundations to the Environment-Energy-Economics Nexus* with minor modifications to the introduction and conclusion to fit this chapter of the thesis.

Beyond the testing and evaluation of the parameter space exploration on the Mark-0 model, one major avenue of research would be the more widespread application of the parameter space exploration to other models. In particular, I am interested in the applications of ABM to the economy-environment-energy nexus. Anthropogenic climate change is one of the major global challenges we face as a society today, with widespread environmental, social and economic effects (IPCC, 2022). To assess the economic impact of climate change and develop economic policies, economists have developed Integrated Assessment Models (IAMs) that link economic dynamics with environmental aspects such as increasing CO₂ concentrations or changes in land-use dynamics. These models are central in assessments of climate change mitigation strategies and their implications

(e.g. see IPCC, 2018). However, existing models have been subject to critique, especially Policy Optimization models as suggested in Farmer et al. (2015). Agent-based Models have recently come to the fore as an alternative framework for macroeconomic modeling, as well as environment-energy-economics modeling (Balint et al., 2017; Ciarli and Savona, 2019; Lamperti et al., 2019b). In this section, I compare four Agent-based Integrated Assessment Models (ABIAMs) and how they respond to the critiques of the currently mainstream equilibrium-based IAMs. Previous reviews of the application of complexity economics and agent-based modeling to climate issues (e.g. Castro et al., 2020; Balint et al., 2017; Lamperti et al., 2019b) have focused on the general benefits of these approaches, but there has been no systematic comparison of existing ABIAMs.¹

For the purposes of this section, I consider an IAM to be a model with endogenous and linked climate and economic modules. That is, there is a feedback from the climate system back into the economy, generally in the form of a *damage* function. This is a rather restricted definition as there is a wide scope of these models (Krey et al., 2019). Furthermore, I consider primarily the cost-benefit type IAMs following the Dynamic Integrated Climate Economics (DICE) models of Nordhaus (1992) and the process-driven IAMs used in IPCC (2014, 2018, 2022).

The majority of the IAMs meeting this criteria consider a first-best economic system grounded in equilibrium (general or partial) and perfect foresight (Forster et al., 2018; Keppo et al., 2021).² It is these models that have faced strong critique for their grounding in equilibrium, foresight and optimization by representative agents in the macroeconomic modules underlying IAM (see Ackerman et al., 2009; Pindyck, 2013, 2017; Stern, 2013, 2016; Weitzman, 2013; Revesz et al., 2014; Farmer et al., 2015).³ A recent review by Krey et al. (2019) considers several distinct areas of critique specifically for IAMs:

1. The absence of heterogeneity of actors and within groups of actors, which is key to societal transitions due to social processes emerging from interactions and coordination (e.g. lifestyle change, political actions). The absence of heterogeneity also strongly limits the analysis of distributional effects despite their importance (Diffenbaugh and Burke, 2019).
2. Technology and its diffusion is misrepresented either by being exogenous, or, when endogenous, being too optimistic (or pessimistic) in advances and diffusion (see Mercure et al., 2019; Gambhir et al., 2019).
3. A lacking representation of the financial system, though it faces large risks (e.g. see Monasterolo, 2020; Monasterolo et al., 2019) and at the same time may be a

¹Lamperti et al. (2019b) showcased the usefulness of the ABIAM approach by highlighting the work of Lamperti et al. (2018a) and Wolf et al. (2013), but did not compare models in detail.

²A main source of the critique is centered around the IAMs coming from the economics literature, and that are based on the work of Nordhaus (1992). The IPCC review models tend to have a higher degree of complexity and detail.

³I refer here to the introduction of this thesis for a slightly wider discussion of the DSGE approach to macroeconomics.

- driving force in the green transition (see Campiglio et al., 2018; Caldecott, 2018).
4. Energy-economy feedback loops are not fully represented, with missing energy system, material and economy linkages (Pauliuk et al., 2017) and an unrealistic decoupling of economic growth from energy usage or emissions (Nieto et al., 2020).
 5. Adding to this list, IAMs, particularly those based on Nordhaus (1992), are frequently criticized for their "ad-hoc" (Lamperti et al., 2018a) representation of damages from increases in the temperature anomaly over pre-industrial levels, thus underestimating the costs of climate change and the benefits to a low-carbon economy (Stern, 2016; Pindyck, 2017). In these models, the gradual deterministic reactions to mean surface temperatures omit the emergence of tipping-points, rare events, increasing variability and decreasing predictability of climate conditions (Wright and Erickson, 2003; Farmer et al., 2015).

The accumulation of concerns and critiques suggests that it might be beneficial to consider alternative frameworks for representing the economy and the environment-economy feedback loops. ABAIMs aim to provide an alternative methodology for modeling the interactions of the socioeconomic system with the biosphere. Interest in the application of Agent-based modeling to the climate-economy-energy nexus dates back to Moss et al. (2001); Moss (2002), who argued that it can serve as a well validated description of social and natural systems. The actual applications of Agent-based Models to the environment-energy-economy nexus has recently been reviewed in Balint et al. (2017) and Castro et al. (2020), and includes topics such as carbon and electricity markets, technology diffusion models, and coalition formation. In relation to the biosphere-economy interaction, ABIAMs offer several distinct advantages such as the built in heterogeneity, a more granular endogenous innovation and diffusion process, a granular representation of the financial system, and agent-specific damage functions. Their modularity and detail of agent-based models also allows for a closer interaction with stakeholders, such as policymakers.

The remainder of this section proceeds as follows: first, I briefly presents the models under consideration and their macroeconomic backbones, before comparing the energy, resource and climate modules of these ABIAMs. Following this, I consider the policy studies and recommendations that these models have been used for. Finally, I consider several next steps in the ABIAM research stream. Tables E.1-E.4 in the Appendix give a detailed comparison of the different models in the spirit of Dawid and Delli Gatti (2018), with extensions for the energy (Table E.2) and climate (Table E.3) modules. These tables may serve as a reference point for the current state of the models going forward.

Agent-based Integrated Assessment Models: Four Candidates

To consider how Agent-based Models may address some of the critiques levied at Integrated Assessment Models, I consider four ABIAM models: the Dystopian Schumpeter meets Keynes model (DSK hereafter, Lamperti et al., 2018a, 2019a, 2020, 2021), the AB-

MIAM model of Safarzynska and van den Bergh (2022), the model of Czupryna et al. (2020) (CFHS hereafter), and the model of Gerdes et al. (2022) (GRSW hereafter). This section introduces the models, their purposes and general macroeconomic structure.

The DSK model is a climate-extension of the well-established Keynes+Schumpeter model (Dosi et al., 2010). The purpose of the DSK presented in Lamperti et al. (2018a, 2020) is to provide an alternative to the specification of climate damages due to increases in atmospheric CO_2 concentration. The model consists of heterogeneous firms in the consumption and capital goods sectors, both of which are powered by electricity from an energy sector. Capital goods firms produce machines, with an innovation and imitation process leading to newer capital vintages (more energy efficient, higher labor productivity) based on R&D investments. Consumption good firms invest in different machines to produce a generic consumption good purchased by the households. These investments are financed by through imperfect capital markets, where firms are subject to credit limitations. The energy sector produces electricity using stylized green (renewable) and brown (fossil) power plants, which differ in their cost structure and emissions. Again power plants may invest in R&D and pursue an imitation-innovation strategy. Finally, a climate module connects CO_2 emissions by all sectors to a carbon cycle, and thus to an increase in the temperature anomaly over pre-industrial times. This leads to a disaster generating distribution with increasing average damage, and a higher probability of extreme damages to firms' capital stock and labor productivity as CO_2 concentration rises.

In Lamperti et al. (2019a, 2021), the DSK model is extended with a detailed set of heterogeneous banks to study the effects of changes in financial regulation on the transition to a low-carbon economy in the presence of financial instability. Their study responds to a growing body of literature suggesting that climate change feedback loops could be increased by financial market instability, especially with a banking system exposed to physical risks (Monasterolo, 2020). In doing so, they also simplify the climate module, thus for the remainder of the section I will distinguish between the main branch of the DSK model and the extension to a financial system, referred to as DSK-FIN.

The second model, the ABMIAM proposed by Safarzynska and van den Bergh (2022) studies how heterogeneous agents and emerging inequalities in labor and capital income can lead to larger estimates of the social cost of carbon than the aggregate Dynamic Integrated Climate Economy (DICE) models based on Nordhaus (1992) that have been subject to strong critique. Building on the work of Safarzynska and van den Bergh (2017b,a), the ABMIAM considers a set of consumption good firms with quality differentiated products, which households purchase based on a combination of quality, price and peer-effects. Over time, firms can improve their production technology by investing in R&D. To expand, firms may request loans from a bank which is subject to credit regulation and an inter-bank market. To produce goods, the firms apply labor, capital and electricity. The energy sector is based on heterogeneous power plants fueled by gas, coal or renewable energy with different cost structures, efficiencies and emission intensities. The emissions of the energy sector lead to increases in the atmospheric stock of carbon,

and consequent damage feedback. These were modeled like the cost-benefit DICE IAM to compare calculations of the social cost of carbon, but unlike DICE the shocks are distributed among agents based on their relative share of wealth similar to Dennig et al. (2015).

Czupryna et al. (2020) propose a global ABIAM model including ten regions and multiple different goods. In each region consumer and capital firms produce their respective goods, and households purchase with preferences for price and a minimum necessary consumption. To expand their capacities, firms receive investments from households rather than a financial system, paying back a share of profits as returns. As in the previous models, production requires the use of electricity which is supplied by a two-layer energy sector. In the first instance, fuel-extraction firms mine from finite regional stocks. The collected fossil fuels are then sold to a heterogeneous set of power plants, who provide electricity regionally. Over time, production technology improves at an exogenously given rate, with renewable technology costs also dropping. The combustion of fossil fuels increases the mean temperature anomaly over pre-industrial times, which leads to region-specific agricultural shocks (reduced productivity), labor productivity reductions, and natural disasters (capital damage).

Finally, Gerdes et al. (2022) propose a fourth ABIAM model.⁴ The GRSW model considers two regions, one with a mining sector (a stylized global south) and one with a capital goods sector (a stylized global north). Both regions have local labor and consumption good markets, with a lower wage in the south than the north. Households purchase a homogeneous consumption good from their regional consumption firms, who produce using capital and labor. The capital goods sector additionally requires resources from the mining sector in the south. In the GRSW model, the mines cause localized pollution that reduces workers health and thus labor productivity. Simultaneously, the capital goods sector emits CO_2 when producing machines, which causes natural disaster shocks to hit firms across both regions. Using this model, the authors explore unequal exchange and a double burden faced by the global south (economic dependence on the north, resulting in lower wages, and heavy damage to the local environment).

In terms of the economic systems represented, these four models differ primarily in their representations of the financial system, with more detail in the DSK-FIN and ABMIAM, innovation (endogenous methods in DSK, ABMIAM and GRSW), spatial scope (country-level vs. global), energy system (degree of detail), and policy institutions. With respect to the climate modules, the structure in each is similar and oriented to CO_2 emissions, they differ in the types of damages and their distribution.

By construction, each of the four main models considered in this review respond to the IAM critique of missing heterogeneity and interactions. Even when initialized with perfectly similar agents, over time the interactions and market protocols lead to emerging heterogeneity, such as through differentiated innovation. Many of the models, such as

⁴The authors themselves do not consider it to be an IAM, however, they do link the economy and environment both ways, so it falls under the definition of an IAM as applied here.

the ABMIAM or CFHS initially impose some form of heterogeneity in order to study its effects and to calibrate the model more closely to empirically observed heterogeneity. The results indicate that this heterogeneity is indeed important. For example, Safarzyńska and van den Bergh (2022) find that reducing inequalities in labor income leads to a larger social cost of carbon, while inequalities in capital income (rent) lead to a larger share of the population driven into poverty, thus reducing GDP and emissions, and consequently the social cost of carbon.

In terms of geographic and spatial heterogeneity, the considered ABIAMs lag behind their process-based counterparts. Only the CFHS model represents multiple consumption sectors and a larger number of regions, though both are still small in comparison to some other more heterodox IAMs such as ACCLIMATE (Otto et al., 2017) or E3ME (Mercure et al., 2018). Incorporating this might lead to a more detailed description of the production process and the global value chains that underlie it. Agent-based Models have already been applied successfully to production networks, and could be integrated here as well (e.g. see the models of Wolf et al., 2013; Gualdi and Mandel, 2019).

Across the DSK, ABIAM and GRSW, innovation and technology diffusion are also endogenized. For instance, through the innovation-imitation process of different firms that lead to new technologies but also to the imitation of technology that are in use by industry leaders. However, technologies still primarily affect the productivity coefficients of the firms' production function, and are thus not necessarily related to specific technology paths such as those considered in process-based IAMs. It is thus unclear whether these models are overly optimistic or pessimistic, as noted in the critiques of Krey et al. (2019).

Finally, in terms of the financial system, both the DSK-FIN and the ABMIAM incorporate a detailed representation of imperfect capital markets that may lead to bank failures. In particular, the DSK-FIN extension is focused around studying the financial fragility emanating from increasing climate damages. The ABMIAM also offers a heterogeneous banking sector, which also includes interbank loans, thus creating an implicit interbank network that may be subject to cascading crises and the types of systemic risk that is not captured in more common IAMs (Monasterolo, 2020).

Energy and Resource Modules

In this section, I address how the four considered models treat the energy-resource-environment nexus. The representations range from very stylized (GRSW and DSK) to more sophisticated (CFHS) in terms of the variety of energy sources used in electricity production.

Beginning with the raw sources of materials and primary energy sources, the DSK and ABMIAM models consider stylized primary energy types, involving an extracted fossil fuel (gas and coal in ABMIAM) and a form of renewable energy that functions without the requirement for fuel inputs. In both these models, the extraction and provision of these fuels are outside of the model's boundary, and are thus infinitely available with

exogenously given price processes. While the GRSW model does not have an energy sector and primary energy sources, it does have a mining sector that extracts material resources in the southern region and sells them to the capital goods sector in the northern, yet here too there are no dynamics of mine-depletion as the purpose is primarily to study unequal exchange between north and south in the presence of local and global pollution.

CFHS goes beyond this, by considering not only the conversion of primary energy to electricity, but also the extraction of fossil fuels. In particular, they consider seven distinct fuel types (coal, gas, oil, nuclear, hydro, wind and solar), with a regional fuel-extraction sector. The presence of a fuel extraction sector is important, as the dynamics of reserve depletion through extraction and the shifts in regional extraction will impact the agents decisions about which type of power plant to invest in, what degree of emissions are possible, and how technological innovation interacts with depletion (see debates in Höök and Tang, 2013; Capellán-Pérez et al., 2014). In the CFHS model, each region has an empirically determined reserve of all fossil fuels. The marginal cost of their extraction is increasing in the cumulative amount of extraction (modelled by a Rogner curve as in Nordhaus and Boyer (2003)), reflecting the principle that the easiest-to-extract reserves are captured first. There is thus a regional shift as local reserves become depleted, favoring those regions with more abundant resources. This is an important driver in evolving regional inequalities, in the sense of the unequal exchange studied in the GRSW model (though they do not include depletion dynamics), as investments and employment will increase in those regions with the cheapest to extract resources.

Turning to the transformation of primary energy into energy carriers, the DSK, ABMIAM and CFHS models all consider only electricity as a final energy carrier. Notably, the CFHS model has a demand for fuel by households, but considers this to be part of the primary energy sector (there is no transformation to liquid fuels). All three of these models consider heterogeneous power plants in terms of their fuel type, cost structure, energy and emission intensity. For the DSK and ABMIAM the production is centrally cleared in a given period: sectors present their electricity demand and power plants are activated until this demand is met. The price is then determined as a markup on the operating cost of the final power plant to be activated such that supply meets demand. The CFHS model takes this one step further by considering the variability in renewable electricity generation. They split each simulation step into a series of substeps with alternating zero and peak production of solar. Markets are then cleared in the same procedure as the DSK and ABMIAM but on a subperiod base, which leads to an average price per simulation step. Furthermore, fossil and nuclear plants are also considered stochastic, activating with a beta distribution proportional to their capacity factors in order to simulate the effects of maintenance which can lead to price jumps.⁵

Across all models, investment in the replacement of power plants once they become obsolete and the decision about the type of power plant to construct are based on

⁵Such as simultaneous refueling and maintenance of nuclear plants in France in 2022 that increased the electricity price in the context of the gas crisis.

the net present value of each powerplant option, the candidate plant with the highest expected value (lowest lifetime cost) is then selected. Where the models differ is in how innovation in the energy-sector occurs. The DSK applies an endogenous innovation framework, which allows for reductions in the fixed costs of renewable power plants while for fossil-fuel plants there is an improvement in emission intensities. By contrast, the CFHS considers an exogenous noisy decrease in the investment cost for solar and wind, to a specified minimum, as well as an increase in the substitutability of electricity for fuels in consumer demands. While the technology process of the DSK is endogenous, it is unclear whether the reviewed technological innovations are overly optimistic or pessimistic, and thus whether they adequately respond to similar critiques of process-based IAMs (Keppo et al., 2021).

In summary, the energy sector remains rather stylized in three out of the four models, thus limiting the description of energy source substitution and technological advances in this respect, as well as the dynamics of fuel extraction and the multi-scale nature of energy flows through the economy (Giampietro et al., 2012). However, the CFHS model showed that a more detailed representation of the energy sector is feasible within ABIAM, including an implementation of resource extraction in a multi-regional context. Agent-based approaches have already been applied to energy-sector specific questions elsewhere, suggesting that more detailed energy-sector representations in ABIAM should be feasible (see Castro et al., 2020, for more information).

Climate Modules

This section addresses how ABIAMs model the environment and crucially the feedback mechanisms between socioeconomic system and environment. Table E.3 summarizes the climate and environmental modules of the four reviewed ABIAMs. The climate modules of the four reviewed models use similar simplified models of the carbon cycles as in the IAMs based on Nordhaus (2017). These climate boxes have been critiqued as oversimplifying the carbon cycle, in particular with relation to irreversible climate tipping points that lock-in certain temperature changes (Dietz et al., 2021).

Where the ABIAMs differentiate themselves is in modeling the damage feedback from increasing temperature anomalies, and in the case of the GRSW model also the effects of localised pollution. As a reference point, consider the case of Nordhaus (1992), where a deterministic (quadratic) function of temperature gives a multiplier that is applied to the total production. For example, a given temperature might map to a 5% reduction in total GDP compared to an “absence-of-damages” world. This approach has been quite heavily critiqued as underestimating damage impacts and their spatial heterogeneity, and ignoring non-linear tipping points (Lenton et al., 2019; Steffen et al., 2018). In contrast, each of the ABIAMs has a specific treatment of the size of the damage as well as its distribution amongst the model agents. With the exception of the ABMIAM model, there are also multiple types of damage, generally including damage to capital stocks and labor reductions (the ABMIAM considers only shocks to the consumers’ budgets) that are applied at the agent-level. This approach alone already differs from

the cost-benefit IAMs.

Turning to the calculation of damage magnitude and distribution, there is significant heterogeneity between the models dependent on their intended purpose. The DSK model directly aims to address the issue of the increasing variability and risk of rare climate events omitted in deterministic damage functions. In this spirit, the increase in temperature anomaly changes the parameters of a beta distribution, leading to a higher mean damage and a stronger skew. With each firm hit by a random draw affecting capital and labor productivity, Lamperti et al. (2019b) conclude that “under the same ‘business-as-usual’ emission scenario, roughly adherent to the Representative Concentration Pathway 8.5, the average climate shock from DSK and the damage function in DICE2013r (Nordhaus, 2014) are vaguely similar: the 2100- μ -level shock averages 5.4% while the DICE damage function implies a GDP loss of approximately 5.2%. However, aggregate impacts are radically different, with end-of-century projected output being around 90% of the ‘without-climate-change’ scenario in DICE while amounting to 15% in DSK (with shocks assumed to target labor productivity)”.

The choices made across models differ. The DSK-FIN uses the deterministic damage function from Nordhaus (2017), applying the damage perturbed with noise to all firms, in order to facilitate comparison to the cost-benefit IAM research. Meanwhile, the ABMIAM and the CFHS model aim to understand how heterogeneity in agents and shocks can lead to different results than when using the representative agent and aggregate damages. In particular, the ABMIAM looks at the social cost of carbon and its associated optimal tax rate when shocks are distributed amongst individual agents in relation to their wealth, finding significant impacts of wealth and rent-inequalities on the social cost of carbon, such as increased poverty rates, leading to a deterioration of the economy. CFHS aim to verify whether the aggregate damage functions arise when considering different agents and empirically calibrated deterministic micro-shocks. To do so they take literature estimates of agricultural and labor shocks, and empirically estimate the damage of natural disasters for each of their 10 sub-regions. Running the model, they then study the shape of the aggregate damage function to compare it to cost-benefit type deterministic functions. They find that depending on remaining fuel availability and the speed of renewable energy growth, the aggregate damage function can take different shapes, and have differing magnitudes. Finally, the GRSW model proposes that damages are proportional to the total stock of capital, while the frequency of shocks per period increases with the temperature anomaly. Thus increasing the probability that one firm may be hit multiple times per period in a non-linear way. This is in conjunction with being the only model to consider localised pollution as well. In the mining sector, workers’ health, proxied by their labor productivity, decreases with the period of their employment in the mine, thus requiring ever more new labor. They find that in an accelerated climate scenario, damages outstrip efforts to repair and innovate, leading to a long-term collapse of the economy.

In summary, ABIAMs have begun addressing the underestimation of climate damages suggested by Stern (2016). The benefit of the modular structure of ABMs is that many

different forms and distributions of damage can be considered, such as regional damages, as applied in CFHS, or agent-specific damages as in the ABMIAM. Thus models can easily incorporate on-going research on the micro-impacts of climate change (Carleton and Hsiang, 2016) and contribute to the ongoing debate on climate damage modelling (Diaz and Moore, 2017).

Policy Implications

The considered models differ markedly in the scope of implemented and tested policy options, depending on the respective authors' stated purpose. With the exception of CFHS, each model considers some form of policy aimed at mitigating carbon emissions and hence the negative feedback effects from increases in CO_2 concentration. Over-archingly, the benefit of the modular ABIAMs is that policies that have been suggested in the literature but are not typically possible in cost-benefit IAMs can be tested. To highlight this, I consider the case of green financial regulation in the DSK-FIN model and the implementation of a global markets institution in the GRSW model.

The literature on the climate-finance nexus suggests that an active role of financial regulators can help shape climate related risk management (Monasterolo, 2020) and a low-carbon transition (Campiglio et al., 2018). The DSK-FIN model has a detailed set of heterogeneous banks that are subject to credit regulation. In this context, the authors test three regulatory policies: (i) green Basel requirements which excludes loans to green firms are excluded from banks capital requirements, (ii) green credit guarantees where the government guarantees the value of loans to green firms, and (iii) carbon-adjusted credit ratings where greener companies receive an improved credit rating. In the context of the model, the authors find that individually none of the policies manage the trade-off between climate mitigation and stable growth, with either a mitigation at the cost of more frequent banking crises, or a stable economy until climate damages become extreme enough to lead to collapse. In addition to the new financial regulations, the DSK model also includes standard leaning-against-the wind policies such as unemployment benefits and taxes on firm profits, as well as credit-multipliers to limit total debt.

In contrast to the financial system regulation, Gerdes et al. (2022) use the GRSW model to explore the implications of a "civilized market institution" that has the power to fine capital-goods companies for emissions and mines for local pollution, and redistribute the collected fines in the form of subsidies for retrofitting firms' capital to reduce their pollution and emission intensities. Their results suggest that fines alone, which may include items such as carbon taxes, are insufficient to counter natural disaster damage nor unequal exchange, regardless of the severity of the fines. On the other hand, when these fines are used to subsidize innovation activities in the form of retrofitting capital, does the value chain become sustainable over the long run. All cases are cost-neutral, as subsidies are funded by sanctions, thus converging to zero as emissions are mitigated. They find that the most promising results include a North-South transfer as capital firms in the north still invest in carbon reductions at the same rate, but significantly boosts mitigation in mines. In a similar vein to the DSK, the model contains unemployment

benefits and taxes on firm profits as well.

The ABMIAM of Safarzynska and van den Bergh (2022) considers the more common policy of carbon taxation. Specifically, they use an adapted form of the analytically derived optimal policy by Rezai and Van der Ploeg (2016), with revenues being redistributed as equal lump-sum payments to citizens. They then analyze to what degree these taxes may reduce inequalities versus the same scenarios in their absence. Their findings suggest that income inequality can be addressed by these policies, while consumption and wealth inequality are not strongly affected.⁶

The policies highlighted in the DSK-FIN and GRSW model show the potential for ABIAM in policy analysis because ABIAM are able to also incorporate regulatory policies and a wide array of policy mixes. This would allow for consultation with policymakers to study in more expansive terms policy interactions and implications. It also allows to address some of the critiques that a focus on carbon pricing omits the interaction with innovation and diffusion processes (Rosenbloom et al., 2020), as well as the interactions and trade-offs with other sustainable development goals such as equity (Geels et al., 2016). At the same time, Keppo et al. (2021) note that policies alternative to the commonplace carbon taxation or trading are also possible in process-based IAMs, and suggest that the focus is driven by how models are meant to be used and what policymakers have demanded.

The Path Ahead

While ABIAMs have made advances in revealing the micro-impacts of climate shocks, the effects of inequality on policy, and the interlinkages between the economy, the financial system and climate feedbacks, there is much yet to be done to match the detail of process-based IAMs. In first place, the current set of ABIAMs should continue developing to incorporate features common to process-based IAMs, such as a fine-grained representation of the energy system, as started in CFHS, a multi-sectoral structure beyond stylized consumption and capital goods, and a connection with land-use and change models. Several of these areas have already been addressed by agent-based models, such as electricity markets (see Castro et al., 2020, for references) or multi-sectoral and multi-regional structures as in the LAGOM models (Wolf et al., 2013), ACCLIMATE (Otto et al., 2017) or Caiani et al. (2018, 2019a). Beyond the addition of detail and combination of these models, there are several areas worth highlighting.

Bardi and Pereira (2022) point out that “[i]n recent years, global warming and climate change have become the main focus of the environmental movement. That may have led to the importance of resource depletion being neglected - more evidence of the importance of an integrated approach.” A review by Pollitt et al. (2010) confirms these suspicions by reviewing 60 energy-environment-economics models and coming to the

⁶Consumption inequality is addressed in an unjust economy where damages are allocated to the lowest wealth households, as this essentially functions as a non-damaged income such that households can maintain their consumption where otherwise they would be in poverty.

conclusion that “consumption of material inputs is largely unexplored within a dynamic macroeconomic framework”. More recent reviews by Pauliuk et al. (2017) find that IAMs have missing material-energy-economy linkages in their descriptions of installed capital and infrastructures. Indeed, also in the reviewed ABIAMs only GRSW consider the extraction of a physical resource to construct capital, while also not treating its end-of-lifecycle implications. Given the scale of resource extraction and the different required resources, even a saturation of installed capital levels would require large quantities of new virgin materials in developing countries, and for the energy transition (Wiedenhofer et al., 2019, 2021; Watari et al., 2019), which may make a full transition infeasible when aiming to maintain or increase current consumption levels (Michaux, 2021). It is worth noting here that the extraction and processing of raw materials is typically an emission-intensive process. Finally, also end-of-life waste and recycling may become an important factor in a circular economy system with reduced virgin material extraction, this is also underappreciated in current models (McCarthy et al., 2018). The detailed and heterogeneous structure of ABMs would also allow modelers to incorporate insights from industrial ecology to assess not just the socioeconomic but also the material dimensions of a climate transition. In particular, I suggest considering modeling at the stock-flow-practices nexus (Haberl et al., 2021), where the practices and demands of individual consumers become important in the requirements for various materials. Modeling individual agents also allows for interactions and social phenomena on the demand-side that may change the path of the climate transition (e.g. see the review of Castro et al., 2020, for some ABM-based approaches), where there are unexplored policy avenues (e.g. see Fitzpatrick et al., 2022). All of the ABIAMs reviewed here show that there are distributional effects of climate change, especially when considering the differential impacts of climate feedbacks both regionally in the CFHS and GRSW models, and among individuals in ABMIAM and DSK. More broadly, ABM-type models may be able to incorporate more indicators and analyses relating to the interaction among different Sustainable Development Goals such that policymakers and citizens may assess and decide on trade-offs.

Finally, additions and enlargements of ABMs can run into computational issues. Thus one wishes to carefully choose where to model heterogeneity to capture key aspects of the overall system behavior (Keppo et al., 2021). To facilitate this, a process of *agentization* should be considered (Guerrero and Axtell, 2011). This involves beginning the modeling exercise with a model that captures overall system behavior at an aggregate level, and then slowly replacing sectors or aggregations with individual agents. This would facilitate both the computational exercise, validation, and communicability of the results of ABMs. In particular, I am interested in the study of climate financial risk, whose importance has been highlighted by Battiston et al. (2021a, 2023, 2021b), Monasterolo (2020) and Monasterolo et al. (2022a,b), and how one can design policy, as was done in the DSK-FIN, to mitigate the effects of physical and transition risks (Monasterolo et al., 2019). To this end, there are several models incorporating climate and finance linkages already that would be excellent starting points for agentization, such as the EIRIN flow-of-funds

model of Monasterolo and Raberto (2018)⁷ or the model of Dafermos et al. (2017, 2018).). Even in a non-agentized setup, these models already have a high number of parameters, which may make the introduction of more agents appear to create a black-box.⁸ But, as Part I of this thesis showed, only a handful of critical directions are relevant to the observed dynamics and for the calibration of these models, suggesting this as a fruitful avenue for research that connects macroeconomic dynamics with micro-agent behavior in the context of climate change.

Conclusions

In sum, there are many avenues of future research to follow with respect to the exploration and understanding of the parameter space in Agent-based Models: probabilistic approaches, parameter space reduction, and visualization, together with applications to the modeling of the economy-environment-energy nexus. To this I would like to add a hope that future developments of ABM by their respective authors consider their differentiable implementation. Many models are currently written in an object-oriented manner using Java or using C++ to improve their execution speed. However, this complicates their analysis, even for simple tasks such as computing finite difference derivatives, and prevents the use of the statistical ecosystems of Python or Julia, which provide differentiable libraries that may lead to the next leap in the exploration and calibration of MABM.

11.3 Inflationary Dynamics and their Drivers

Part II of this thesis presented three different models of confidence and collective behavior in a macroeconomic setting. Each of these models, as is the case with all models, is generally incomplete and ripe for various extensions. For instance, including credit dynamics in the Dynamic Solow Model to study higher frequency business cycles through a Minskian instability approach, or including active policy in the modified DSGE of Chapter 8 and extending it to a scope more alike to Smets and Wouters (2007). However, the continuing research I am most interested in is the extension of the post-COVID study using Mark-0 as presented in Chapter 9. This is because, as mentioned for the ABIAM approaches, energy and resource price dynamics will likely grow in importance as part of the green energy transition and our collective response to anthropogenic climate change.

The simplest extension to the model is the consideration of a dual mandate for the Central Bank, and how this may address or fail to address spikes in the unemployment rate. More critically would be to expand on the firm dynamics and with it the labor market by introducing a firm-to-firm network that may dampen or amplify the shocks we have seen. In particular, for the ongoing debate on supply-chain disruptions, having

⁷See also Naqvi and Monasterolo (2021); Dunz et al. (2021); Gourdel et al. (2022) for recent advances using the EIRIN model

⁸more than 100 for the EIRIN version used in Gourdel et al. (2022) and more than 70 for Dafermos et al. (2017).

a firm-to-firm network with embedded industries would be an asset, as one can see how these types of shocks can propagate and possible effects of profit-driven inflation. Together with a more detailed firm-to-firm interaction can come a more detailed sectoral employment and labor market dynamic, that may also consider a network-like structure as in del Rio-Chanona et al. (2020). As it stands, the analysis presented in Chapter 9 is based on data for the United States, where, in contrast to the classic unemployment-inflation tradeoff, labor markets have remained tight. The reasons behind this are mixed: job openings have outgrown the number of unemployed rapidly, with variations by state (Jefferson and Fuller, 2023) and industry (Birinci and Ngan, 2023). Simultaneously, the U.S. labor force participation rate has remained below its pre-pandemic level following a longer-term decline since the Great Financial Crisis of 2008 (Lee et al., 2023). Recent research by Lee et al. (2023) suggests the negative impact of 2008-09 has been on the extensive margin, as people are continuously leaving the labor market, most notably young males without a bachelor's degree. However, the pandemic-specific drop in annual hours worked was primarily in the intensive margin: with educated prime-age workers reducing the number of hours worked. Chapters 6 and 9 highlighted how critical the labor-market dynamics and parameterization is to the overall dynamic of the Mark-0 model. Incorporating these longer-term trends and understanding these mechanics within the model would shed light on the high-inflation high-interest environment with a low unemployment rate. To do so, I believe that incorporating a labor flow network in addition to the firm-to-firm network would be the best approach, as it allows us to maintain an otherwise aggregated household sector.

Finally, as with the ABIAMs mentioned above, one aspect I am interested in is the financial risk aspect of these shocks. The Mark-0 model already incorporates bankruptcies, where the bankruptcy threshold can cause a phase transition out of the full employment regime (Gualdi et al., 2015). This is already relevant as during the COVID-19 pandemic many countries extended additional credit lines or lifted bankruptcy requirements, raising fears of a bankruptcy wave. However, the financial system is a key channel for monetary policy to pass through to the economy, yet the bank in Mark-0 is currently very stylized, setting loan-rates based on the aggregate default rates and no credit rationing. One can extend this setup to study more closely the pass-through of monetary policy through credit issuance by adopting the approach presented in the Post-Keynesian and Stock-Flow Consistent (SFC) modeling literature (e.g. see the models in Caverzasi and Godin, 2015). This would involve following a setup such as Godley and Lavoie (2006) or Lavoie and Godley (2006), to have a bank agent that lends to the individual firms at firm-specific lending rates, and is subject to prudential regulation. These constraints then determine the degree to which the interest rates are passed through, and how firms react to changes in the credit environment.

Bibliography

- Michael C. Abbott and Benjamin B. Machta. Far from Asymptopia: Unbiased High-Dimensional Inference Cannot Assume Unlimited Data. *Entropy*, 25(3):434, March 2023. ISSN 1099-4300. doi: 10.3390/e25030434.
- Frank Ackerman, Stephen J. DeCanio, Richard B. Howarth, and Kristen Sheeran. Limitations of integrated assessment models of climate change. *Climatic Change*, 95(3):297–315, August 2009. ISSN 1573-1480. doi: 10.1007/s10584-009-9570-x.
- Philippe Aghion and Peter Howitt. A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–351, 1992. ISSN 0012-9682. doi: 10.2307/2951599.
- Philippe Aghion and Peter Howitt. Joseph Schumpeter Lecture Appropriate Growth Policy: A Unifying Framework. *Journal of the European Economic Association*, 4(2-3):269–314, 2006. ISSN 1542-4774. doi: 10.1162/jeea.2006.4.2-3.269.
- George A. Akerlof and Robert J. Shiller. *Animal Spirits*. Princeton University Press, 2010.
- Peter S. Albin and Duncan K. Foley. *Barriers and Bounds to Rationality: Essays on Economic Complexity and Dynamics in Interactive Systems*. Princeton Studies in Complexity. Princeton University Press, Princeton, N.J, 1998. ISBN 978-0-691-02676-3.
- Jorge Alvarez, John C Bluedorn, Niels-Jakob Hansen, Youyou Huang, Evgenia Pugacheva, and Alexandre Sollaci. Wage-price spirals: What is the historical evidence? 2022.
- Philipp Andelfinger. Towards Differentiable Agent-Based Simulation. *ACM Transactions on Modeling and Computer Simulation*, 32(4):27:1–27:26, January 2023. ISSN 1049-3301. doi: 10.1145/3565810.
- P. W. Anderson. More Is Different: Broken symmetry and the nature of the hierarchical structure of science. *Science*, 177(4047):393–396, August 1972. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.177.4047.393.
- Philip W. Anderson, Kenneth Arrow, and David Pines. *The Economy as an Evolving Complex System: The Proceedings of the Evolutionary Paths of the Global Economy Workshop*. Number 5 in Santa Fe Institute Studies in the Sciences of Complexity Proceedings. Addison-Wesley, Redwood, CA, 1988. ISBN 978-0-201-15685-0.

-
- George-Marios Angeletos and Jennifer La'O. Sentiments. *Econometrica*, 81(2):739–779, 2013. ISSN 1468-0262. doi: 10.3982/ECTA10008.
- George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Quantifying Confidence. *Econometrica*, 86(5):1689–1726, 2018. ISSN 1468-0262. doi: 10.3982/ECTA13079.
- George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Business-Cycle Anatomy. *American Economic Review*, 110(10):3030–3070, October 2020. ISSN 0002-8282. doi: 10.1257/aer.20181174.
- Joshua F. Apgar, David K. Witmer, Forest M. White, and Bruce Tidor. Sloppy models, parameter uncertainty, and the role of experimental design. *Molecular BioSystems*, 6(10):1890–1900, 2010. doi: 10.1039/B918098B.
- Kenneth J. Arrow. The Economic Implications of Learning by Doing. *The Review of Economic Studies*, 29(3):155–173, 1962. ISSN 0034-6527. doi: 10.2307/2295952.
- W. Brian Arthur, Steven N. Durlauf, David A. Lane, and SFI Economics Program, editors. *The Economy as an Evolving Complex System II*. Number v. 27 in Proceedings Volume ... Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley, Advanced Book Program, Reading, Mass, 1997. ISBN 978-0-201-95988-8 978-0-201-32823-3.
- Quamrul Ashraf, Boris Gershman, and Peter Howitt. Banks, Market Organization, and Macroeconomic Performance: An Agent-Based Computational Analysis. *Journal of Economic Behavior and Organization*, 135:143–180, March 2017. ISSN 0167-2681, 0167-2681. doi: 10.1016/j.jebo.2016.12.023.
- Tiziana Assenza and Domenico Delli Gatti. The Financial Transmission of Shocks in a Simple Hybrid Macroeconomic Agent Based Model. *Journal of Evolutionary Economics*, 29(1):265–297, March 2019. ISSN 0936-9937, 0936-9937. doi: 10.1007/s00191-018-0559-3.
- Tiziana Assenza, Domenico Delli Gatti, and Jakob Grazzini. Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control*, 50:5–28, January 2015. ISSN 01651889. doi: 10.1016/j.jedc.2014.07.001.
- Maria Grazia Attinasi, Mirco Balatti, Michele Mancini, Luca Metelli, et al. Supply chain disruptions and the effects on the global economy. *Economic Bulletin Boxes*, 8, 2022.
- Dimitrios Bakas, Georgios Chortareas, and Georgios Magkonis. Volatility and growth: A not so straightforward relationship. *Oxford Economic Papers*, 71(4):874–907, October 2019. ISSN 0030-7653. doi: 10.1093/oenp/gpy065.
- Tomas Balint, Francesco Lamperti, Antoine Mandel, Mauro Napoletano, Andrea Roventini, and Alessandro Sapia. Complexity and the Economics of Climate Change: A Survey and a Look Forward. *Ecological Economics*, 138:252–265, August 2017. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2017.03.032.
- Opeoluwa Banwo, Paul Harrald, and Francesca Medda. Understanding the Consequences of Diversification on Financial Stability. *Journal of Economic Interaction and Coordination*, 14(2):273–292, June 2019. ISSN 1860-711X, 1860-711X. doi: 10.1007/s11403-018-0216-9.
- David Rezza Baqaee and Emmanuel Farhi. The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem. *Econometrica*, 87(4):1155–1203, 2019. ISSN 0012-9682. doi: 10.3982/ECTA15202.

BIBLIOGRAPHY

- Sylvain Barde. Direct Comparison of Agent-Based Models of Herding in Financial Markets. *Journal of Economic Dynamics and Control*, 73:329–353, December 2016. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2016.10.005.
- Ugo Bardi and Carlos Alvarez Pereira. *Limits and Beyond: 50 years on from The Limits to Growth, what did we learn and what's next?* Exapt Press, Verlagsort nicht ermittelbar, 2022. ISBN 978-1-914549-03-8.
- Robert B Barsky and Eric R Sims. Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review*, 102(4):1343–1377, June 2012. ISSN 0002-8282. doi: 10.1257/aer.102.4.1343.
- Stefano Battiston, Yannis Dafermos, and Irene Monasterolo. Climate risks and financial stability. *Journal of Financial Stability*, 54:100867, June 2021a. ISSN 15723089. doi: 10.1016/j.jfs.2021.100867.
- Stefano Battiston, Irene Monasterolo, Keywan Riahi, and Bas J. Van Ruijven. Accounting for finance is key for climate mitigation pathways. *Science*, 372(6545):918–920, May 2021b. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.abf3877.
- Stefano Battiston, Antoine Mandel, Irene Monasterolo, and Alan Roncoroni. Climate credit risk and corporate valuation, April 2023.
- Mr Thomas Baunsgaard and Nate Vernon. *Taxing Windfall Profits in the Energy Sector*. International Monetary Fund, 2022.
- Paul Beaudry and Franck Portier. News-Driven Business Cycles: Insights and Challenges. *Journal of Economic Literature*, 52(4):993–1074, 2014. ISSN 0022-0515.
- Paul Beaudry, Dana Galizia, and Franck Portier. Putting the Cycle Back into Business Cycle Analysis. *American Economic Review*, 110(1):1–47, January 2020. ISSN 0002-8282. doi: 10.1257/aer.20190789.
- Jess Benhabib and Roger E. A. Farmer. Chapter 6 Indeterminacy and sunspots in macroeconomics. In *Handbook of Macroeconomics*, volume 1, pages 387–448. Elsevier, January 1999. doi: 10.1016/S1574-0048(99)01009-5.
- Ben Bernanke and Olivier Blanchard. What caused the U.S. pandemic-era inflation? March 2023.
- Ben Bernanke, Mark Gertler, and Simon Gilchrist. Chapter 21 The financial accelerator in a quantitative business cycle framework. In *Handbook of Macroeconomics*, volume 1, pages 1341–1393. Elsevier, January 1999. doi: 10.1016/S1574-0048(99)10034-X.
- Ben Bernanke, Timothy F. Geithner, and Henry M. Paulson Jr. *Firefighting: The Financial Crisis and Its Lessons*. Penguin Books, New York, 2019. ISBN 978-0-14-313448-0.
- Serdar Birinci and Tran Khanh Ngan. Labor Market Tightness after the COVID-19 Recession: Differences across Industries, 2023.
- Olivier Blanchard. Output, the stock market, and interest rates. *The American Economic Review*, 71(1):132–143, 1981. ISSN 00028282.
- Olivier Blanchard. Consumption and the Recession of 1990-1991. *The American Economic Review*, 83(2):270–274, 1993.

-
- Olivier Blanchard. Where Danger Lurks. *Finance and Development (IMF)*, 2014.
- Olivier Blanchard. On the future of macroeconomic models. *Oxford Review of Economic Policy*, 34(1-2):43–54, January 2018. ISSN 0266-903X, 1460-2121. doi: 10.1093/oxrep/grx045.
- Alan S. Blinder. The Anatomy of Double-Digit Inflation in the 1970s. In *Inflation: Causes and Effects*, pages 261–282. University of Chicago Press, 1982.
- Alan S. Blinder and Jeremy B. Rudd. The Supply-Shock Explanation of the Great Stagflation Revisited. *NBER Chapters*, pages 119–175, 2013.
- Michele Boldrin and Michael Woodford. Equilibrium models displaying endogenous fluctuations and chaos: A survey. *Journal of Monetary Economics*, 25(2):189–222, March 1990. ISSN 0304-3932. doi: 10.1016/0304-3932(90)90013-T.
- Andrea Borsato. Secular Stagnation and innovation dynamics: An agent-based SFC model. Part I. September 2020.
- Andrea Borsato. An Agent-based Model for Secular Stagnation in the USA: Theory and Empirical Evidence. March 2021.
- Z. I. Botev, J. F. Grotowski, and D. P. Kroese. Kernel density estimation via diffusion. *The Annals of Statistics*, 38(5):2916–2957, October 2010. ISSN 0090-5364, 2168-8966. doi: 10.1214/10-AOS799.
- Alberto Botta, Eugenio Caverzasi, and Alberto Russo. When complexity meets finance: A contribution to the study of the macroeconomic effects of complex financial systems. *Research Policy*, page 103990, May 2020. ISSN 00487333. doi: 10.1016/j.respol.2020.103990.
- Jean-Philippe Bouchaud. Crises and Collective Socio-Economic Phenomena: Simple Models and Challenges. *Journal of Statistical Physics*, 151(3-4):567–606, 2013. ISSN 0022-4715.
- Jean-Philippe Bouchaud. Radical complexity. *Entropy. An International and Interdisciplinary Journal of Entropy and Information Studies*, 23(12):1676, 2021.
- Jean-Philippe Bouchaud, Stanislao Gualdi, Marco Tarzia, and Francesco Zamponi. Optimal inflation target: Insights from an agent-based model. *Economics: The Open-Access, Open-Assessment E-Journal*, 2018. ISSN 1864-6042. doi: 10.5018/economics-ejournal.ja.2018-15.
- George E. P. Box. Science and Statistics. *Journal of the American Statistical Association*, 71(356):791–799, December 1976. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.1976.10480949.
- William A. Branch and Bruce McGough. Heterogeneous Expectations and Micro-Foundations in Macroeconomics. In *Handbook of Computational Economics*, volume 4, pages 3–62. Elsevier, 2018. ISBN 978-0-444-64131-1. doi: 10.1016/bs.hescom.2018.03.001.
- Jan-Niklas Brenneisen. Monetary policy under imperfect information and consumer confidence. 2020.
- William Brock and Steven Durlauf. Interactions-based models. In *Handbook of Econometrics*, pages 3297–3380. Elsevier, 2001.
- K S Brown, C C Hill, G A Calero, C R Myers, K H Lee, J P Sethna, and R A Cerione. The statistical mechanics of complex signaling networks: Nerve growth factor signaling. *Physical*

BIBLIOGRAPHY

- Biology*, 1(3):184–195, October 2004. ISSN 1478-3967, 1478-3975. doi: 10.1088/1478-3967/1/3/006.
- Kevin S. Brown and James P. Sethna. Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*, 68(2):021904, August 2003. ISSN 1063-651X, 1095-3787. doi: 10.1103/PhysRevE.68.021904.
- Michael Bruno and William Easterly. Inflation crises and long-run growth. *Journal of Monetary Economics*, 41(1):3–26, February 1998. ISSN 0304-3932. doi: 10.1016/S0304-3932(97)00063-9.
- Willem Buiter. The unfortunate uselessness of most 'state of the art' academic monetary economics. *VOX EU*, June 2009.
- Alessandro Caiani, Ermanno Catullo, and Mauro Gallegati. The effects of fiscal targets in a monetary union: A multi-country agent-based stock flow consistent model. *Industrial and Corporate Change*, 27(6):1123–1154, December 2018. ISSN 0960-6491. doi: 10.1093/icc/dty016.
- Alessandro Caiani, Ermanno Catullo, and Mauro Gallegati. The Effects of Alternative Wage Regimes in a Monetary Union: A Multi-country Agent Based-Stock Flow Consistent Model. *Journal of Economic Behavior and Organization*, 162:389–416, June 2019a. ISSN 0167-2681, 0167-2681. doi: 10.1016/j.jebo.2018.12.023.
- Alessandro Caiani, Alberto Russo, and Mauro Gallegati. Does Inequality Hamper Innovation and Growth? An AB-SFC Analysis. *Journal of Evolutionary Economics*, 29(1):177–228, March 2019b. ISSN 0936-9937, 0936-9937. doi: 10.1007/s00191-018-0554-8.
- Ben Caldecott, editor. *Stranded Assets and the Environment: Risk, Resilience and Opportunity*. Routledge Explorations in Environmental Studies. Routledge, Taylor & Francis Group, London ; New York, 2018. ISBN 978-1-138-12060-0.
- Emanuele Campiglio, Yannis Dafermos, Pierre Monnin, Josh Ryan-Collins, Guido Schotten, and Misa Tanaka. Climate change challenges for central banks and financial regulators. *Nature Climate Change*, 8(6):462–468, 2018. ISSN 1758-6798.
- Iñigo Capellán-Pérez, Margarita Mediavilla, Carlos de Castro, Óscar Carpintero, and Luis Javier Miguel. Fossil fuel depletion and socio-economic scenarios: An integrated approach. *Energy*, 77:641–666, December 2014. ISSN 0360-5442. doi: 10.1016/j.energy.2014.09.063.
- Alberto Cardaci and Francesco Saraceno. Inequality and Imbalances: A Monetary Union Agent-Based Model. *Journal of Evolutionary Economics*, 29(3):853–890, July 2019. ISSN 0936-9937, 0936-9937. doi: 10.1007/s00191-019-00611-4.
- Tamma A. Carleton and Solomon M. Hsiang. Social and economic impacts of climate. *Science*, 353(6304):aad9837, September 2016. doi: 10.1126/science.aad9837.
- Wendy Carlin and David Soskice. Stagnant productivity and low unemployment: Stuck in a Keynesian equilibrium. *Oxford Review of Economic Policy*, 34(1-2):169–194, January 2018. ISSN 0266-903X. doi: 10.1093/oxrep/grx060.
- Charles T. Carlstrom, Timothy S. Fuerst, and Matthias Paustian. Targeting Long Rates in a Model with Segmented Markets. *American Economic Journal: Macroeconomics*, 9(1):205–242, January 2017. ISSN 1945-7707. doi: 10.1257/mac.20150179.

- Adrian Carro, Marc Hinterschweiger, Arzu Uluc, and J Doyne Farmer. Heterogeneous effects and spillovers of macroprudential policy in an agent-based model of the UK housing market. *Bank of England Staff Working Paper*, 2022.
- Christopher D. Carroll, Jeffrey C. Fuhrer, and David W. Wilcox. Does Consumer Sentiment Forecast Household Spending? If So, Why? *The American Economic Review*, 84(5):1397–1408, 1994. ISSN 0002-8282.
- David Cass and Karl Shell. Do Sunspots Matter? *Journal of Political Economy*, 91(2):193–227, 1983. ISSN 0022-3808.
- Juana Castro, Stefan Drews, Filippos Exadaktylos, Joël Foramitti, Franziska Klein, Théo Konc, Ivan Savin, and Jeroen van den Bergh. A review of agent-based modeling of climate-energy policy. *WIREs Climate Change*, 11(4):e647, July 2020. ISSN 1757-7780. doi: 10.1002/wcc.647.
- Gennaro Catapano, Francesco Franceschi, Michele Loberto, and Valentina Michelangeli. Macroprudential Policy Analysis via an Agent Based Model of the Real Estate Sector. *Banca d’Italia Working Papers*, 2021. ISSN 1556-5068. doi: 10.2139/ssrn.3891583.
- Alberto Cavallo and Oleksiy Kryvtsov. What can stockouts tell us about inflation? Evidence from online micro data. *Journal of International Economics*, page 103769, 2023.
- Eugenio Caverzasi and Antoine Godin. Post-Keynesian stock-flow-consistent modelling: A survey. *Cambridge Journal of Economics*, 39(1):157–187, January 2015. ISSN 0309-166X. doi: 10.1093/cje/beu021.
- Eugenio Caverzasi and Alberto Russo. Toward a new microfounded macroeconomics in the wake of the crisis. *Industrial and Corporate Change*, 27(6):999–1014, December 2018. ISSN 0960-6491, 1464-3650. doi: 10.1093/icc/dty043.
- V. V. Chari, Patrick J. Kehoe, and Ellen R. McGrattan. New Keynesian Models: Not Yet Useful for Policy Analysis. *American Economic Journal: Macroeconomics*, 1(1):242–266, January 2009. ISSN 1945-7707. doi: 10.1257/mac.1.1.242.
- Lucy Chennells. The windfall tax. *Fiscal Studies*, 18(3):279–291, 1997.
- Enkhbayar Chojijil, Christian Espinosa Méndez, Wing-Keung Wong, João Paulo Vieito, and Munkh-Ulzii Batmunkh. Thirty years of herd behavior in financial markets: A bibliometric analysis. *Research in International Business and Finance*, 59:101506, January 2022. ISSN 02755319. doi: 10.1016/j.ribaf.2021.101506.
- Ayush Chopra, Alexander Rodríguez, Jayakumar Subramanian, Balaji Krishnamurthy, B. Aditya Prakash, and Ramesh Raskar. Differentiable agent-based epidemiological modeling for end-to-end learning. In *ICML 2022 Workshop AI for Agent-Based Modelling*, July 2022.
- Dimitris Christelis, Dimitris Georgarakos, Tullio Jappelli, and Maarten Van Rooij. Trust in the central bank and inflation expectation. *Available at SSRN 3540974*, 2020.
- Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1):1–45, 2005. ISSN 0022-3808. doi: 10.1086/426038.
- Lawrence J. Christiano, Martin S. Eichenbaum, and Mathias Trabandt. On DSGE Models. *Journal of Economic Perspectives*, 32(3):113–140, August 2018. ISSN 0895-3309. doi: 10.1257/jep.32.3.113.

BIBLIOGRAPHY

- Tommaso Ciarli and Maria Savona. Modelling the Evolution of Economic Structure and Climate Change: A Review. *Ecological Economics*, 158:51–64, April 2019. ISSN 0921-8009, 0921-8009. doi: 10.1016/j.ecolecon.2018.12.008.
- Silvano Cincotti, Marco Raberto, and Andrea Teglio. Credit Money and Macroeconomic Instability in the Agent-based Model and Simulator Eurace. *Economics: The Open-Access, Open-Assessment E-Journal*, 4(2010-26):1, 2010. ISSN 1864-6042. doi: 10.5018/economics-ejournal.ja.2010-26.
- Silvano Cincotti, Marco Raberto, and Andrea Teglio. Why do we need agent-based macroeconomics? *Review of Evolutionary Political Economy*, 3(1):5–29, April 2022. ISSN 2662-6144. doi: 10.1007/s43253-022-00071-w.
- David Colander, Richard Holt, and J. Barkley Rosser. The Changing Face of Mainstream Economics. *Review of Political Economy*, 16(4), 2004.
- David Colander, Peter Howitt, Alan Kirman, Axel Leijonhufvud, and Perry Mehrling. Beyond DSGE Models: Toward an Empirically Based Macroeconomics. *American Economic Review*, 98(2):236–40, May 2008. ISSN 0002-8282. doi: 10.1257/aer.98.2.236.
- T.F. Cooley and E.C. Prescott. Economic growth and business cycle. In *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, N.J, 1995.
- Russell Cooper and Andrew John. Coordinating Coordination Failures in Keynesian Models. *The Quarterly Journal of Economics*, 103(3):441–463, 1988. ISSN 0033-5533. doi: 10.2307/1885539.
- A Cornea, Cars H Hommes, and Domenico Massaro. Behavioral heterogeneity in U.S. inflation dynamics. *Journal of Business & Economic Statistics*, 37(2):288–300, 2019.
- Andrew Crooks, Christian Castle, and Michael Batty. Key challenges in agent-based modelling for geo-spatial simulation. *Computers, Environment and Urban Systems*, 32(6):417–430, November 2008. ISSN 01989715. doi: 10.1016/j.compenvurbsys.2008.09.004.
- Imre Csiszár and Paul C. Shields. Information Theory and Statistics: A Tutorial. *Foundations and Trends™ in Communications and Information Theory*, 1(4):417–528, 2004. ISSN 1567-2190, 1567-2328. doi: 10.1561/0100000004.
- Richard M. Cyert and James G. March. *Behavioral Theory of the Firm*. Wiley-Blackwell, Cambridge, Mass., USA, 2nd edition edition, July 1992. ISBN 978-0-631-17451-6.
- Marcin Czupryna, Christian Franzke, Sascha Hokamp, and Jürgen Scheffran. An Agent-Based Approach to Integrated Assessment Modelling of Climate Change. *Journal of Artificial Societies and Social Simulation*, 23(3):7, 2020. ISSN 1460-7425.
- Zhi Da, Xing Huang, and Lawrence J. Jin. Extrapolative beliefs in the cross-section: What can we learn from the crowds? *Journal of Financial Economics*, 140(1):175–196, April 2021. ISSN 0304-405X. doi: 10.1016/j.jfineco.2020.10.003.
- Yannis Dafermos, Maria Nikolaidi, and Giorgos Galanis. A stock-flow-fund ecological macroeconomic model. *Ecological Economics*, 131:191–207, January 2017. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2016.08.013.
- Yannis Dafermos, Maria Nikolaidi, and Giorgos Galanis. Climate Change, Financial Stability

- and Monetary Policy. *Ecological Economics*, 152:219–234, October 2018. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2018.05.011.
- Bryan C Daniels, Yan-Jiun Chen, James P Sethna, Ryan N Gutenkunst, and Christopher R Myers. Sloppiness, robustness, and evolvability in systems biology. *Current Opinion in Biotechnology*, 19(4):389–395, August 2008. ISSN 09581669. doi: 10.1016/j.copbio.2008.06.008.
- Herbert Dawid and Domenico Delli Gatti. Chapter 2 - Agent-Based Macroeconomics. In Cars Hommes and Blake LeBaron, editors, *Handbook of Computational Economics*, volume 4 of *Handbook of Computational Economics*, pages 63–156. Elsevier, January 2018. doi: 10.1016/bs.hescom.2018.02.006.
- François de Soyres, Ana Maria Santacreu, and Henry Young. Demand-Supply imbalance during the Covid-19 pandemic: The role of fiscal policy. *International Finance Discussion Papers (IFDP)*, August 2022a.
- François de Soyres, Ana Maria Santacreu, and Henry Young. Fiscal policy and excess inflation during Covid-19: A cross-country view. Technical report, July 2022b.
- Marco Del Negro, Marc P. Giannoni, and Frank Schorfheide. Inflation in the Great Recession and New Keynesian Models. *American Economic Journal: Macroeconomics*, 7(1):168–196, January 2015. ISSN 1945-7707. doi: 10.1257/mac.20140097.
- R. Maria del Rio-Chanona, Penny Mealy, Mariano Beguerisse-Díaz, François Lafond, and J. Doyne Farmer. Occupational mobility and automation: A data-driven network model. *Journal of The Royal Society Interface*, 18(174):20200898, 2020. doi: 10.1098/rsif.2020.0898.
- Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, Alberto Russo, and Joseph Stiglitz. The financial accelerator in an evolving credit network. *Journal of Economic Dynamics and Control*, 34(9):1627–1650, September 2010. ISSN 0165-1889. doi: 10.1016/j.jedc.2010.06.019.
- Pinar Deniz and Erhan Aslanoglu. Consumer Confidence in a DSGE Model for Turkey. In *Allied Social Science Associations Annual Meeting*, Philadelphia, PA, 2014. Allied Social Science Associations.
- Francis Dennig, Mark B. Budolfson, Marc Fleurbaey, Asher Siebert, and Robert H. Socolow. Inequality, climate impacts on the future poor, and carbon prices. *Proceedings of the National Academy of Sciences*, 112(52):15827–15832, December 2015. doi: 10.1073/pnas.1513967112.
- Théo Dessertaine, José Moran, Michael Benzaquen, and Jean-Philippe Bouchaud. Out-of-equilibrium dynamics and excess volatility in firm networks. *Journal of Economic Dynamics and Control*, 138:104362, May 2022. ISSN 0165-1889. doi: 10.1016/j.jedc.2022.104362.
- Delavane Diaz and Frances Moore. Quantifying the economic risks of climate change. *Nature Climate Change*, 7(11):774–782, November 2017. ISSN 1758-6798. doi: 10.1038/nclimate3411.
- Simon Dietz, Frederick van der Ploeg, Armon Rezai, and Frank Venmans. Are Economists Getting Climate Dynamics Right and Does It Matter? *Journal of the Association of Environmental and Resource Economists*, 8(5):895–921, September 2021. ISSN 2333-5955, 2333-5963. doi: 10.1086/713977.
- Noah S. Diffenbaugh and Marshall Burke. Global warming has increased global economic inequality. *Proceedings of the National Academy of Sciences*, 116(20):9808–9813, May 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1816020116.

BIBLIOGRAPHY

- Paola D’Orazio. Income Inequality, Consumer Debt, and Prudential Regulation: An Agent-Based Approach to Study the Emergence of Crises and Financial Instability. *Economic Modelling*, 82:308–331, November 2019. ISSN 0264-9993, 0264-9993. doi: 10.1016/j.econmod.2019.01.015.
- Paola D’Orazio and Marco Valente. The role of finance in environmental innovation diffusion: An evolutionary modeling approach. *Journal of Economic Behavior & Organization*, 162: 417–439, June 2019. ISSN 0167-2681. doi: 10.1016/j.jebo.2018.12.015.
- G Dosi, M C Pereira, A Roventini, and M E Virgillito. What If Supply-Side Policies Are Not Enough? The Perverse Interaction of Flexibility and Austerity. *Journal of Economic Behavior and Organization*, 162:360–388, June 2019. ISSN 0167-2681, 0167-2681. doi: 10.1016/j.jebo.2018.11.026.
- Giovanni Dosi and Andrea Roventini. More is different ... and complex! the case for agent-based macroeconomics. *Journal of Evolutionary Economics*, 29(1):1–37, March 2019. ISSN 1432-1386. doi: 10.1007/s00191-019-00609-y.
- Giovanni Dosi, Giorgio Fagiolo, and Andrea Roventini. Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control*, 34(9):1748–1767, September 2010. ISSN 0165-1889. doi: 10.1016/j.jedc.2010.06.018.
- Giovanni Dosi, Giorgio Fagiolo, Mauro Napoletano, and Andrea Roventini. Income distribution, credit and fiscal policies in an agent-based Keynesian model. *Journal of Economic Dynamics and Control*, 37(8):1598–1625, August 2013. ISSN 0165-1889. doi: 10.1016/j.jedc.2012.11.008.
- Giovanni Dosi, Giorgio Fagiolo, Mauro Napoletano, Andrea Roventini, and Tania Treibich. Fiscal and monetary policies in complex evolving economies. *Journal of Economic Dynamics and Control*, 52:166–189, March 2015. ISSN 0165-1889. doi: 10.1016/j.jedc.2014.11.014.
- Giovanni Dosi, Marcelo C Pereira, Andrea Roventini, and Maria Enrica Virgillito. The effects of labour market reforms upon unemployment and income inequalities: An agent-based model. *Socio-Economic Review*, 16(4):687–720, October 2018. ISSN 1475-1461, 1475-147X. doi: 10.1093/ser/mwx054.
- Giovanni Dosi, Francesco Lamperti, Mariana Mazzucato, Mauro Napoletano, and Andrea Roventini. Mission-Oriented Policies and the ”Entrepreneurial State” at Work: An Agent-Based Exploration. May 2021.
- Nepomuk Dunz, Asjad Naqvi, and Irene Monasterolo. Climate sentiments, transition risk, and financial stability in a stock-flow consistent model. *Journal of Financial Stability*, 54:100872, June 2021. ISSN 1572-3089. doi: 10.1016/j.jfs.2021.100872.
- Joel Dyer, Patrick Cannon, J. Doynne Farmer, and Sebastian Schmon. Black-box Bayesian inference for economic agent-based models, February 2022.
- Freeman Dyson. A meeting with Enrico Fermi. *Nature*, 427(6972):297–297, January 2004. ISSN 1476-4687. doi: 10.1038/427297a.
- E Epanechnikov. Non-parametric estimation of a multivariate probability density. *Theory of Probability and its Applications*, 14, 1969.
- Joshua M. Epstein. Agent-based computational models and generative social science. *Complexity*, 4(5):41–60, 1999. ISSN 1099-0526. doi: 10.1002/(SICI)1099-0526(199905/06)4:5<41::AID-CPLX9>3.0.CO;2-F.

-
- Joshua M. Epstein. Chapter 34 Remarks on the Foundations of Agent-Based Generative Social Science. In L. Tesfatsion and K.L. Judd, editors, *Handbook of Computational Economics*, Handbook of Computational Economics, pages 1585–1604. 2006. ISBN 978-0-444-51253-6. doi: 10.1016/S1574-0021(05)02034-4.
- Joshua M. Epstein. Why Model? *Journal of Artificial Societies and Social Simulation*, 11(4), October 2008.
- Giorgio Fagiolo and Andrea Roventini. Macroeconomic Policy in DSGE and Agent-Based Models Redux: New Developments and Challenges Ahead. *Journal of Artificial Societies and Social Simulation*, 20(1):1, 2017. ISSN 1460-7425. doi: 10.18564/jasss.3280.
- Giorgio Fagiolo, Mattia Guerini, Francesco Lamperti, Alessio Moneta, and Andrea Roventini. Validation of Agent-Based Models in Economics and Finance. In Claus Beisbart and Nicole J. Saam, editors, *Computer Simulation Validation: Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*, Simulation Foundations, Methods and Applications, pages 763–787. Springer International Publishing, Cham, 2019. ISBN 978-3-319-70766-2. doi: 10.1007/978-3-319-70766-2_31.
- Giorgio Fagiolo, Daniele Giachini, and Andrea Roventini. Innovation, finance, and economic growth: An agent-based approach. *Journal of Economic Interaction and Coordination*, 15(3): 703–736, July 2020. ISSN 1860-7128. doi: 10.1007/s11403-019-00258-1.
- Ray C. Fair. Has macro progressed? *Journal of Macroeconomics*, 34(1):2–10, March 2012. ISSN 0164-0704. doi: 10.1016/j.jmacro.2010.09.009.
- Ray C Fair. Some important macro points. *Oxford Review of Economic Policy*, 36(3):556–578, December 2020. ISSN 0266-903X. doi: 10.1093/oxrep/graa011.
- J. Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460 (7256):685–686, August 2009. ISSN 0028-0836, 1476-4687. doi: 10.1038/460685a.
- J. Doyne Farmer, Cameron Hepburn, Penny Mealy, and Alexander Teytelboym. A Third Wave in the Economics of Climate Change. *Environmental and Resource Economics*, 62(2):329–357, October 2015. ISSN 1573-1502. doi: 10.1007/s10640-015-9965-2.
- Thomas Ferguson and Servaas Storm. Myth and Reality in the Great Inflation Debate: Supply Shocks and Wealth Effects in a Multipolar World Economy. Technical report, Institute for New Economic Thinking Working Paper Series, January 2023.
- John G. Fernald. Reassessing Longer-Run U.S. Growth: How Low? *Federal Reserve Bank of San Francisco, Working Paper Series*, pages 01–25, August 2016. doi: 10.24148/wp2016-18.
- Nick Fitzpatrick, Timothée Parrique, and Inês Cosme. Exploring degrowth policy proposals: A systematic mapping with thematic synthesis. *Journal of Cleaner Production*, 365:132764, September 2022. ISSN 0959-6526. doi: 10.1016/j.jclepro.2022.132764.
- P Forster, D Huppmann, E Kriegler, L Mundaca, C Smith, J Rogelj, and R Seferian. Mitigation Pathways Compatible with 1.5°C in the Context of Sustainable Development Supplementary Material. 2018.
- Benjamin L. Francis and Mark K. Transtrum. Unwinding the model manifold: Choosing similarity measures to remove local minima in sloppy dynamical systems. *Physical Review E*, 100 (1), 2019. doi: 10.1103/PhysRevE.100.012206.

BIBLIOGRAPHY

- Reiner Franke and Frank Westerhoff. Taking Stock: A Rigorous Modelling of Animal Spirits in Macroeconomics. *Journal of Economic Surveys*, 31(5):1152–1182, 2017. ISSN 1467-6419. doi: 10.1111/joes.12219.
- Milton Friedman. The optimum quantity of money, and other essays. Technical report, 1969.
- Martin Fukac and Adrian Pagan. Issues In Adopting Dsge Models For Use In The Policy Process. *CAMA Working Papers*, (2006-10), March 2006.
- Xavier Gabaix. A Behavioral New Keynesian Model. *American Economic Review*, 110(8):2271–2327, August 2020. ISSN 0002-8282. doi: 10.1257/aer.20162005.
- Giampaolo Gabbi, Giulia Iori, Saqib Jafarey, and James Porter. Financial Regulations and Bank Credit to the Real Economy. *Journal of Economic Dynamics and Control*, 50:117–143, January 2015. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2014.07.002.
- Edoardo Gaffeo, Domenico Delli Gatti, Saul Desiderio, and Mauro Gallegati. Adaptive Microfoundations for Emergent Macroeconomics. *Eastern Economic Journal*, 34(4):441–463, October 2008. ISSN 1939-4632. doi: 10.1057/ej.2008.27.
- Jordi Galí. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton, N.J, 2 edition, 2015. ISBN 978-0-691-13316-4.
- Joao Gama Batista, Jean-Philippe Bouchaud, and Damien Challet. Sudden Trust Collapse in Networked Societies. *Eur. Phys. J. B*, 88(3), 2015.
- Ajay Gambhir, Isabela Butnar, Pei-Hao Li, Pete Smith, and Neil Strachan. A Review of Criticisms of Integrated Assessment Models and Proposed Approaches to Address These, through the Lens of BECCS. *Energies*, 12(9):1747, January 2019. ISSN 1996-1073. doi: 10.3390/en12091747.
- Giancarlo Gandolfo. *Economic Dynamics*. Springer-Verlag, Berlin Heidelberg, 4 edition, 2009. ISBN 978-3-642-03862-4.
- Domenico Delli Gatti, Mauro Gallegati, Bruce C. Greenwald, Alberto Russo, and Joseph Stiglitz. Business fluctuations and bankruptcy avalanches in an evolving network economy. *Journal of Economic Interaction and Coordination*, 4(2):195–212, November 2009. ISSN 1860-7128. doi: 10.1007/s11403-009-0054-x.
- Domenico Delli Gatti, Saul Desiderio, Edoardo Gaffeo, Pasquale Cirillo, and Mauro Gallegati. *Macroeconomics from the Bottom-up*. Springer Science & Business Media, April 2011. ISBN 978-88-470-1971-3.
- Frank W. Geels, Frans Berkhout, and Detlef P. van Vuuren. Bridging analytical approaches for low-carbon transitions. *Nature Climate Change*, 6(6):576–583, June 2016. ISSN 1758-6798. doi: 10.1038/nclimate2980.
- Lena Gerdes, Bernhard Rengs, and Manuel Scholz-Wäckerle. Labor and environment in global value chains: An evolutionary policy study with a three-sector and two-region agent-based macroeconomic model. *Journal of Evolutionary Economics*, 32(1):123–173, January 2022. ISSN 1432-1386. doi: 10.1007/s00191-021-00750-7.
- Fabio Ghironi. Macro needs micro. *Oxford Review of Economic Policy*, 34(1-2):195–218, January 2018. ISSN 0266-903X. doi: 10.1093/oxrep/grx050.

-
- Mario Giampietro, Kozo Mayumi, and Alevgul H. Sorman. *The Metabolic Pattern of Societies: Where Economists Fall Short*. Routledge, 2012.
- Gerd Gigerenzer and Henry Brighton. Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, 1(1):107–143, 2009. ISSN 1756-8765. doi: 10.1111/j.1756-8765.2008.01006.x.
- Andrew Glover, José Mustre-del-Río, and Alice von Ende-Becker. How Much Have Record Corporate Profits Contributed to Recent Inflation? *The Federal Reserve Bank of Kansas City Economic Review*, January 2023. ISSN 01612387. doi: 10.18651/ER/v108n1GloverMustredelRiovonEndeBecker.
- Dhananjay K. Gode and Shyam Sunder. Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1):119–137, 1993. ISSN 0022-3808.
- W. Godley and M. Lavoie. *Monetary Economics: An Integrated Approach to Credit, Money, Income, Production and Wealth*. Springer, December 2006. ISBN 978-0-230-62654-6.
- Richard Goodwin. The Nonlinear Accelerator and the Persistence of Business Cycles. *Econometrica*, 19(1):1–17, 1951. doi: 10.2307/1907905.
- Regis Gourdel, Irene Monasterolo, Nepomuk Dunz, Andrea Mazzocchetti, and Laura Parisi. The double materiality of climate physical and transition risks in the euro area. 2022.
- Jakob Grazzini, Matteo G Richiardi, and Mike Tsionas. Bayesian Estimation of Agent-Based Models. *Journal of Economic Dynamics and Control*, 77:26–47, April 2017. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2017.01.014.
- Megan Greene. Joe Biden’s experiment could revolutionise economic thinking. *Financial Times*, May 2021.
- Robin Greenwood and Andrei Shleifer. Expectations of Returns and Expected Returns. *The Review of Financial Studies*, 27(3):714–746, March 2014. ISSN 0893-9454. doi: 10.1093/rfs/hht082.
- Vladas Griskevicius, Robert Cialdini, and Noah Goldstein. Applying (and Resisting) Peer Influence. *MIT Sloan Management Review*, 49(2), 2008.
- Stanislao Gualdi and Antoine Mandel. Endogenous Growth in Production Networks. *Journal of Evolutionary Economics*, 29(1):91–117, March 2019. ISSN 0936-9937, 0936-9937. doi: 10.1007/s00191-018-0552-x.
- Stanislao Gualdi, Marco Tarzia, Francesco Zamponi, and Jean-Philippe Bouchaud. Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, 50(C):29–61, 2015. ISSN 0165-1889.
- Stanislao Gualdi, Marco Tarzia, Francesco Zamponi, and Jean-Philippe Bouchaud. Monetary policy and dark corners in a stylized agent-based model. *Journal of Economic Interaction and Coordination*, 12(3):507–537, October 2017. ISSN 1860-711X, 1860-7128. doi: 10.1007/s11403-016-0174-z.
- Omar A. Guerrero and Robert L. Axtell. Using Agentization for Exploring Firm and Labor Dynamics. In Sjoukje Osinga, Gert Jan Hofstede, and Tim Verwaart, editors, *Emergent Results of Artificial Economics*, Lecture Notes in Economics and Mathematical Sys-

BIBLIOGRAPHY

- tems, pages 139–150. Springer, Berlin, Heidelberg, 2011. ISBN 978-3-642-21108-9. doi: 10.1007/978-3-642-21108-9_12.
- Maxim Gusev and Dimitri Kroujiline. A Simple Economic Model with Interactions. SSRN Scholarly Paper ID 3552539, Social Science Research Network, Rochester, NY, February 2020.
- Maxim Gusev, Dimitri Kroujiline, Boris Govorkov, Sergey V. Sharov, Dmitry Ushanov, and Maxim Zhilyaev. Predictable markets? A news-driven model of the stock market. *Algorithmic Finance*, 4(1-2):5–51, 2015. ISSN 21585571. doi: 10.3233/AF-150042.
- Ryan N. Gutenkunst. *Sloppiness, Modeling, and Evolution in Biochemical Networks*. PhD thesis, Cornell University, Ithaca, New York, 2008.
- Ryan N. Gutenkunst, Joshua J. Waterfall, Fergal P. Casey, Kevin S. Brown, Christopher R. Myers, and James P. Sethna. Universally Sloppy Parameter Sensitivities in Systems Biology Models. *PLOS Computational Biology*, 3(10):e189, October 2007. ISSN 1553-7358. doi: 10.1371/journal.pcbi.0030189.
- G Haag and W Weidlich. *Concepts and Models of a Quantitative Sociology: The Dynamics of Interacting Populations*. Springer, 1 edition, 1983.
- Helmut Haberl, Martin Schmid, Willi Haas, Dominik Wiedenhofer, Henrike Rau, and Verena Winiwarter. Stocks, flows, services and practices: Nexus approaches to sustainable social metabolism. *Ecological Economics*, 182:106949, April 2021. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2021.106949.
- Sarah Hafner, Annela Anger-Kraavi, Irene Monasterolo, and Aled Jones. Emergence of New Economics Energy Transition Models: A Review. *Ecological Economics*, 177:106779, November 2020. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2020.106779.
- David R. Hagen, Jacob K. White, and Bruce Tidor. Convergence in parameters and predictions using computational experimental design. *Interface Focus*, 3(4):20130008, August 2013. doi: 10.1098/rsfs.2013.0008.
- Andrew G Haldane and Arthur E Turrell. Drawing on Different Disciplines: Macroeconomic Agent-Based Models. *Journal of Evolutionary Economics*, 29(1):39–66, March 2019. ISSN 0936-9937, 0936-9937. doi: 10.1007/s00191-018-0557-5.
- David F Hendry and John N J Muellbauer. The future of macroeconomics: Macro theory and models at the Bank of England. *Oxford Review of Economic Policy*, 34(1-2):287–328, January 2018. ISSN 0266-903X, 1460-2121. doi: 10.1093/oxrep/grx055.
- J. R. Hicks. Mr. Keynes and the "Classics"; A Suggested Interpretation. *Econometrica*, 5(2):147–159, 1937. ISSN 0012-9682. doi: 10.2307/1907242.
- John R. Hicks. *A Contribution to the Theory of the Trade Cycle*. Oxford University Press, Oxford, 1950. ISBN 978-0-19-828416-1.
- Tomohiro Hirano and Joseph Stiglitz. The Wobbly Economy; Global Dynamics with Phase Transitions and State Transitions. *Discussion Papers*, (2204), January 2022.
- A Steven Holland. The changing responsiveness of wages to price-level shocks: Explicit and implicit indexation. *Economic Inquiry*, 26(2):265–279, 1988.

-
- Cars Hommes. Heterogeneous Agent Models in Economics and Finance. In *Handbook of Computational Economics*, pages 1109–1186. Elsevier, 2006.
- Cars Hommes. *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press, Cambridge, 2013. ISBN 978-1-107-01929-4. doi: 10.1017/CBO9781139094276.
- Cars Hommes. Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach. *Journal of Economic Literature*, 59(1):149–219, March 2021. ISSN 0022-0515. doi: 10.1257/jel.20191434.
- Cars Hommes and Blake Dean LeBaron. *Handbook of Computational Economics. Volume 4, Heterogeneous Agent Modeling*, volume 4 of *Handbook of Computational Economics*. North Holland, Amsterdam, 2018. ISBN 978-0-444-64132-8.
- Cars Hommes and Joep Lustenhouwer. Managing Unanchored, Heterogeneous Expectations and Liquidity Traps. *Journal of Economic Dynamics and Control*, 101:1–16, April 2019. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2019.01.004.
- Cars Hommes, Mario He, Sebastian Poledna, Melissa Siqueira, and Yang Zhang. CANVAS: A Canadian Behavioral Agent-Based Model. *Bank of Canada Working Paper*, 2022.
- Mikael Höök and Xu Tang. Depletion of fossil fuels and anthropogenic climate change—A review. *Energy Policy*, 52:797–809, January 2013. ISSN 0301-4215. doi: 10.1016/j.enpol.2012.10.046.
- IPCC. Climate change 2014: Mitigation of climate change. Technical report, 2014.
- IPCC. Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. Technical report, 2018.
- IPCC. *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Cambridge University Press, Cambridge, UK; New York, 2022.
- Ernst Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, February 1925. ISSN 0044-3328. doi: 10.1007/BF02980577.
- Prem Jagadeesan, Karthik Raman, and Arun K. Tangirala. Sloppiness: Fundamental study, new formalism and its application in model assessment. *PLOS ONE*, 18(3):e0282609, March 2023. ISSN 1932-6203. doi: 10.1371/journal.pone.0282609.
- Nathan Jefferson and Jack Fuller. A State-Level Look at U.S. Labor Market Supply and Demand, 2023.
- Harold Jeffreys. *Theory of Probability*. Oxford Classic Texts in the Physical Sciences. Clarendon Press ; Oxford University Press, Oxford [Oxfordshire] : New York, 2nd edition, 1948. ISBN 978-0-19-850368-2.
- Robert Calvert Jump and Paul Levine. Behavioural New Keynesian models. *Journal of Macroeconomics*, 59:59–77, March 2019. ISSN 0164-0704. doi: 10.1016/j.jmacro.2018.11.002.
- Nicholas Kaldor. A Model of the Trade Cycle. *The Economic Journal*, 50(197):78–92, 1940. doi: 10.2307/2225740.

BIBLIOGRAPHY

- Michał Kalecki. A Theory of the Business Cycle. *The Review of Economic Studies*, 4(2):77–97, February 1937. ISSN 0034-6527. doi: 10.2307/2967606.
- Greg Kaplan and Giovanni L. Violante. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239, 2014. ISSN 0012-9682.
- Greg Kaplan and Giovanni L. Violante. Microeconomic Heterogeneity and Macroeconomic Shocks. *Journal of Economic Perspectives*, 32(3):167–194, August 2018. ISSN 0895-3309. doi: 10.1257/jep.32.3.167.
- Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743, March 2018. ISSN 0002-8282. doi: 10.1257/aer.20160042.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling Laws for Neural Language Models, January 2020.
- Patrick J. Kehoe, Virgiliu Midrigan, and Elena Pastorino. Evolution of Modern Business Cycle Models: Accounting for the Great Recession. *Journal of Economic Perspectives*, 32(3):141–166, August 2018. ISSN 0895-3309. doi: 10.1257/jep.32.3.141.
- I. Keppo, I. Butnar, N. Bauer, M. Caspani, O. Edelenbosch, J. Emmerling, P. Fragkos, C. Guivarch, M. Harmsen, J. Lefèvre, T. Le Gallic, M. Leimbach, W. McDowall, J.-F. Mercure, R. Schaeffer, E. Trutnevite, and F. Wagner. Exploring the possibility space: Taking stock of the diverse capabilities and gaps in integrated assessment models. *Environmental Research Letters*, 16(5):053006, April 2021. ISSN 1748-9326. doi: 10.1088/1748-9326/abe5d8.
- John Maynard Keynes. *The General Theory of Employment, Interest, and Money*. Macmillan, London, 1936. ISBN 978-3-319-70344-2.
- Mervyn King and John Kay. *Radical Uncertainty: Decision-making for an Unknowable Future*. Hachette UK, 2020.
- Alan Kirman. Whom or What Does the Representative Individual Represent. *Journal of Economic Perspectives*, 6(2):117–136, 1992. ISSN 0895-3309.
- Alan Kirman. Ants, Rationality, and Recruitment. *The Quarterly Journal of Economics*, 108(1):137–156, February 1993. ISSN 0033-5533. doi: 10.2307/2118498.
- Alan Kirman. Heterogeneity in Economics. *Journal of Economic Interaction and Coordination*, 1(1):89–117, May 2006. ISSN 1860-711X, 1860-7128. doi: 10.1007/s11403-006-0005-8.
- Alan Kirman. The Economic Crisis is a Crisis for Economic Theory. *CEifo Economic Studies*, 56(4):498–535, 2010. ISSN 1610-241X.
- Max Sina Knicker, Karl Naumann-Woleske, Jean-Philippe Bouchaud, and Francesco Zamponi. Post-COVID Inflation & the Monetary Policy Dilemma: An Agent-Based Scenario Analysis. June 2023.
- Frank H. Knight. *Risk, Uncertainty and Profit*. Courier Corporation, March 2012. ISBN 978-0-486-14793-2.
- Gerhard Koekemoer and Jan W.H. Swanepoel. Transformation Kernel Density Estimation with

-
- Applications. *Journal of Computational and Graphical Statistics*, 17(3):750–769, 2008. ISSN 1061-8600.
- A Korinek. Thoughts on DSGE macroeconomics: Matching the moment, but missing the point? In *Economic Theory and Public Policies: Joseph Stiglitz and the Teaching of Economics*. Columbia University Press, 2017.
- Andrey V. Korotayev and Sergey V. Tsirel. A Spectral Analysis of World GDP Dynamics: Kondratieff Waves, Kuznets Swings, Juglar and Kitchin Cycles in Global Economic Development, and the 2008–2009 Economic Crisis. *Structure and Dynamics*, 4(1), January 2010.
- Mark A. Kramer, Herschel Rabitz, and Joseph M. Calo. Sensitivity analysis of oscillatory systems. *Applied Mathematical Modelling*, 8(5):328–340, October 1984. ISSN 0307-904X. doi: 10.1016/0307-904X(84)90146-X.
- Volker Krey, Fei Guo, Peter Kolp, Wenji Zhou, Roberto Schaeffer, Aayushi Awasthy, Christoph Bertram, Harmen-Sytze de Boer, Panagiotis Fragkos, Shinichiro Fujimori, Chenmin He, Gokul Iyer, Kimon Keramidias, Alexandre C. Köberle, Ken Oshiro, Lara Aleluia Reis, Bianka Shoai-Tehrani, Saritha Vishwanathan, Pantelis Capros, Laurent Drouet, James E. Edmonds, Amit Garg, David E.H.J. Gernaat, Kejun Jiang, Maria Kannavou, Alban Kitous, Elmar Kriegler, Gunnar Luderer, Ritu Mathur, Matteo Muratori, Fuminori Sano, and Detlef P. van Vuuren. Looking under the hood: A comparison of techno-economic assumptions across national and global integrated assessment models. *Energy*, 172:1254–1267, April 2019. ISSN 03605442. doi: 10.1016/j.energy.2018.12.131.
- Dimitri Kroujiline, Maxim Gusev, Dmitry Ushanov, Sergey V. Sharov, and Boris Govorkov. Forecasting stock market returns over multiple time horizons. *Quantitative Finance*, 16(11): 1695–1712, November 2016. ISSN 1469-7688, 1469-7696. doi: 10.1080/14697688.2016.1176241.
- Dimitri Kroujiline, Maxim Gusev, Dmitry Ushanov, Sergey V. Sharov, and Boris Govorkov. An endogenous mechanism of business cycles. *Algorithmic Finance*, pages 1–22, October 2019. ISSN 21585571, 21576203. doi: 10.3233/AF-190292.
- Paul Krugman. The Profession and the Crisis. *Eastern Economic Journal*, 37(3):307–312, May 2011. ISSN 0094-5056, 1939-4632. doi: 10.1057/eej.2011.8.
- Theresa Kuchler and Basit Zafar. Personal Experiences and Expectations about Aggregate Outcomes. *The Journal of Finance*, 74(5):2491–2542, 2019. ISSN 1540-6261. doi: 10.1111/jofi.12819.
- S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, March 1951. ISSN 0003-4851. doi: 10.1214/aoms/1177729694.
- Finn E. Kydland and Edward C. Prescott. Time to Build and Aggregate Fluctuations. *Econometrica*, 50(6):1345, November 1982. ISSN 00129682. doi: 10.2307/1913386.
- F. Lamperti, G. Dosi, M. Napoletano, A. Roventini, and A. Sapio. Faraway, So Close: Coupled Climate and Economic Dynamics in an Agent-based Integrated Assessment Model. *Ecological Economics*, 150:315–339, August 2018a. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2018.03.023.
- F. Lamperti, G. Dosi, M. Napoletano, A. Roventini, and A. Sapio. Climate change and green transitions in an agent-based integrated assessment model. *Technological Forecasting and Social Change*, 153:119806, April 2020. ISSN 0040-1625. doi: 10.1016/j.techfore.2019.119806.

BIBLIOGRAPHY

- Francesco Lamperti, Andrea Roventini, and Amir Sani. Agent-Based Model Calibration Using Machine Learning Surrogates. *Journal of Economic Dynamics and Control*, 90:366–389, May 2018b. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2018.03.011.
- Francesco Lamperti, Valentina Bosetti, Andrea Roventini, and Massimo Tavoni. The public costs of climate-induced financial instability. *Nature Climate Change*, 9(11):829–833, November 2019a. ISSN 1758-6798. doi: 10.1038/s41558-019-0607-5.
- Francesco Lamperti, Antoine Mandel, Mauro Napoletano, Alessandro Sapio, Andrea Roventini, Tomas Balint, and Igor Khorenzhenko. Towards agent-based integrated assessment models: Examples, challenges, and future developments. *Regional Environmental Change*, 19(3):747–762, March 2019b. ISSN 1436-3798, 1436-378X. doi: 10.1007/s10113-018-1287-9.
- Francesco Lamperti, Valentina Bosetti, Andrea Roventini, Massimo Tavoni, and Tania Treibich. Three green financial policies to address climate risks. *Journal of Financial Stability*, 54:100875, June 2021. ISSN 1572-3089. doi: 10.1016/j.jfs.2021.100875.
- Raima Larter, Herschel Rabitz, and Mark Kramer. Sensitivity analysis of limit cycles with application to the Brusselator. *The Journal of Chemical Physics*, 80(9):4120–4128, May 1984. ISSN 0021-9606. doi: 10.1063/1.447293.
- Marc Lavoie and Wynne Godley. Features of a Realistic Banking System within a Post-Keynesian Stock-flow Consistent Model. In *Complexity, Endogenous Money and Macroeconomic Theory*, chapter Complexity, Endogenous Money and Macroeconomic Theory. Edward Elgar Publishing, September 2006. ISBN 978-1-84720-311-3.
- Blake LeBaron and Leigh Tesfatsion. Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents. *American Economic Review*, 98(2):246–250, May 2008. ISSN 0002-8282. doi: 10.1257/aer.98.2.246.
- Dain Lee, Jinhyeok Park, and Yongseok Shin. Where Are the Workers? From Great Resignation to Quiet Quitting, January 2023.
- Axel Leijonhufvud. Towards a Not-Too-Rational Macroeconomics. *Southern Economic Journal*, 60(1), 1993.
- Axel Leijonhufvud. *Macroeconomic Instability and Coordination: Selected Essays of Axel Leijonhufvud*. Economists of the Twentieth Century. Edward Elgar, Cheltenham, UK ; Northampton, MA, USA, 2000. ISBN 978-1-85278-967-1.
- Axel Leijonhufvud. Chapter 36 Agent-Based Macro. In L. Tesfatsion and K. L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1625–1637. Elsevier, January 2006. doi: 10.1016/S1574-0021(05)02036-8.
- Timothy M. Lenton, Johan Rockström, Owen Gaffney, Stefan Rahmstorf, Katherine Richardson, Will Steffen, and Hans Joachim Schellnhuber. Climate tipping points — too risky to bet against. *Nature*, 575(7784):592–595, November 2019. doi: 10.1038/d41586-019-03595-0.
- Roberto Leombruni and Matteo Richiardi. Why are economists sceptical about agent-based simulations? *Physica A: Statistical Mechanics and its Applications*, 355(1):103–109, September 2005. ISSN 0378-4371. doi: 10.1016/j.physa.2005.02.072.
- Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944. ISSN 0033-569X, 1552-4485. doi: 10.1090/qam/10666.

-
- J. Lindé, F. Smets, and R. Wouters. Challenges for Central Banks' Macro Models. In *Handbook of Macroeconomics*, volume 2, pages 2185–2262. Elsevier, 2016. ISBN 978-0-444-59487-7. doi: 10.1016/bs.hesmac.2016.04.009.
- Andrew W. Lo. Reading about the Financial Crisis: A Twenty-One-Book Review. *Journal of Economic Literature*, 50(1):151–178, March 2012. ISSN 0022-0515. doi: 10.1257/jel.50.1.151.
- Eduardo López, Omar A. Guerrero, and Robert L. Axtell. A network theory of inter-firm labor flows. *EPJ Data Science*, 9(1):1–41, December 2020. ISSN 2193-1127. doi: 10.1140/epjds/s13688-020-00251-w.
- Hans-Walter Lorenz. *Nonlinear Dynamical Economics and Chaotic Motion*. Springer-Verlag, Berlin Heidelberg, 2 edition, 1993. ISBN 978-3-642-78326-5. doi: 10.1007/978-3-642-78324-1.
- Robert E Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2): 103–124, April 1972. ISSN 00220531. doi: 10.1016/0022-0531(72)90142-1.
- Robert E Lucas. Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46, 1976.
- Robert E. Lucas. *Models of Business Cycles*. Yrjö Jahnsson Lectures. B. Blackwell, Oxford, OX, UK ; Cambridge, Mass., USA, 1987. ISBN 978-0-631-14789-3.
- Robert E Lucas and Thomas J. Sargent. After Keynesian Macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review*, 3(2):1, 1979.
- Thomas Lux. Estimation of Agent-Based Models Using Sequential Monte Carlo Methods. *Journal of Economic Dynamics and Control*, 91:391–408, June 2018. ISSN 0165-1889, 0165-1889.
- Thomas Lux. Bayesian Estimation of Agent-Based Models via Adaptive Particle Markov Chain Monte Carlo. *Computational Economics*, 60(2):451–477, August 2022. ISSN 1572-9974. doi: 10.1007/s10614-021-10155-0.
- C M Macal. Everything you need to know about agent-based modelling and simulation. *Journal of Simulation*, 10(2):144–156, May 2016. ISSN 1747-7778, 1747-7786. doi: 10.1057/jos.2016.7.
- Benjamin B. Machta, Ricky Chachra, Mark K. Transtrum, and James P. Sethna. Parameter Space Compression Underlies Emergent Theories and Predictive Models. *Science*, 342(6158): 604–607, November 2013. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1238723.
- Dougal Maclaurin. *Modeling, Inference and Optimization With Composable Differentiable Procedures*. PhD thesis, Harvard University, 2016.
- Nicholas Magliocca, Virginia McConnell, and Margaret Walls. Integrating Global Sensitivity Approaches to Deconstruct Spatial and Temporal Sensitivities of Complex Spatial Agent-Based Models. *Journal of Artificial Societies and Social Simulation*, 21(1):12, 2018. ISSN 1460-7425. doi: 10.18564/jasss.3625.
- Donald W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, June 1963. ISSN 0368-4245, 2168-3484. doi: 10.1137/0111030.
- John Matsusaka and Argia Sbordone. Consumer Confidence and Economic Fluctuations. *Economic Inquiry*, 33(2):296–318, 1995.

BIBLIOGRAPHY

- H. Damon Matthews, Nathan P. Gillett, Peter A. Stott, and Kirsten Zickfeld. The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248):829–832, June 2009. ISSN 1476-4687. doi: 10.1038/nature08047.
- H. Damon Matthews, Susan Solomon, and Raymond Pierrehumbert. Cumulative carbon as a policy framework for achieving climate stabilization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1974):4365–4379, September 2012. doi: 10.1098/rsta.2012.0064.
- Henry H. Mattingly, Mark K. Transtrum, Michael C. Abbott, and Benjamin B. Machta. Maximizing the information learned from finite data selects a simple model. *Proceedings of the National Academy of Sciences*, 115(8):1760–1765, February 2018. doi: 10.1073/pnas.1715306115.
- Jürgen Mayer, Khaled Khairy, and Jonathon Howard. Drawing an elephant with four complex parameters. *American Journal of Physics*, 78(6):648–649, June 2010. ISSN 0002-9505. doi: 10.1119/1.3254017.
- Andrew McCarthy, Rob Dellink, and Ruben Bibas. The Macroeconomics of the Circular Economy Transition: A Critical Review of Modelling Approaches. OECD Environment Working Papers 130, April 2018.
- Alisdair McKay and Ricardo Reis. The Role of Automatic Stabilizers in the U.S. Business Cycle. *Econometrica*, 84(1):141–194, 2016. ISSN 1468-0262. doi: 10.3982/ECTA11574.
- Patrick Mellacher and Timon Scheuer. Wage Inequality, Labor Market Polarization and Skill-Biased Technological Change: An Evolutionary (Agent-Based) Approach. *Computational Economics*, pages 1–46, August 2020. ISSN 1572-9974. doi: 10.1007/s10614-020-10026-0.
- Jean-Francois Mercure, Hector Pollitt, Neil R. Edwards, Philip B. Holden, Unnada Chewpreecha, Pablo Salas, Aileen Lam, Florian Knobloch, and Jorge E. Vinuales. Environmental impact assessment for climate change policy with the simulation-based integrated assessment model E3ME-FTT-GENIE. *Energy Strategy Reviews*, 20:195–208, April 2018. ISSN 2211-467X. doi: 10.1016/j.esr.2018.03.003.
- Jean-Francois Mercure, Florian Knobloch, Hector Pollitt, Leonidas Paroussos, S. Serban Scricieiu, and Richard Lewney. Modelling innovation and the macroeconomics of low-carbon transitions: Theory, perspectives and practical use. *Climate Policy*, 19(8):1019–1037, September 2019. ISSN 1469-3062. doi: 10.1080/14693062.2019.1617665.
- Silvia Merler. The DSGE Model Quarrel (Again), 2017.
- Bence MÉRŐ, András Borsos, Zsuzsanna Hosszú, Zsolt Oláh, and Nikolett Vágó. A High-Resolution Agent-based Model of the Hungarian Housing Market. *MNB Working Papers*, 7, 2022.
- Simon Michaux. The Mining of Minerals and the Limits to Growth. Technical Report 16, Geological Survey of Finland, Finland, 2021.
- David Miles, Ugo Panizza, Ricardo Reis, and Ángel Ubide. *And yet It Moves: Inflation and the Great Recession*. International Center for Monetary and Banking Studies/Centre for Economic ..., 2017.
- C. Chet Miller and Laura B. Cardinal. Strategic Planning and Firm Performance: A Synthesis of More than Two Decades of Research. *The Academy of Management Journal*, 37(6):1649–1665, 1994. ISSN 0001-4273. doi: 10.2307/256804.

-
- Hyman P. Minsky. *John Maynard Keynes*. Springer, 1976.
- Hyman P. Minsky. *Stabilizing an Unstable Economy*. McGraw-Hill, New York, [NY], neuaufl. edition, 2008. ISBN 978-0-07-159299-4.
- Irene Monasterolo. Climate Change and the Financial System. *Annual Review of Resource Economics*, 12:299–320, 2020. doi: 10.1146/annurev-resource-110119-031134.
- Irene Monasterolo and Marco Raberto. The EIRIN Flow-of-funds Behavioural Model of Green Fiscal Policies and Green Sovereign Bonds. *Ecological Economics*, 144:228–243, February 2018. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2017.07.029.
- Irene Monasterolo, Andrea Roventini, and Tim J Foxon. Uncertainty of Climate Policies and Implications for Economics and Finance: An Evolutionary Economics Approach. *Ecological Economics*, 163:177–182, September 2019. ISSN 0921-8009, 0921-8009. doi: 10.1016/j.ecolecon.2019.05.012.
- Irene Monasterolo, Nepomuk Dunz, Andrea Mazzocchetti, and Régis Gourdel. Derisking the low-carbon transition: Investors’ reaction to climate policies, decarbonization and distributive effects. *Review of Evolutionary Political Economy*, 3(1):31–71, April 2022a. ISSN 2662-6144. doi: 10.1007/s43253-021-00062-3.
- Irene Monasterolo, Andrea Mazzocchetti, and Andrea Vismara. A Question of Trust? Banks’ Climate Sentiments, Lending Behavior and the Low-Carbon Transition, December 2022b.
- José Moran, Antoine Fosset, Michael Benzaquen, and Jean-Philippe Bouchaud. Schrödinger’s ants: A continuous description of Kirman’s recruitment model. *Journal of Physics: Complexity*, 1(3):035002, August 2020. ISSN 2632-072X. doi: 10.1088/2632-072X/aba115.
- José Moran, Antoine Fosset, Alan Kirman, and Michael Benzaquen. From ants to fishing vessels: A simple model for herding and exploitation of finite resources. *Journal of Economic Dynamics and Control*, 129:104169, August 2021. ISSN 0165-1889. doi: 10.1016/j.jedc.2021.104169.
- Federico Morelli, Michael Benzaquen, Marco Tarzia, and Jean-Philippe Bouchaud. Confidence Collapse in a Multi-Household, Self-Reflexive DSGE Model. *PNAS*, 2020.
- Federico Morelli, Karl Naumann-Woleske, Michael Benzaquen, Marco Tarzia, and Jean-Philippe Bouchaud. Economic Crises in a Model with Capital Scarcity and Self-Reflexive Confidence, September 2021.
- Scott Moss. Agent Based Modelling for Integrated Assessment. *Integrated Assessment*, 3(1): 63–77, March 2002. ISSN 1389-5176. doi: 10.1076/iaij.3.1.63.7407.
- Scott Moss, Claudia Pahl-Wostl, and Thomas Downing. Agent-based integrated assessment modelling: The example of climate change. *Integrated Assessment*, 2(1):17–30, March 2001. ISSN 1573-1545. doi: 10.1023/A:1011527523183.
- Daniel Murphy. Excess capacity in a fixed-cost economy. *European Economic Review*, 91:245–260, January 2017. ISSN 00142921. doi: 10.1016/j.eurocorev.2016.11.002.
- John F. Muth. Rational Expectations and the Theory of Price Movements. *Econometrica*, 29 (3):315, July 1961. ISSN 00129682. doi: 10.2307/1909635.
- Mauro Napoletano. A Short Walk on the Wild Side: Agent-Based Models and their Implications

BIBLIOGRAPHY

- for Macroeconomic Analysis. *Revue de l'OFCE*, 157(3):257, 2018. ISSN 1265-9576, 1777-5647. doi: 10.3917/reof.157.0257.
- Asjad Naqvi and Irene Monasterolo. Assessing the cascading impacts of natural disasters in a multi-layer behavioral network framework. *Scientific Reports*, 11(1):20146, October 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-99343-4.
- Karl Naumann-Woleske. Agent-based Integrated Assessment Models: Alternative Foundations to the Environment-Energy-Economics Nexus. January 2023.
- Karl Naumann-Woleske, Michael Benzaquen, Maxim Gusev, and Dimitri Kroujiline. Capital Demand Driven Business Cycles: Mechanism and Effects. *Review of Behavioral Economics*, 2022. doi: 10.1561/105.00000162.
- Karl Naumann-Woleske, Max Sina Knicker, Michael Benzaquen, and Jean-Philippe Bouchaud. Exploration of the Parameter Space in Macroeconomic Agent-Based Models. In *Handbook of Complexity Economics*. 2023.
- Richard R. Nelson and Sidney Winter. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA, 1982. ISBN 978-0-674-04143-1.
- Harald Niederreiter. Low-discrepancy and low-dispersion sequences. *Journal of Number Theory*, 30(1):51–70, September 1988. ISSN 0022314X. doi: 10.1016/0022-314X(88)90025-X.
- Jaime Nieto, Óscar Carpintero, Luis J. Miguel, and Ignacio de Blas. Macroeconomic modelling under energy constraints: Global low carbon transition scenarios. *Energy Policy*, 137:111090, February 2020. ISSN 03014215. doi: 10.1016/j.enpol.2019.111090.
- William Nordhaus. An Optimal Transition Path for Controlling Greenhouse Gases. *Science*, 258(5086):1315–1319, November 1992. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.258.5086.1315.
- William Nordhaus. Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches. *Journal of the Association of Environmental and Resource Economists*, 1(1/2):273–312, March 2014. ISSN 2333-5955, 2333-5963. doi: 10.1086/676035.
- William Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, February 2017. doi: 10.1073/pnas.1609244114.
- William Nordhaus and Joseph Boyer. *Warming the World: Economic Models of Global Warming*. The MIT Press, Cambridge, Mass., revised ed. edition edition, August 2003. ISBN 978-0-262-64054-1.
- C Otto, S N Willner, L Wenz, K Frieler, and A Levermann. Modeling Loss-Propagation in the Global Supply Network: The Dynamic Agent-Based Model Acclimate. *Journal of Economic Dynamics and Control*, 83:232–269, October 2017. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2017.08.001.
- Tolga Ozden. Heterogeneous Expectations and the Business Cycle at the Effective Lower Bound. 2021.
- Scott E. Page. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton Univ. Press, Princeton, NJ, 3. print., and 1. paperback print., with a new preface edition, 2007. ISBN 978-0-691-13854-1 978-0-691-12838-2.

-
- Scott E. Page. *The Model Thinker: What You Need to Know to Make Data Work for You*. Basic Books, New York, NY, 2018. ISBN 978-0-465-09462-2 978-0-465-09463-9.
- Elisa Palagi, Mauro Napoletano, Andrea Roventini, and Jean-Luc Gaffard. Inequality, Redistributive Policies and Multiplier Dynamics in an Agent-based Model with Credit Rationing. *Italian Economic Journal*, 3(3):367–387, November 2017. ISSN 2199-322X, 2199-3238. doi: 10.1007/s40797-017-0055-1.
- Elisa Palagi, Mauro Napoletano, Andrea Roventini, and Jean-Luc Gaffard. An agent-based model of trickle-up growth and income inequality. 2021.
- Georgios Papadopoulos. Income Inequality, Consumption, Credit and Credit Risk in a Data-Driven Agent-Based Model. *Journal of Economic Dynamics and Control*, 104:39–73, July 2019. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2019.05.002.
- Stefan Pauliuk, Anders Arvesen, Konstantin Stadler, and Edgar G. Hertwich. Industrial ecology in integrated assessment models. *Nature Climate Change*, 7(1):13–20, January 2017. ISSN 1758-6798. doi: 10.1038/nclimate3148.
- Gerhard Petschel-Held, Hans-Joachim Schellnhuber, Thomas Bruckner, Ferenc L. Tóth, and Klaus Hasselmann. The Tolerable Windows Approach: Theoretical and Methodological Foundations. *Climatic Change*, 41(3):303–331, March 1999. ISSN 1573-1480. doi: 10.1023/A:1005487123751.
- A.S. Pikovsky and J Kurths. Coherence resonance in a noise-driven excitable system. *Physical Review Letters*, 78:775–778, 1997.
- Robert S. Pindyck. Climate Change Policy: What Do the Models Tell Us? *Journal of Economic Literature*, 51(3):860–872, September 2013. ISSN 0022-0515. doi: 10.1257/jel.51.3.860.
- Robert S. Pindyck. The Use and Misuse of Models for Climate Policy. *Review of Environmental Economics and Policy*, 11(1):100–114, January 2017. ISSN 1750-6816. doi: 10.1093/reep/rew012.
- Donovan Platt. A comparison of economic agent-based model calibration methods. *Journal of Economic Dynamics and Control*, 113:103859, April 2020. ISSN 01651889. doi: 10.1016/j.jedc.2020.103859.
- Donovan Platt. Bayesian Estimation of Economic Simulation Models Using Neural Networks. *Computational Economics*, February 2021. ISSN 1572-9974. doi: 10.1007/s10614-021-10095-9.
- Sebastian Poledna, Michael Gregor Miess, Cars Hommes, and Katrin Rabitsch. Economic forecasting with an agent-based model. *European Economic Review*, 151:104306, January 2023. ISSN 0014-2921. doi: 10.1016/j.euroecorev.2022.104306.
- J. Gareth Polhill, Matthew Hare, Tom Bauermann, David Anzola, Erika Palmer, Doug Salt, and Patrycja Antosz. Using Agent-Based Models for Prediction in Complex and Wicked Systems. *Journal of Artificial Societies and Social Simulation*, 24(3):2, 2021. ISSN 1460-7425.
- Hector Pollitt, Anthony Barker, Jennifer Barton, Elke Pirgmaier, Christine Polzin, Stephan Lutter, Friedrich Hinterberger, and Andrea Stocker. A scoping study on the macroeconomic view of sustainability. *Final report for the European Commission, DG Environment, Cambridge Econometrics and Sustainable Europe Research Institute (July 2010)*, 2010.

BIBLIOGRAPHY

- Linda Ponta, Marco Raberto, Andrea Teglio, and Silvano Cincotti. An Agent-Based Stock-Flow Consistent Model of the Sustainable Transition in the Energy Sector. *Ecological Economics*, 145:274–300, March 2018. ISSN 0921-8009, 0921-8009. doi: 10.1016/j.ecolecon.2017.08.022.
- Lilit Popoyan, Mauro Napoletano, and Andrea Roventini. Taming macroeconomic instability: Monetary and macro-prudential policy interactions in an agent-based model. *Journal of Economic Behavior & Organization*, 134:117–140, February 2017. ISSN 01672681. doi: 10.1016/j.jebo.2016.12.017.
- Lilit Popoyan, Mauro Napoletano, and Andrea Roventini. Winter Is Possibly Not Coming: Mitigating Financial Instability in an Agent-Based Model with Interbank Market. *Journal of Economic Dynamics and Control*, 117, August 2020. ISSN 0165-1889, 0165-1889. doi: 10.1016/j.jedc.2020.103937.
- Christoph Priesmeier and Nikolai Stähler. Long dark shadows or innovative spirits? the effects of (smoothing) business cycles on economic growth: A survey of the literature. *Journal of Economic Surveys*, 25(5):898–912, December 2011. ISSN 0950-0804. doi: 10.1111/j.1467-6419.2011.00694.x.
- Ilya Prigogine and Isabelle Stengers. *Order out of Chaos: Man's New Dialogue with Nature*. Bantam Books, Toronto, 5. [print.] edition, 1988. ISBN 978-0-553-34082-2 978-0-553-34363-2.
- Frederic L. Pryor. *Economic Evolution and Structure: The Impact of Complexity on the U.S. Economic System*. Cambridge University Press, Cambridge, 1996. ISBN 978-1-139-17460-2.
- Arnau Quera-Bofarull, Ayush Chopra, Anisoara Calinescu, Michael Wooldridge, and Joel Dyer. Bayesian calibration of differentiable agent-based models, May 2023a.
- Arnau Quera-Bofarull, Joel Dyer, Anisoara Calinescu, and Michael Wooldridge. Some challenges of calibrating differentiable agent-based models, July 2023b.
- Katherine N. Quinn, Colin B. Clement, Francesco De Bernardis, Michael D. Niemack, and James P. Sethna. Visualizing probabilistic models and data with Intensive Principal Component Analysis. *Proceedings of the National Academy of Sciences*, 116(28):13762–13767, July 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1817218116.
- Katherine N. Quinn, Michael Abbot, Mark K. Transtrum, Benjamin B. Machta, and James P. Sethna. Information geometry for multiparameter models: New perspectives on the origin of simplicity. *Reports on Progress in Physics*, 86, 2023.
- Marco Raberto, Andrea Teglio, and Silvano Cincotti. Integrating Real and Financial Markets in an Agent-Based Economic Model: An Application to Monetary Policy Design. *Computational Economics*, 32(1):147–162, September 2008. ISSN 1572-9974. doi: 10.1007/s10614-008-9138-2.
- Marco Raberto, Bulent Ozel, Linda Ponta, Andrea Teglio, and Silvano Cincotti. From financial instability to green finance: The role of banking and credit market regulation in the Eurace model. *Journal of Evolutionary Economics*, 29(1):429–465, March 2019. ISSN 1432-1386. doi: 10.1007/s00191-018-0568-2.
- Ricardo Reis. Losing the inflation anchor. *Brookings Papers on Economic Activity*, Fall:307–361, September 2021.
- Ricardo Reis. The burst of high inflation in 2021–22: How and why did we get here? 2022.

-
- Severin Reissl. Heterogeneous expectations, forecasting behaviour and policy experiments in a hybrid Agent-based Stock-flow-consistent model. *Journal of Evolutionary Economics*, 31(1): 1–49, 2021. ISSN 1432-1386. doi: 10.1007/s00191-020-00683-7.
- Bernhard Rengs and Manuel Scholz-Wackerle. Consumption and Class in Evolutionary Macroeconomics. *Journal of Evolutionary Economics*, 29(1):229–263, March 2019. ISSN 0936-9937, 0936-9937. doi: 10.1007/s00191-018-0592-2.
- Richard L. Revesz, Peter H. Howard, Kenneth Arrow, Lawrence H. Goulder, Robert E. Kopp, Michael A. Livermore, Michael Oppenheimer, and Thomas Sterner. Global warming: Improve economic models of climate change. *Nature*, 508(7495):173–175, April 2014. ISSN 1476-4687. doi: 10.1038/508173a.
- Armon Rezai and Frederick Van der Ploeg. Intergenerational Inequality Aversion, Growth, and the Role of Damages: Occam’s Rule for the Global Carbon Tax. *Journal of the Association of Environmental and Resource Economists*, 3(2):493–522, June 2016. ISSN 2333-5955. doi: 10.1086/686294.
- Luca Riccetti, Alberto Russo, and Mauro Gallegati. Financial Regulation and Endogenous Macroeconomic Crises. *Macroeconomic Dynamics*, 22(4):896–930, June 2018. ISSN 1365-1005, 1365-1005.
- Luca Riccetti, Alberto Russo, and Mauro Gallegati. Firm–bank credit network, business cycle and macroprudential policy. *Journal of Economic Interaction and Coordination*, February 2021. ISSN 1860-7128. doi: 10.1007/s11403-021-00317-6.
- Paul Romer. *The Trouble With Macroeconomics*. 2016.
- Daniel Rosenbloom, Jochen Markard, Frank W. Geels, and Lea Fuenfschilling. Why carbon pricing is not sufficient to mitigate climate change—and how “sustainability transition policy” can help. *Proceedings of the National Academy of Sciences*, 117(16):8664–8668, April 2020. doi: 10.1073/pnas.2004093117.
- J. Barkley Rosser. On the Complexities of Complex Economic Dynamics. *Journal of Economic Perspectives*, 13(4):169–192, December 1999. ISSN 0895-3309. doi: 10.1257/jep.13.4.169.
- Jeremy B. Rudd. Why do we think that inflation expectations matter for inflation? (And should we?). *Review of Keynesian Economics*, 10(1):25–45, January 2022. ISSN 2049-5331, 2049-5323. doi: 10.4337/roke.2022.01.02.
- Alberto Russo. An Agent Based Macroeconomic Model with Social Classes and Endogenous Crises. *Italian Economic Journal*, 3(3):285–306, November 2017. ISSN 2199-3238. doi: 10.1007/s40797-017-0060-4.
- Alberto Russo, Luca Riccetti, and Mauro Gallegati. Increasing inequality, consumer credit and financial fragility in an agent based macroeconomic model. *Journal of Evolutionary Economics*, 26(1):25–47, March 2016. ISSN 1432-1386. doi: 10.1007/s00191-015-0410-z.
- Karolina Safarzynska. Modeling the rebound effect in two manufacturing industries. *Technological Forecasting and Social Change*, 79(6):1135–1154, July 2012. ISSN 00401625. doi: 10.1016/j.techfore.2012.01.004.
- Karolina Safarzynska and Jeroen C. J. M. van den Bergh. Integrated crisis-energy policy: Macro-evolutionary modelling of technology, finance and energy interactions. *Technology*

BIBLIOGRAPHY

- ical Forecasting and Social Change*, 114:119–137, January 2017a. ISSN 0040-1625. doi: 10.1016/j.techfore.2016.07.033.
- Karolina Safarzynska and Jeroen C J M van den Bergh. ABM-IAM: Optimal climate policy under bounded rationality and multiple inequalities. *Environmental Research Letters*, 17(9): 094022, September 2022. ISSN 1748-9326. doi: 10.1088/1748-9326/ac8b25.
- Karolina Safarzynska and Jeroen C.J.M. van den Bergh. Industry evolution, rational agents and the transition to sustainable electricity production. *Energy Policy*, 39(10):6440–6452, October 2011. ISSN 03014215. doi: 10.1016/j.enpol.2011.07.046.
- Karolina Safarzynska and Jeroen C.J.M. van den Bergh. Financial stability at risk due to investing rapidly in renewable energy. *Energy Policy*, 108:12–20, September 2017b. ISSN 03014215. doi: 10.1016/j.enpol.2017.05.042.
- Isabelle Salle and Murat Yıldızoğlu. Efficient Sampling and Meta-Modeling for Computational Economic Models. *Computational Economics*, 44(4):507–536, December 2014. ISSN 0927-7099, 1572-9974. doi: 10.1007/s10614-013-9406-7.
- Thomas J. Sargent. *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures*. Oxford Univ. Press, Oxford, 1993. ISBN 978-0-19-828869-5.
- José A. Scheinkman. Nonlinearities in Economic Dynamics. *The Economic Journal*, 100(400): 33–48, April 1990. ISSN 0013-0133. doi: 10.2307/2234182.
- Joseph A Schumpeter. *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process*. McGraw Hill, New York, 1939.
- Joseph A. Schumpeter. *Capitalism, Socialism, and Democracy*. Harper and Brothers, New York, 1942.
- Dhruv Sharma, Jean-Philippe Bouchaud, Stanislao Gualdi, Marco Tarzia, and Francesco Zamponi. V-, U-, L-, or W-shaped recovery after COVID? Insights from an Agent Based Model. *PLoS One*, 16(3), July 2020.
- William F. Sharpe. Mutual Fund Performance. *The Journal of Business*, 39(1):119–138, 1966. ISSN 0021-9398.
- Robert J. Shiller. *Irrational Exuberance*. Currency/Doubleday, New York, 2nd ed edition, 2005. ISBN 978-0-7679-2363-7.
- Robert J. Shiller. *Narrative Economics*. 2019. ISBN 978-0-691-18229-2.
- Nate Silver. *The Signal and the Noise*. April 2013.
- Christopher A. Sims. Macroeconomics and Reality. *Econometrica*, 48(1):1, January 1980. ISSN 00129682. doi: 10.2307/1912017.
- Ekaterina Sinitskaya and Leigh Tesfatsion. Macroeconomies as constructively rational games. *Journal of Economic Dynamics and Control*, 61:152–182, December 2015. ISSN 0165-1889. doi: 10.1016/j.jedc.2015.09.011.
- Frantisek Slanina. *Essentials of Econophysics Modelling*. Oxford University Press, Oxford, 2013. ISBN 978-0-19-929968-3. doi: 10.1093/acprof:oso/9780199299683.001.0001.

-
- Frank Smets and Raf Wouters. Forecasting with a Bayesian DSGE model: An application to the euro area. *JCMS: Journal of Common Market Studies*, 42(4):841–867, 2004.
- Frank Smets and Rafael Wouters. Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American economic review*, 97(3):586–606, 2007.
- Sobol. Distribution of Points in a Cube and Approximate Evaluation of Integrals. *Zh. Vych. Mat. Mat. Fiz.*, 7, 1967.
- Robert M. Solow. A Contribution to the Theory of Economic Growth. *The Quarterly Journal of Economics*, 70(1):65, February 1956. ISSN 00335533. doi: 10.2307/1884513.
- Will Steffen, Johan Rockström, Katherine Richardson, Timothy M. Lenton, Carl Folke, Diana Liverman, Colin P. Summerhayes, Anthony D. Barnosky, Sarah E. Cornell, Michel Crucifix, Jonathan F. Donges, Ingo Fetzer, Steven J. Lade, Marten Scheffer, Ricarda Winkelmann, and Hans Joachim Schellnhuber. Trajectories of the Earth System in the Anthropocene. *Proceedings of the National Academy of Sciences*, 115(33):8252–8259, August 2018. doi: 10.1073/pnas.1810141115.
- Rémi Stellan, Gabriel I. Penagos, and Jenny P. Danna-Buitrago. Firms in financial distress: Evidence from inter-firm payment networks with volatility driven by ‘animal spirits’. *Journal of Economic Interaction and Coordination*, 16(1):59–101, January 2021. ISSN 1860-7128. doi: 10.1007/s11403-020-00285-3.
- John Sterman, Thomas Fiddaman, Travis Franck, Andrew Jones, Stephanie McCauley, Philip Rice, Elizabeth Sawin, and Lori Siegel. Climate interactive: The C-ROADS climate policy model. *System Dynamics Review*, 28(3):295–305, 2012. ISSN 1099-1727. doi: 10.1002/sdr.1474.
- John Sterman, Thomas Fiddaman, Travis Franck, Andrew Jones, Stephanie McCauley, Philip Rice, Elizabeth Sawin, and Lori Siegel. Management flight simulators to support climate negotiations. *Environmental Modelling & Software*, 44:122–135, June 2013. ISSN 1364-8152. doi: 10.1016/j.envsoft.2012.06.004.
- Nicholas Stern. The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models. *Journal of Economic Literature*, 51(3):838–859, September 2013. ISSN 0022-0515. doi: 10.1257/jel.51.3.838.
- Nicholas Stern. Economics: Current climate models are grossly misleading. *Nature*, 530(7591):407–409, February 2016. ISSN 1476-4687. doi: 10.1038/530407a.
- Joseph Stiglitz. Where modern macroeconomics went wrong. *Oxford Review of Economic Policy*, 34(1-2):70–106, 2018.
- Joseph Stiglitz and Ira Regmi. The Causes of and Responses to Today’s Inflation. Technical report, Roosevelt Institute, 2022.
- James H. Stock and Mark W. Watson. Business Cycle Fluctuations in U.S. Macroeconomic Time Series. In John B Taylor and Michael Woodford, editors, *Handbook of Macroeconomics*, volume 1A, pages 3–64. Elsevier, New York, 1999. doi: 10.3386/w6528.
- James H. Stock and Mark W. Watson. Chapter 10 Forecasting with Many Predictors. In G. Elliott, C. W. J. Granger, and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 1, pages 515–554. Elsevier, January 2006. doi: 10.1016/S1574-0706(05)01010-4.

BIBLIOGRAPHY

- James P. Stodder. The Evolution of Complexity in Primitive Economies: Theory. *Journal of Comparative Economics*, 20(1):1–31, 1995.
- James P. Stodder. Complexity Aversion: Simplification in the Herrnstein and Allais Behaviors. *Eastern Economic Journal*, 23(1), 1997.
- Lawrence H. Summers. Some skeptical observations on real business cycle theory. *Quarterly Review*, 10(Fall):23–27, 1986.
- Masuo Suzuki and Ryogo Kubo. Dynamics of the Ising Model near the Critical Point. I. *Journal of the Physical Society of Japan*, 24(1):51–60, January 1968. ISSN 0031-9015. doi: 10.1143/JPSJ.24.51.
- Guus ten Broeke, George van Voorn, Arend Ligtenberg, and Jaap Molenaar. The Use of Surrogate Models to Analyse Agent-Based Models. *Journal of Artificial Societies and Social Simulation*, 24(2):3, 2021. ISSN 1460-7425.
- Leigh Tesfatsion. Agent-Based Computational Economics: Growing Economies from the Bottom Up. *SSRN Electronic Journal*, 2002. ISSN SSRN Electronic Journal.
- Leigh Tesfatsion and Kenneth Judd. Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics, April 2006.
- Jan C. Thiele, Winfried Kurth, and Volker Grimm. Facilitating Parameter Estimation and Sensitivity Analysis of Agent-Based Models: A Cookbook Using NetLogo and R. *Journal of Artificial Societies and Social Simulation*, 17(3):11, 2014. ISSN 1460-7425.
- Allan Timmermann. Chapter 4 Forecast Combinations. In G. Elliott, C. W. J. Granger, and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 1, pages 135–196. Elsevier, January 2006. doi: 10.1016/S1574-0706(05)01004-9.
- Mark K. Transtrum. Model Reduction by Manifold Boundaries. *Physical Review Letters*, 113(9), 2014. doi: 10.1103/PhysRevLett.113.098701.
- Mark K. Transtrum, Benjamin B. Machta, and James P. Sethna. Why are Nonlinear Fits to Data so Challenging? *Physical Review Letters*, 104(6):060201, February 2010. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.104.060201.
- Mark K. Transtrum, Benjamin B. Machta, and James P. Sethna. Geometry of nonlinear least squares with applications to sloppy models and optimization. *Physical Review E*, 83(3):036701, March 2011. ISSN 1539-3755, 1550-2376. doi: 10.1103/PhysRevE.83.036701.
- Amos Tversky and Daniel Kahneman. Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131, September 1974. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.185.4157.1124.
- Hiroshi Ugai et al. Effects of the quantitative easing policy: A survey of empirical analyses. *Monetary and economic studies-Bank of Japan*, 25(1):1, 2007.
- J. Van Den Bergh, J. Castro, S. Drews, F. Exadaktylos, J. Foramitti, F. Klein, T. Konc, and I. Savin. Designing an effective climate-policy mix: Accounting for instrument synergy. *Climate Policy*, 21(6):745–764, July 2021. ISSN 1469-3062, 1752-7457. doi: 10.1080/14693062.2021.1907276.

-
- Sander van der Hoog. Surrogate Modelling in (and of) Agent-Based Models: A Prospectus. *Computational Economics*, 53(3):1245–1263, March 2019. ISSN 0927-7099, 0927-7099. doi: 10.1007/s10614-018-9802-0.
- Andrea Vandin, Daniele Giachini, Francesco Lamperti, and Francesca Chiaromonte. Automated and Distributed Statistical Analysis of Economic Agent-Based Models. *Journal of Economic Dynamics and Control*, 143, 2022. doi: 10.1016/j.jedc.2022.104458.
- David Vines and Samuel Wills. The rebuilding macroeconomic theory project: An analytical assessment. *Oxford Review of Economic Policy*, 34(1-2):1–42, January 2018. ISSN 0266-903X, 1460-2121. doi: 10.1093/oxrep/grx062.
- David Vines and Samuel Wills. The rebuilding macroeconomic theory project part II: Multiple equilibria, toy models, and policy models in a new macroeconomic paradigm. *Oxford Review of Economic Policy*, 36(3):427–497, December 2020. ISSN 0266-903X. doi: 10.1093/oxrep/graa066.
- M. P. Wand. Fast Computation of Multivariate Kernel Estimators. *Journal of Computational and Graphical Statistics*, 3(4):433–445, 1994. ISSN 1061-8600. doi: 10.2307/1390904.
- M. P. Wand, J. S. Marron, and D. Ruppert. Transformations in Density Estimation. *Journal of the American Statistical Association*, 86(414):343–353, 1991. ISSN 0162-1459. doi: 10.2307/2290569.
- Takuma Watari, Benjamin C. McLellan, Damien Giurco, Elsa Dominish, Eiji Yamasue, and Keisuke Nansai. Total material requirement for the global energy transition to 2050: A focus on transport and electricity. *Resources, Conservation and Recycling*, 148:91–103, September 2019. ISSN 0921-3449. doi: 10.1016/j.resconrec.2019.05.015.
- Joshua J. Waterfall, Fergal P. Casey, Ryan N. Gutenkunst, Kevin S. Brown, Christopher R. Myers, Piet W. Brouwer, Veit Elser, and James P. Sethna. Sloppy-Model Universality Class and the Vandermonde Matrix. *Physical Review Letters*, 97(15):150601, October 2006. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.97.150601.
- Isabella Weber. Big Oil’s Profits and Inflation: Winners and Losers. *Challenge*, pages 1–9, 2022.
- Isabella Weber and Evan Wasner. Sellers’ Inflation, Profits and Conflict: Why can Large Firms Hike Prices in an Emergency? *Economics Department Working Paper Series*, January 2023.
- Isabella Weber, Jesús Lara, Lucas Teixeira, and Luiza Nassif Pires. Inflation in Times of Overlapping Emergencies: Systemically Significant Prices from an Input-output Perspective. December 2022.
- Martin L. Weitzman. Tail-Hedge Discounting and the Social Cost of Carbon. *Journal of Economic Literature*, 51(3):873–882, September 2013. ISSN 0022-0515. doi: 10.1257/jel.51.3.873.
- Dominik Wiedenhofer, Tomer Fishman, Christian Lauk, Willi Haas, and Fridolin Krausmann. Integrating Material Stock Dynamics Into Economy-Wide Material Flow Accounting: Concepts, Modelling, and Global Application for 1900–2050. *Ecological Economics*, 156:121–133, February 2019. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2018.09.010.
- Dominik Wiedenhofer, Tomer Fishman, Barbara Plank, Alessio Miatto, Christian Lauk, Willi Haas, Helmut Haberl, and Fridolin Krausmann. Prospects for a saturation of humanity’s resource use? An analysis of material stocks and flows in nine world regions from 1900 to

BIBLIOGRAPHY

2035. *Global Environmental Change*, 71:102410, November 2021. ISSN 0959-3780. doi: 10.1016/j.gloenvcha.2021.102410.
- A. Katharina Wilkins, Bruce Tidor, Jacob White, and Paul I. Barton. Sensitivity Analysis for Oscillating Dynamical Systems. *SIAM Journal on Scientific Computing*, 31(4):2706–2732, January 2009. ISSN 1064-8275. doi: 10.1137/070707129.
- Sarah Wolf, Steffen Fürst, Antoine Mandel, Wiebke Lass, Daniel Lincke, Federico Pablo-Martí, and Carlo Jaeger. A multi-agent model of several economic regions. *Environmental Modelling & Software*, 44:25–43, June 2013. ISSN 1364-8152. doi: 10.1016/j.envsoft.2012.12.012.
- Michael Woodford. Convergence in Macroeconomics: Elements of the New Synthesis. *American Economic Journal: Macroeconomics*, 1(1):267–279, January 2009. ISSN 1945-7707. doi: 10.1257/mac.1.1.267.
- Simon Wren-Lewis. Ending the microfoundations hegemony. *Oxford Review of Economic Policy*, 34(1-2):55–69, January 2018. ISSN 0266-903X, 1460-2121. doi: 10.1093/oxrep/grx054.
- Evelyn L. Wright and Jon D. Erickson. Incorporating Catastrophes into Integrated Assessment: Science, Impacts, and Adaptation. *Climatic Change*, 57(3):265–286, 2003. ISSN 01650009. doi: 10.1023/A:1022829706609.
- Takeshi Yagihashi. DSGE Models Used by Policymakers: A Survey. *Discussion papers*, (ron333), October 2020.
- Behnam Zakeri, Katsia Paulavets, Leonardo Barreto-Gomez, Luis Gomez Echeverri, Shonali Pachauri, Benigna Boza-Kiss, Caroline Zimm, Joeri Rogelj, Felix Creutzig, Diana Ürges-Vorsatz, et al. Pandemic, war, and global energy transitions. *Energies*, 15(17):6114, 2022.
- Yi Zhang, Zhe Li, and Yongchao Zhang. Validation and Calibration of an Agent-Based Model: A Surrogate Approach. <https://www.hindawi.com/journals/ddns/2020/6946370/>, January 2020. ISSN 1026-0226.

Appendices

Appendix A

Derivations for Quantifying Sloppiness

This appendix shows the derivations referenced in Chapter 3.

A.1 Derivation of the Hessian for the MSE Loss

Recalling the MSE-Loss function from Eq. (3.5),

$$\mathcal{L}^{MSE}(\Phi, \Phi + \delta\Phi) = \frac{1}{2TS} \sum_{t \in T} \sum_{s \in S} (\mathbf{x}_{s,t}(\Phi) - \mathbf{x}_{s,t}(\Phi + \delta\Phi))^\top \Sigma^{-1} (\mathbf{x}_{s,t}(\Phi) - \mathbf{x}_{s,t}(\Phi + \delta\Phi)),$$

I now compute the Hessian with respect to the log parameters $\log \Phi$, whose definition is (Eq. (3.1))

$$H_{i,j}^{\mathcal{L}}(\Phi) = \Phi_i \Phi_j \left. \frac{d^2 \mathcal{L}^{MSE}(\Phi, \Phi + \delta\Phi)}{d\Phi_i d\Phi_j} \right|_{\delta\Phi=0}.$$

Define first the $K \times P$ Jacobian matrix

$$J_{s,t}(\Phi) = \left[\frac{d\mathbf{x}_{s,t}(\Phi)}{d\Phi_1}, \dots, \frac{d\mathbf{x}_{s,t}(\Phi)}{d\Phi_P} \right] = \begin{bmatrix} \frac{d\mathbf{x}_{1,s,t}(\Phi)}{d\Phi_1} & \dots & \frac{d\mathbf{x}_{1,s,t}(\Phi)}{d\Phi_P} \\ \vdots & \ddots & \vdots \\ \frac{d\mathbf{x}_{K,s,t}(\Phi)}{d\Phi_1} & \dots & \frac{d\mathbf{x}_{K,s,t}(\Phi)}{d\Phi_P} \end{bmatrix}$$

Using this definition, the first derivative is

$$\left. \frac{d}{d\Phi} \mathcal{L}^{MSE}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} = \frac{1}{TS} \sum_{t \in T} \sum_{s \in S} J_{s,t}^\top(\Phi) \Sigma^{-1} (\mathbf{x}_{s,t}(\Phi) - \mathbf{x}_{s,t}(\Phi + \delta\Phi))$$

which leads to a second derivative

$$\left. \frac{d^2}{d\Phi d\Phi} \mathcal{L}^{MSE}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} = \frac{1}{TS} \sum_{t \in T} \sum_{s \in S} J_{s,t}^\top(\Phi) \Sigma^{-1} J_{s,t}(\Phi)$$

where I make use of the fact that $\Phi + \delta\Phi$ is the object of the differentiation while the reference Φ , and thus $\mathbf{x}_{s,t}(\Phi)$, is considered a constant. By definition of $\delta\Phi = 0$ we are also at a minimum implying the second-order terms vanish. This implies

$$H^{MSE}(\Phi) = \Phi\Phi^\top \frac{1}{TS} \sum_{t \in T} \sum_{s \in S} J_{s,t}^\top(\Phi) \Sigma^{-1} J_{s,t}(\Phi),$$

where the first term accounts for the log-derivative.

A.2 Hessian for the sKL divergence

Recalling the symmetrized Kullback-Leibler divergence from Eq. (3.7),

$$\begin{aligned} \mathcal{L}^{sKL}(\Phi, \Phi + \delta\Phi) &= \frac{1}{2} (\mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi) + \mathcal{L}^{KL}(\Phi + \delta\Phi, \Phi)) \\ \mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi) &= \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \log \left(\frac{P(\mathbf{x}|\Phi)}{P(\mathbf{x}|\Phi + \delta\Phi)} \right), \end{aligned}$$

we derive the Hessian matrix

$$H_{i,j}^{\mathcal{L}}(\Phi) = \Phi_i \Phi_j \left. \frac{d^2 \mathcal{L}^{sKL}(\Phi, \Phi + \delta\Phi)}{d\Phi_i d\Phi_j} \right|_{\delta\Phi=0}.$$

Beginning with $\mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi)$ and noting that $P(\mathbf{x}|\Phi)$ is a constant as we work with $\Phi + \delta\Phi$,

$$\begin{aligned} \left. \frac{d}{d\Phi_j} \mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} &= \frac{d}{d\Phi_j} \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \log \left(\frac{P(\mathbf{x}|\Phi)}{P(\mathbf{x}|\Phi + \delta\Phi)} \right) \\ &= - \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{d \log P(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j} \\ &= - \sum_{\mathbf{x}} \frac{P(\mathbf{x}|\Phi)}{P(\mathbf{x}|\Phi + \delta\Phi)} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j}, \end{aligned}$$

from which one obtains as second derivative

$$\begin{aligned} \left. \frac{d^2}{d\Phi_i d\Phi_j} \mathcal{L}^{KL}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} &= - \frac{d}{d\Phi_i} \sum_{\mathbf{x}} \frac{P(\mathbf{x}|\Phi)}{P(\mathbf{x}|\Phi + \delta\Phi)} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j} \\ &= - \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \left(\frac{\frac{d^2 P(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_i d\Phi_j}}{P(\mathbf{x}|\Phi + \delta\Phi)} - \frac{\frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_i} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j}}{P(\mathbf{x}|\Phi + \delta\Phi)^2} \right) \\ &= \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{d \log P(\mathbf{x}|\Phi)}{d\Phi_i} \frac{d \log P(\mathbf{x}|\Phi)}{d\Phi_j} \\ &= g_{i,j}(\Phi) d\Phi_i d\Phi_j, \end{aligned}$$

which is the Fisher Information Matrix. Note that I have used the following property for $\delta\Phi = 0$

$$-\sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \left. \frac{\frac{d^2 P(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_i d\Phi_j}}{P(\mathbf{x}|\Phi + \delta\Phi)} \right|_{\delta\Phi=0} = \frac{d^2}{d\Phi_i d\Phi_j} \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) = \frac{d^2}{d\Phi_i d\Phi_j} 1 = 0.$$

We can now repeat this procedure for the second part of the SKL, $\mathcal{L}^{KL}(\Phi + \delta\Phi, \Phi)$,

$$\begin{aligned} \left. \frac{d}{d\Phi_j} \mathcal{L}^{KL}(\Phi + \delta\Phi, \Phi) \right|_{\delta\Phi=0} &= \frac{d}{d\Phi_j} \sum_{\mathbf{x}} P(\mathbf{x}|\Phi + \delta\Phi) \log \left(\frac{P(\mathbf{x}|\Phi + \delta\Phi)}{P(\mathbf{x}|\Phi)} \right) \\ &= \sum_{\mathbf{x}} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j} \log \left(\frac{P(\mathbf{x}|\Phi + \delta\Phi)}{P(\mathbf{x}|\Phi)} \right) + \underbrace{\sum_{\mathbf{x}} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j}}_{=0} \\ &= \sum_{\mathbf{x}} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j} \log \left(\frac{P(\mathbf{x}|\Phi + \delta\Phi)}{P(\mathbf{x}|\Phi)} \right) \end{aligned}$$

from which one obtains also obtains the FIM as a second derivative

$$\begin{aligned} \left. \frac{d^2}{d\Phi_i d\Phi_j} \mathcal{L}^{KL}(\Phi + \delta\Phi, \Phi) \right|_{\delta\Phi=0} &= \frac{d}{d\Phi_i} \sum_{\mathbf{x}} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j} \log \left(\frac{P(\mathbf{x}|\Phi + \delta\Phi)}{P(\mathbf{x}|\Phi)} \right) \\ &= \sum_{\mathbf{x}} \frac{d^2 P(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_i d\Phi_j} \log \left(\frac{P(\mathbf{x}|\Phi + \delta\Phi)}{P(\mathbf{x}|\Phi)} \right) \\ &\quad + \sum_{\mathbf{x}} \frac{1}{P(\mathbf{x}|\Phi + \delta\Phi)} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_i} \\ &= \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{\frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_i} \frac{dP(\mathbf{x}|\Phi + \delta\Phi)}{d\Phi_j}}{P(\mathbf{x}|\Phi + \delta\Phi) P(\mathbf{x}|\Phi + \delta\Phi)} \\ &= \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{d \log P(\mathbf{x}|\Phi)}{d\Phi_j} \frac{d \log P(\mathbf{x}|\Phi)}{d\Phi_i} \\ &= g_{i,j}(\Phi) d\Phi_i d\Phi_j, \end{aligned}$$

where $\log \left(\frac{P(\mathbf{x}|\Phi + \delta\Phi)}{P(\mathbf{x}|\Phi)} \right) = 0$ for $\delta\Phi = 0$

This leads to a final Hessian

$$\boxed{H_{i,j}^{SKL}(\Phi) = \Phi_i \Phi_j g_{i,j}(\Phi) d\Phi_i d\Phi_j}$$

A.3 Hessian for the Hellinger divergence

The Hellinger distance between to probability distributions is

$$\mathcal{L}^H(\Phi, \Phi + \delta\Phi) = \frac{1}{\sqrt{2}} \sqrt{\sum_{\mathbf{x} \in \mathcal{X}} \left(\sqrt{P(\mathbf{x}|\Phi)} - \sqrt{P(\mathbf{x}|\Phi + \delta\Phi)} \right)^2}. \quad (\text{A.1})$$

For simplicity, I denote temporarily $F = P(\mathbf{x}|\Phi)$ and $G = P(\mathbf{x}|\Phi + \delta\Phi)$, where F is treated as a constant in the derivation (only the $\Phi + \delta\Phi$ is a variable).

Before starting, consider that the Hellinger distance can be re-written as

$$\mathcal{L}^H(\Phi, \Phi + \delta\Phi) = \sqrt{\frac{1}{2} \sum_{\mathbf{x}} (\sqrt{F} - \sqrt{G})^2} = \sqrt{\sum_{\mathbf{x}} \left(\frac{F+G}{2} - \sqrt{FG} \right)} = \sqrt{1 - \sum_{\mathbf{x}} \sqrt{FG}}$$

using the above.

Secondly, note that the squared Hellinger distance, $\mathcal{L}^{H^2}(\Phi, \Phi + \delta\Phi) = (\mathcal{L}^H(\Phi, \Phi + \delta\Phi))^2$, is a f -divergence, as is the KL-divergence that we treated earlier. It turns out to be quite general that these divergences behave quadratically for very small perturbations, and that their local measure is the Fisher Information Metric, $g(\Phi)$. The difference between the squared and non-squared Hellinger divergence is that the regular Hellinger divergence is a metric, it obeys the triangle inequality, while the squared Hellinger distance, much like the KL-divergence does not. To illustrate how one obtains the FIM with the squared Hellinger distance, I derive below the Hessian ($\delta\Phi = 0$).

The first derivative is computed simply,

$$\left. \frac{d}{d\Phi_j} \mathcal{L}^{H^2}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} = \sum_{\mathbf{x}} \frac{\sqrt{F}}{2\sqrt{G}} \frac{dG}{d\Phi_j} = \frac{1}{2} \sum_{\mathbf{x}} \sqrt{FG} \frac{d \log G}{d\Phi_j},$$

and leads to a second derivative

$$\left. \frac{d^2}{d\Phi_j d\Phi_i} \mathcal{L}^{H^2}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} = \frac{d}{d\Phi_i} \frac{1}{2} \sum_{\mathbf{x}} \frac{\sqrt{F}}{\sqrt{G}} \frac{dG}{d\Phi_j}$$

where, as with the KL-divergence derivation, letting $\delta\Phi = 0$ and reverting to our regular notation we obtain

$$\left. \frac{d^2}{d\Phi_j d\Phi_i} \mathcal{L}^{H^2}(\Phi, \Phi + \delta\Phi) \right|_{\delta\Phi=0} = \frac{1}{4} \sum_{\mathbf{x}} P(\mathbf{x}|\Phi) \frac{d \log P(\mathbf{x}|\Phi)}{d\Phi_j} \frac{d \log P(\mathbf{x}|\Phi)}{d\Phi_i},$$

which leads us to the familiar FIM structure (with respect to log-parameters):

$$H_{i,j}^{H^2}(\Phi) = \frac{1}{4} \Phi_i \Phi_j g_{i,j}(\Phi)$$

I have used the same property of the second-derivative as with the KL.

In most of the literature, it appears that for probabilistic models, authors choose the symmetrized KL divergence as a measure rather than the Hellinger distance. In Quinn et al. (2023), the authors chose the KL-divergence over the Hellinger distance when developing embeddings for visualizing model manifolds and their hyperribbon structure. The reason being that as their prediction space dimension increases, distributions become more distinguishable from one another, becoming orthogonal. In this case, using the Hellinger distance, all of the pairwise distances in a larger data-space converge to 1, the upper limit of the Hellinger distance, whereas a SKL approach does not have such a restriction.

Derivations for Kirman's Ants

B.1 The Kullback-Leibler Divergence

Consider Kirman's Ants, where the stationary distribution $P(\mathbf{x}|\Phi) \sim \text{Beta}\left(\frac{\rho}{\mu}, \frac{\rho}{\mu}\right)$, then the Kullback-Leibler divergence for a perturbation $P(\mathbf{x}|\Phi + \delta\Phi)$ is given by:

$$KL(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) = \int_0^1 P(\mathbf{x}|\Phi) \log \frac{P(\mathbf{x}|\Phi)}{P(\mathbf{x}|\Phi^*)},$$

which can be decomposed into the difference between the crossentropy $H(P(\mathbf{x}|\Phi), P(\mathbf{x}|\Phi^*))$ between $P(\mathbf{x}|\Phi)$ and $P(\mathbf{x}|\Phi^*)$ and the entropy $H(P(\mathbf{x}|\Phi))$ of $P(\mathbf{x}|\Phi)$,

$$\begin{aligned} KL(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) &= \int_0^1 P(\mathbf{x}|\Phi) \log \frac{1}{P(\mathbf{x}|\Phi^*)} - \int_0^1 P(\mathbf{x}|\Phi) \log \frac{1}{P(\mathbf{x}|\Phi)} \\ &= H(P(\mathbf{x}|\Phi), P(\mathbf{x}|\Phi^*)) - H(P(\mathbf{x}|\Phi)), \end{aligned}$$

for which there is an analytical solution using the digamma function $\psi(\cdot)$,

$$\begin{aligned} H(P(\mathbf{x}|\Phi)) &= - \int_0^1 P(\mathbf{x}|\Phi) \log P(\mathbf{x}|\Phi) dx \\ &= \log \left(B \left(\frac{\rho}{\mu}, \frac{\rho}{\mu} \right) \right) - 2 \left(\frac{\rho}{\mu} - 1 \right) \psi \left(\frac{\rho}{\mu} \right) + 2 \left(\frac{\rho}{\mu} - 1 \right) \psi \left(2 \frac{\rho}{\mu} \right) \\ H(P(\mathbf{x}|\Phi), P(\mathbf{x}|\Phi^*)) &= - \int_0^1 P(\mathbf{x}|\Phi) \log P(\mathbf{x}|\Phi^*) dx \\ &= \log \left(B \left(\frac{\rho^*}{\mu^*}, \frac{\rho^*}{\mu^*} \right) \right) - 2 \left(\frac{\rho^*}{\mu^*} - 1 \right) \psi \left(\frac{\rho}{\mu} \right) + 2 \left(\frac{\rho^*}{\mu^*} - 1 \right) \psi \left(2 \frac{\rho}{\mu} \right), \end{aligned}$$

where $B(\cdot)$ is the Beta function. This leads to a KL divergence of

$$KL(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) = \log \left(\frac{B\left(\frac{\rho^*}{\mu^*}, \frac{\rho^*}{\mu^*}\right)}{B\left(\frac{\rho}{\mu}, \frac{\rho}{\mu}\right)} \right) + 2 \left(\frac{\rho}{\mu} - \frac{\rho^*}{\mu^*} \right) \left(\psi\left(\frac{\rho}{\mu}\right) - \psi\left(2\frac{\rho}{\mu}\right) \right) \quad (\text{B.1})$$

and a symmetric divergence

$$\begin{aligned} KL_{sym}(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) &= KL(P(\mathbf{x}|\Phi) || P(\mathbf{x}|\Phi^*)) + KL(P(\mathbf{x}|\Phi^*) || P(\mathbf{x}|\Phi)) \\ &= 2 \left(\frac{\rho}{\mu} - \frac{\rho^*}{\mu^*} \right) \left(\psi\left(\frac{\rho}{\mu}\right) - \psi\left(2\frac{\rho}{\mu}\right) + \psi\left(\frac{\rho^*}{\mu^*}\right) - \psi\left(2\frac{\rho^*}{\mu^*}\right) \right). \end{aligned} \quad (\text{B.2})$$

B.2 Hessian Matrices for Kirman's Ants

Hessian for the symmetric KL-divergence

Based on Eq. (3.7) and the continuous formulation of the ants model (recalling that $x \in [0, 1]$), the Hessian matrix for the symmetric KL-divergence can be written as

$$H_{i,j}^{\mathcal{L}}(\Phi) := \frac{d^2 \mathcal{L}(\Phi, \Phi + \delta\Phi)}{d \log \Phi_i d \log \Phi_j} \Big|_{\varepsilon_i=0} = \Phi_i \Phi_j \int_0^1 P(\mathbf{x}|\Phi) \frac{d \log P(\mathbf{x}|\Phi)}{d \Phi_i} \frac{d \log P(\mathbf{x}|\Phi)}{d \Phi_j} dx \quad (\text{B.3})$$

Since we have a continuous formulation, I replace the summation with the integral over the domain of x . Noting here that

$$\begin{aligned} P(\mathbf{x}|\Phi) &= \frac{1}{B\left(\frac{\rho}{\mu}, \frac{\rho}{\mu}\right)} (x(1-x))^{\frac{\rho}{\mu}-1} \\ \log P(\mathbf{x}|\Phi) &= -\log B\left(\frac{\rho}{\mu}, \frac{\rho}{\mu}\right) + \left(\frac{\rho}{\mu} - 1\right) \log(x(1-x)) \end{aligned}$$

where $B(\cdot)$ is the Beta function. Consequently, we can note that the first derivatives of $\log P(\mathbf{x}|\Phi)$ are

$$\partial_i \log P(\mathbf{x}|\Phi) = \left(\partial_i \frac{\rho}{\mu} \right) \left(2\psi\left(2\frac{\rho}{\mu}\right) - 2\psi\left(\frac{\rho}{\mu}\right) + \log(x(1-x)) \right)$$

with $\psi(\cdot)$ being the digamma function. This implies that the Hessian of the ant recruitment model at any point in the (ρ, μ) space is given by

$$\begin{aligned} H_{i,j}^{\mathcal{L}}(\Phi) &= \rho\mu \left(\partial_i \frac{\rho}{\mu} \right) \left(\partial_j \frac{\rho}{\mu} \right) \int_0^1 \left(2\psi\left(2\frac{\rho}{\mu}\right) - 2\psi\left(\frac{\rho}{\mu}\right) + \log(x(1-x)) \right)^2 P(\mathbf{x}|\Phi) dx \\ &= \frac{\rho^2}{\mu^2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \int_0^1 \left(2\psi\left(2\frac{\rho}{\mu}\right) - 2\psi\left(\frac{\rho}{\mu}\right) + \log(x(1-x)) \right)^2 P(\mathbf{x}|\Phi) dx \end{aligned}$$

where the integral term is a constant that scales the hessian.

Defining now

$$C = \psi\left(\frac{\rho}{\mu}\right) - \psi\left(2\frac{\rho}{\mu}\right)$$

I simplify

$$\begin{aligned} & \int_0^1 P(\mathbf{x}|\Phi) (\log(x(1-x)) - 2C)^2 dx \\ &= (4C^2 - 4C\mathbb{E}_{P(\mathbf{x}|\Phi)}(\log x(1-x)) + \mathbb{E}_{P(\mathbf{x}|\Phi)}(\log^2 x(1-x))) \\ &= (-4C^2 + \mathbb{E}_{P(\mathbf{x}|\Phi)}(\log^2 x) + 2\mathbb{E}_{P(\mathbf{x}|\Phi)}(\log x \log(1-x)) + \mathbb{E}_{P(\mathbf{x}|\Phi)}(\log^2(1-x))) \\ &= 2\psi^{(1)}\left(\frac{\rho}{\mu}\right) - 4\psi^{(1)}\left(2\frac{\rho}{\mu}\right) \end{aligned}$$

where $\psi^{(1)}\left(\frac{\rho}{\mu}\right)$ is the trigamma function (first derivative of gamma). Recombining all of these terms yields the explicit Hessian

$$\boxed{H_{i,j}^{\mathcal{L}}(\Phi) = \frac{\rho^2}{\mu^2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \left(2\psi^{(1)}\left(\frac{\rho}{\mu}\right) - 4\psi^{(1)}\left(2\frac{\rho}{\mu}\right) \right)} \quad (\text{B.4})$$

Hessian for the Mean-Squared Error

The Hessian for the Mean-Squared Error in the $S \rightarrow \infty$ case reads

$$H_{i,j}^{MSE}(\Phi) = \Phi_i \Phi_j \frac{d \log V(x|\Phi)}{d\Phi_i} \frac{d \log V(x|\Phi)}{d\Phi_j},$$

Noting that for the symmetric Beta distribution, the variance is

$$\begin{aligned} V(x|\Phi) &= \frac{1}{4} \left(2\frac{\rho}{\mu} + 1 \right)^{-1} \\ \log V(x|\Phi) &= -4 - \log \left(2\frac{\rho}{\mu} + 1 \right) \end{aligned}$$

Starting with the first derivatives,

$$\begin{aligned} \frac{d}{d\rho} \log V(x|\Phi) &= -\frac{1}{2\mu} \left(2\frac{\rho}{\mu} + 1 \right)^{-1} \\ \frac{d}{d\mu} \log V(x|\Phi) &= \frac{\rho}{2\mu^2} \left(2\frac{\rho}{\mu} + 1 \right)^{-1} \end{aligned}$$

the Hessian becomes

$$\boxed{H^{MSE}(\Phi) = \frac{\rho^2}{4\mu^2} \left(2\frac{\rho}{\mu} + 1 \right)^{-2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}} \quad (\text{B.5})$$

Appendix C

Numerical Estimation of MSE-based Hessians

This appendix deals with the tricky problem of estimating the Jacobian matrix $J_{s,t}(\Phi)$ for the Mark-0 model by means of finite differences. The goal is to obtain a good approximation of the Hessian of the underlying data generating process at a particular point in the parameter space. To verify the quality of estimation, I consider the property that when perturbing the model in the directions of the eigenvectors of the observed Hessian, one observes a hierarchy of change: the largest change occurs in the direction of the eigenvector with the largest associated eigenvalue (the well-constrained) direction, and sequentially thereafter. As outlined in Chapter 5, the difficulty lies in the separation of the signal from a change in the parameters from the inherent stochasticity of the model.

To compute the finite difference (FD) Hessian, I use central difference derivatives

$$\frac{d\mathbf{x}_{s,t}(\Phi)}{d\log \Phi_i} = \frac{\mathbf{x}_{s,t}(\Phi^+) - \mathbf{x}_{s,t}(\Phi^-)}{2\varepsilon}, \quad (\text{C.1})$$

where parameters are perturbed in the log parameter space to negate effects from naturally differing scales: $\Phi_j^\pm = \Phi \ \forall j \neq i$ and $\Phi_i^\pm = \exp(\log \Phi_i \pm \varepsilon)$. For this approach to be *correct* one needs to carefully select the perturbation size ε . In models with a smooth manifold around Φ , one may retrieve the expected behavior that as the perturbation size ε decreases one obtains a better estimate of the local Hessian, up to a point where numerical precision becomes an issue. However, for stochastic models like Mark-0 one observes the reverse case because the manifold around a given point in parameter-space is quite rough due to the noise effects in the model.

What are these noise effects I keep mentioning? An example from Mark-0 would be firm entry and exit from the model. Firm entry in particular is a stochastic process, with on average, 10% of currently bankrupt firms reentering the model. Consider now what

happens when the model is perturbed in any direction: some firms that previously went bankrupt may not go bankrupt in this new run, or more likely, go bankrupt at a slightly different point in time. Due to this, also the entry of firms may be shifted, with a firm that would have entered at time t may now do so later, or not at all if it doesn't go bankrupt. In this regard, when there is a regular stream of firms entering and leaving bankruptcy, a change in the parameters might shift around the firms entry and exit, thus changing the dynamics, even if the time-realisation of all noise is fixed for every time and firm. The resulting challenge is that there exists some amount of "baseline loss" from any parameter perturbation, such that in order to obtain a Hessian estimate that identifies the hierarchy of parameter combinations in the model one needs to select a finite difference step-size that perturbs the model beyond its baseline noise amounts.

To show this numerically, consider the same three points as in Chapter 5, and the estimation of their Hessian matrices for variable finite difference step-sizes. Figures C.1, C.2, and C.3 show the MSEs of perturbing the parameters into the first ten eigenvector directions (x-axis) for various distances (vertical panels). Each figure represents, Each box represents 50 perturbations into the $v_i(\Phi)$ direction of a $S = 1$ Hessian using the same random realisation that the Hessian was computed on. Considering the steps with similar magnitudes as the finite difference case, one sees that for a FD step $\varepsilon = 10^{-3}$ only the first eigenvector of FE (center) has a significantly larger impact than the remainder. For both FE and RU, irrespective of the choice of eigenvector, there is a similar degree of change in the model. Increasing the step-size to $\varepsilon = 10^{-2}$, the RU phase begins exhibiting the targeted hierarchy of MSE effects, together with the FE phase. However, it is only once $\varepsilon = 2 \times 10^{-1}$ that across all phases one can see the hierarchy of directional effects. In all cases the hierarchy also exists primarily for the first two eigenvectors moreso than the remainder, whose effects rever to a baseline minimal MSE from any for of parameter change (compare the top panel's magnitudes to the last panel's magnitudes in Figure C.3).

One can see this effect also in the eigenvalues, as shown in Figure C.4, where only for large stepsizes the first eigenvalues begin separating and stabilizing from the pack. This also suggests that a future test for relevance might be to compare the eigenvalue distribution to that of a Random Matrix. That is, developing a hypothesis test for the H_0 that $\lambda_i(\Phi)$ is random by, say, comparing it to a Marchenko-Pastur distribution with similar parameters. This is currently under development.

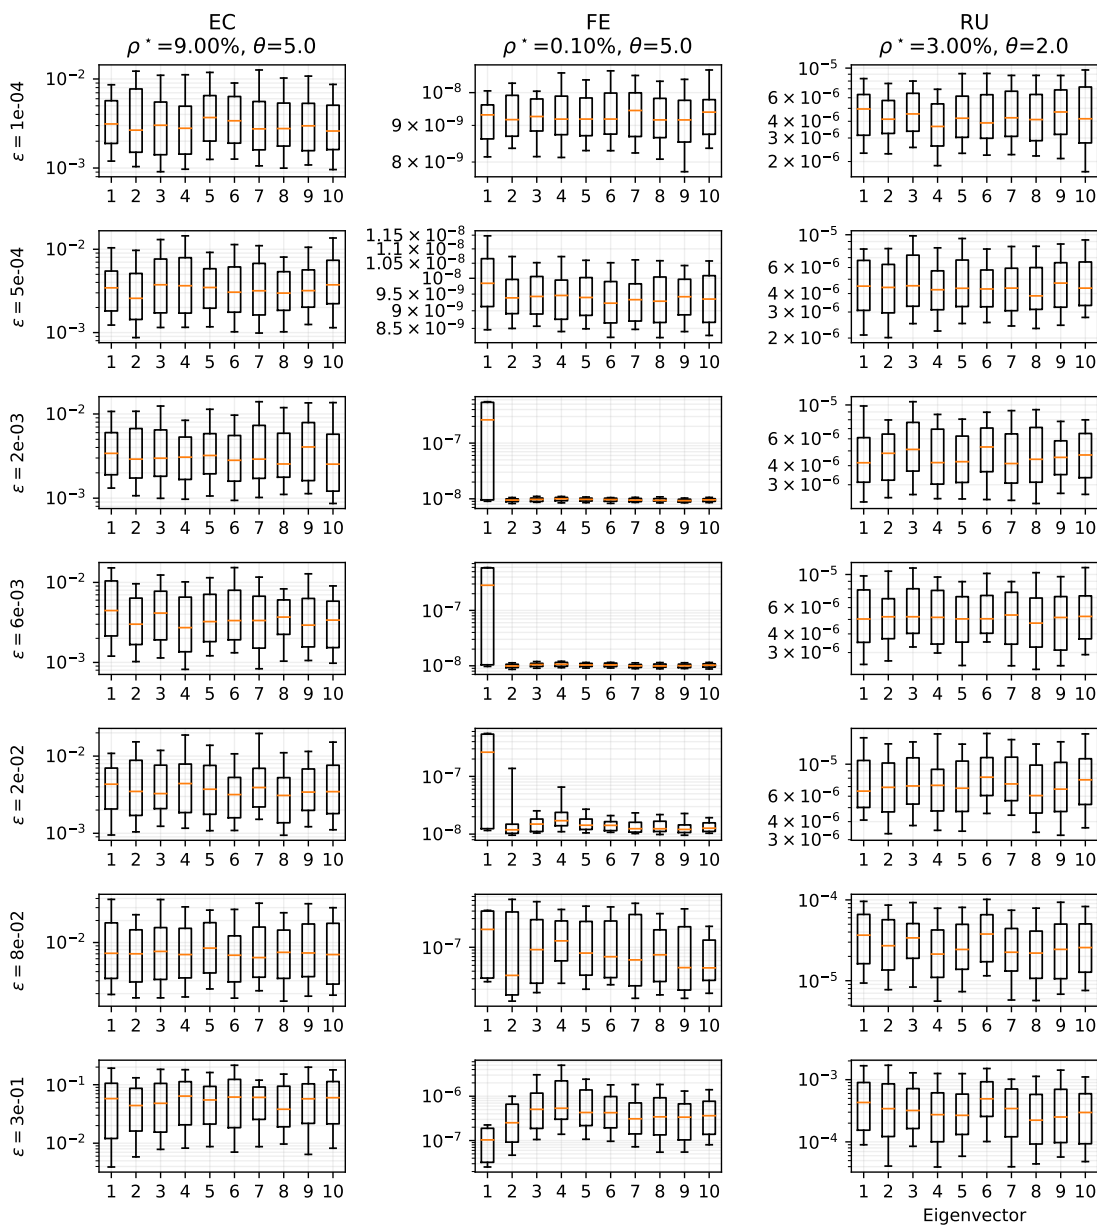


Figure C.1 – Effect of perturbations into eigendirections of the MSE-Hessian using $T = 300$ and $S = 50$ for each of the three parameter-points. Hessians were computed using central difference derivatives with a stepsize of $\epsilon = 10^{-3}$. Plots show the distribution of the mean squared error of the perturbation in $\pm v_i(\Phi)$ (x-axis) over a sample of 50 $S = 1$ Hessians

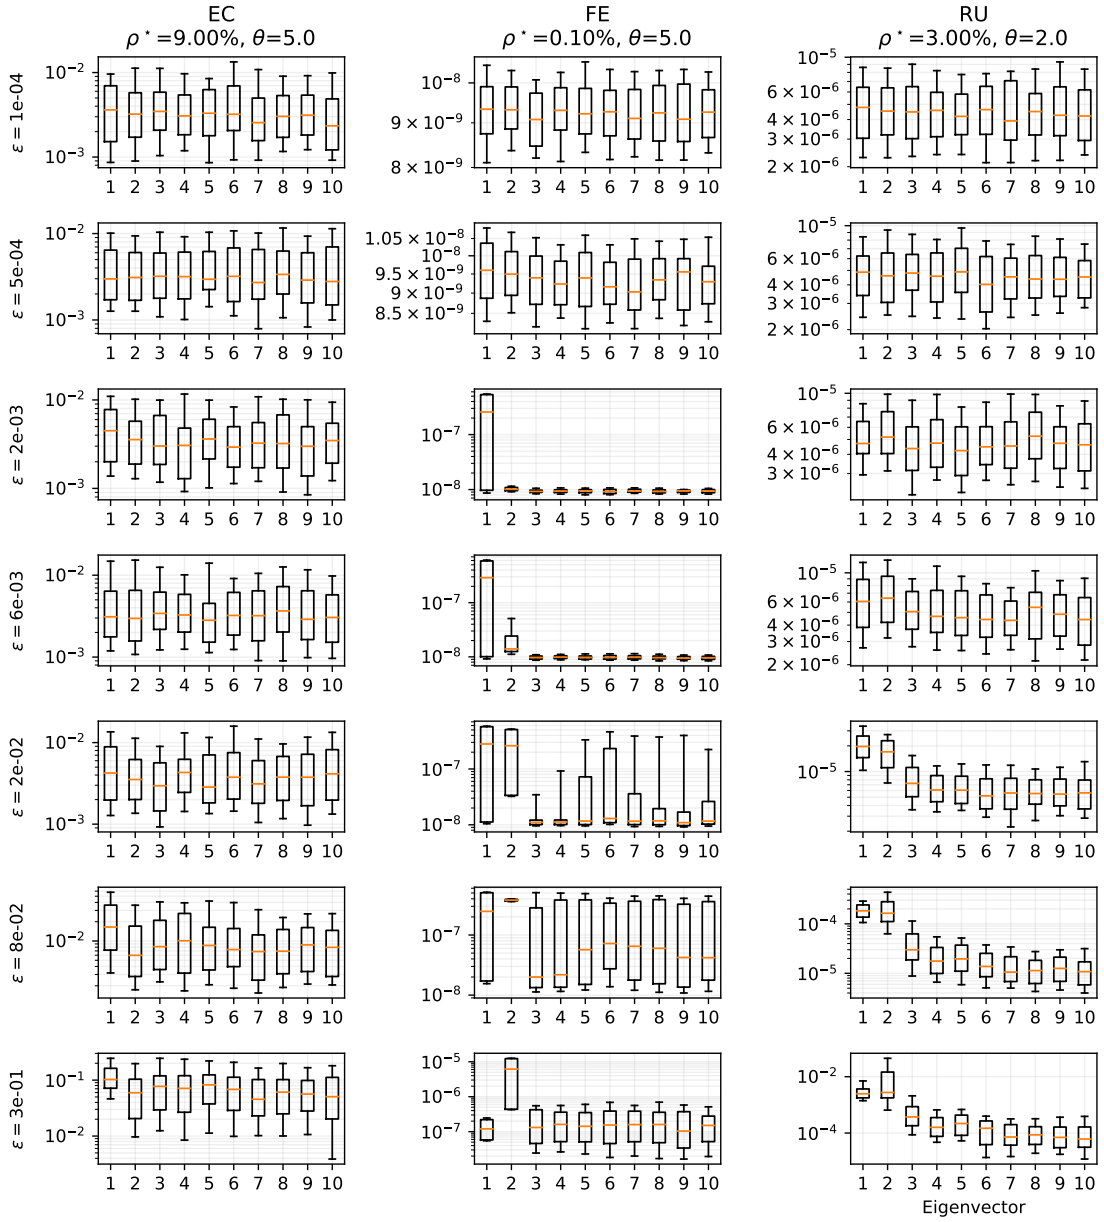


Figure C.2 – Effect of perturbations into eigendirections of the MSE-Hessian using $T = 300$ and $S = 50$ for each of the three parameter-points. Hessians were computed using central difference derivatives with a stepsize of $\epsilon = 10^{-2}$. Plots show the distribution of the mean squared error of the perturbation in $\pm v_i(\Phi)$ (x-axis) over a sample of 50 $S = 1$ Hessians




Figure C.3 – Effect of perturbations into eigendirections of the MSE-Hessian using $T = 300$ and $S = 50$ for each of the three parameter-points. Hessians were computed using central difference derivatives with a stepsize of $\epsilon = 2 \times 10^{-1}$. Plots show the distribution of the mean squared error of the perturbation in $\pm v_i(\Phi)$ (x-axis) over a sample of 50 $S = 1$ Hessians




Figure C.4 – Eigenvalues of a $S = 1$ Hessian as a function of the finite difference step-size ϵ at three different points in the parameter space.

Appendix **D**

Additional Information for the Exploration of the Parameter-space

D.1 Best $pMSE(\mathcal{E}_Q)$ Explorations

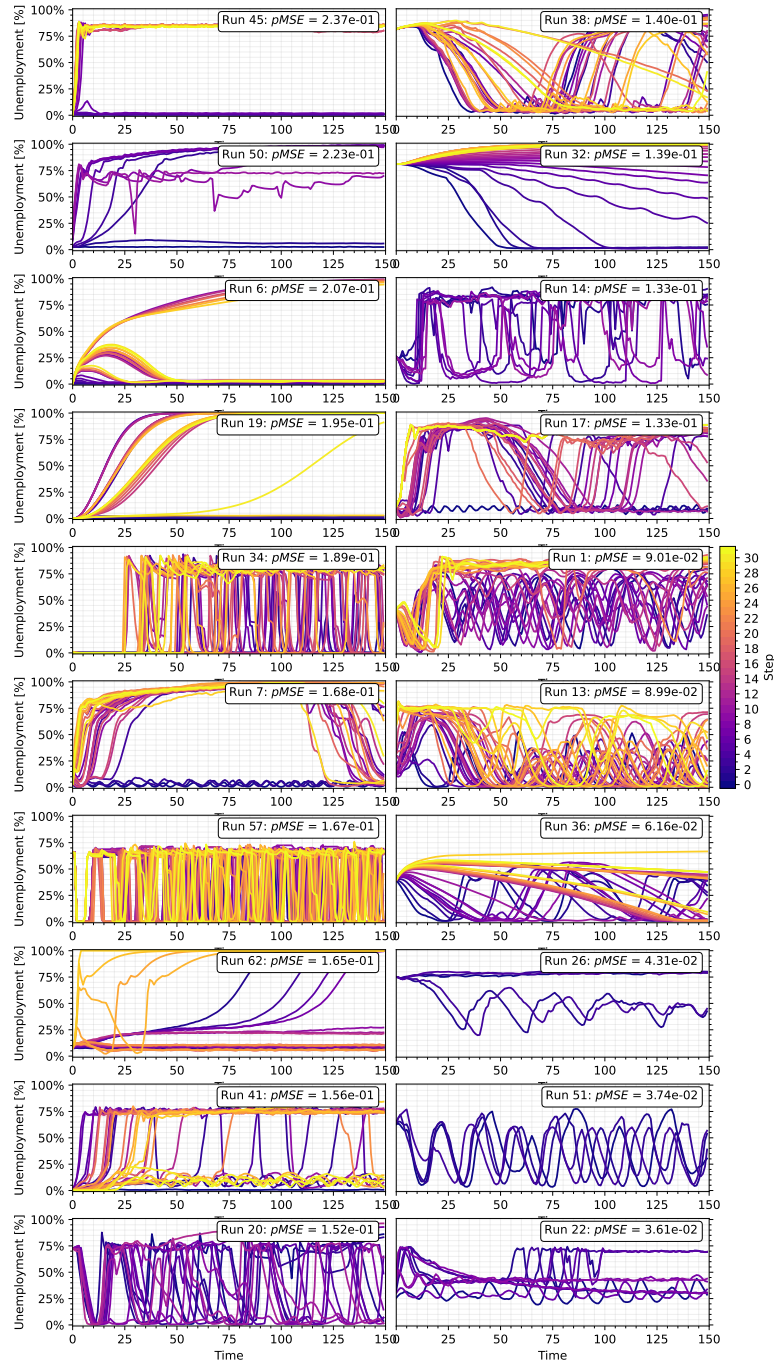


Figure D.1 – Twenty runs with the highest $pMSE(\mathcal{E}_Q)$ for the *SimpleAlgo* approach with $\varepsilon = 10^{-2}$. Initial parameterisations based on Sobol sequence of length 64. Initial conditions are fixed and based on a $T = 1000$ equilibration to the steady state at the initial point. Hessian hyperparameters are $T = 150$ and $S = 1$. Boundaries from Table 5.1

APPENDIX D. ADDITIONAL INFORMATION FOR THE EXPLORATION OF THE PARAMETER-SPACE




Figure D.2 – Twenty runs with the highest $pMSE(\mathcal{E}_Q)$ for the *SimpleAlgoFC* approach with $\varepsilon = 10^{-2}$. Initial parameterisations based on Sobol sequence of length 64. Initial conditions are parameter-point $\Phi(q)$ -specific and based on a $T = 1000$ equilibration to the steady state at that point. Hessian hyperparameters are $T = 150$ and $S = 1$. Boundaries from Table 5.1

D.1. BEST $pMSE(\mathcal{E}_Q)$ EXPLORATIONS

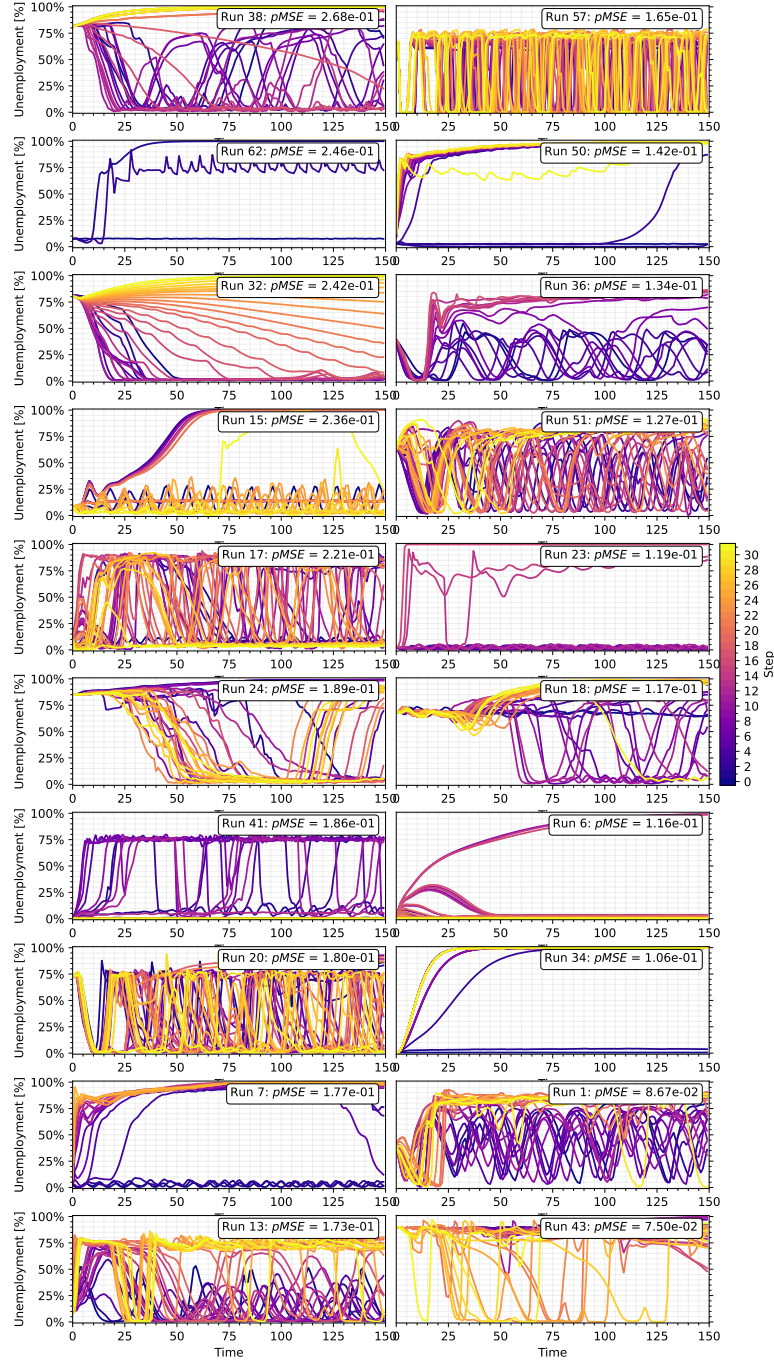


Figure D.3 – Twenty runs with the highest $pMSE(\mathcal{E}_Q)$ for the *ProbeAlgo* approach with $\varepsilon = 10^{-2}$. Initial parameterisations based on Sobol sequence of length 64. Initial conditions are fixed and based on a $T = 1000$ equilibration to the steady state at the initial point. Hessian hyperparameters are $T = 150$ and $S = 1$. Boundaries from Table 5.1

APPENDIX D. ADDITIONAL INFORMATION FOR THE EXPLORATION OF THE PARAMETER-SPACE




Figure D.4 – Twenty runs with the highest $pMSE(\mathcal{E}_Q)$ for the *ProbeAlgoFC* approach with $\varepsilon = 10^{-2}$. Initial parameterisations based on Sobol sequence of length 64. Initial conditions are parameter-point $\Phi(q)$ -specific and based on a $T = 1000$ equilibration to the steady state at that point. Hessian hyperparameters are $T = 150$ and $S = 1$. Boundaries from Table 5.1

D.2 Transition Drivers

APPENDIX D. ADDITIONAL INFORMATION FOR THE EXPLORATION OF THE PARAMETER-SPACE




Figure D.5 – Directions taken by the explorations (vertical) within, into and out of the EC phase in the Mark-0 model, as defined by Galdi et al. (2015). The vertical lines demarcate different exploration runs.

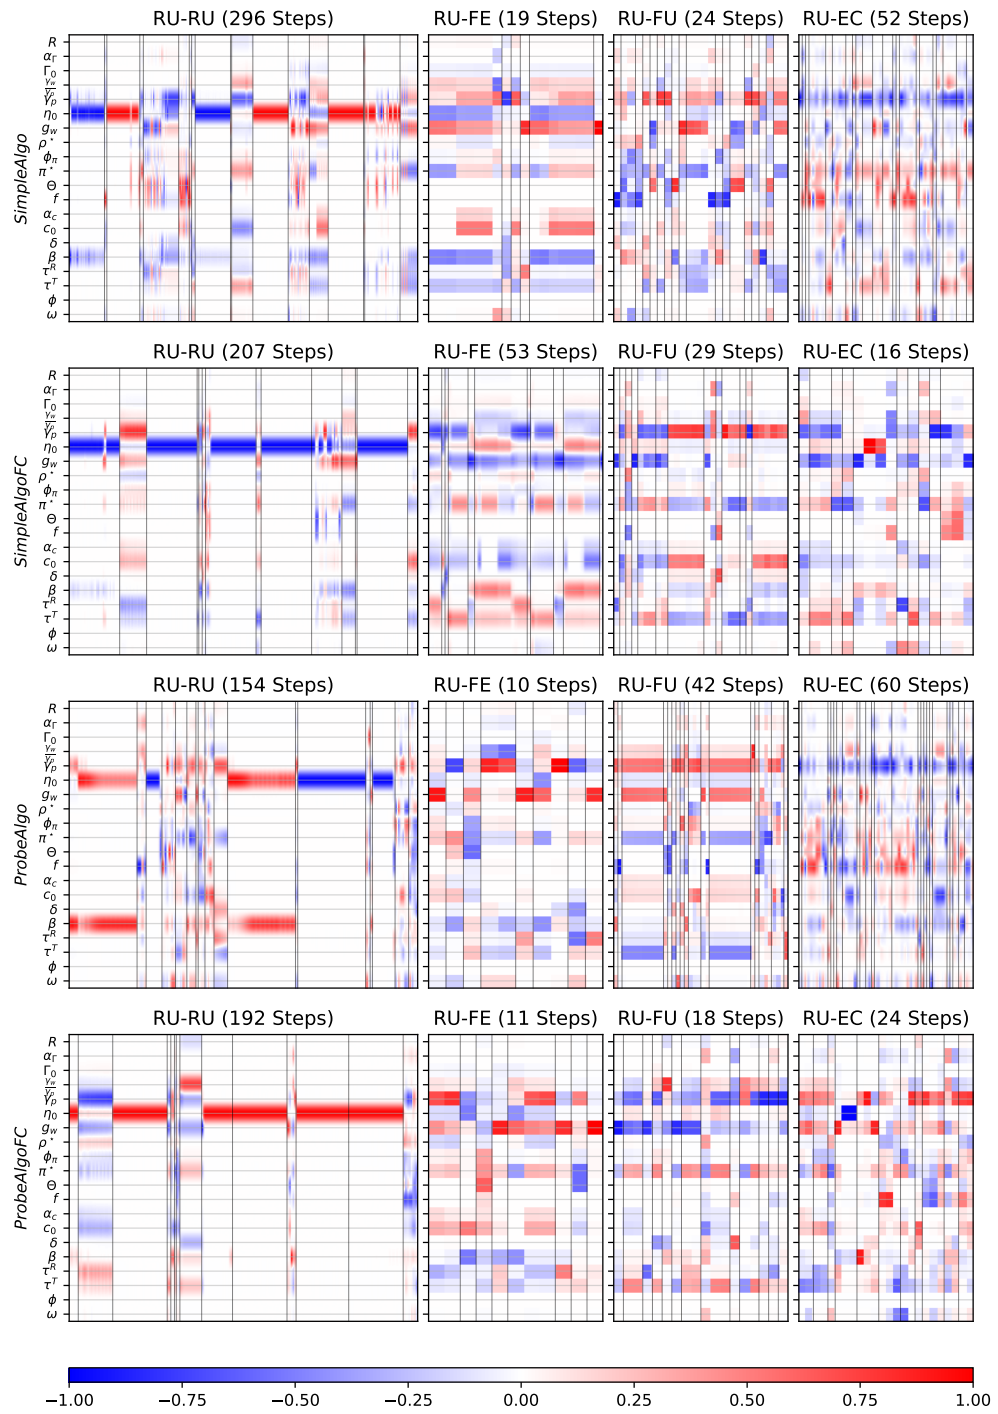


Figure D.6 – Directions taken by the explorations (vertical) within, into and out of the RU phase in the Mark-0 model, as defined by Gualdi et al. (2015). The vertical lines demarkate different exploration runs.

APPENDIX D. ADDITIONAL INFORMATION FOR THE EXPLORATION OF THE PARAMETER-SPACE

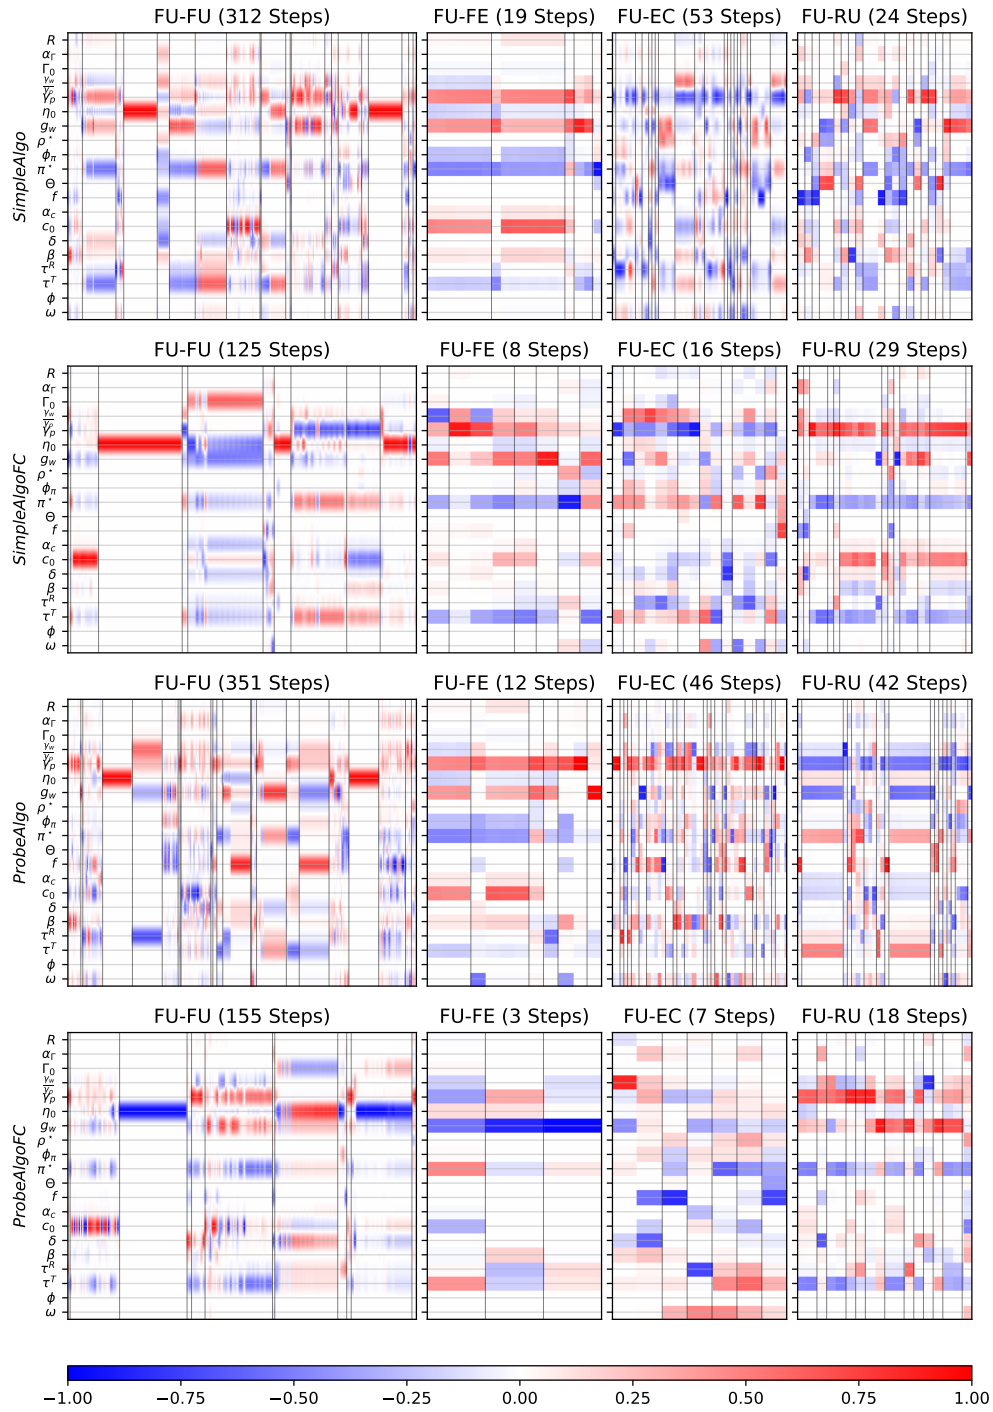


Figure D.7 – Directions taken by the explorations (vertical) within, into and out of the FU phase in the Mark-0 model, as defined by Galdi et al. (2015). The vertical lines demarkate different exploration runs.

Appendix **E**

ABIAM Model Structures

E.1 Macroeconomic Overview

Table E.1 – Overview of the General Structure of four Agent-based Integrated Assessment Models. Details on the energy and resource sector are in Table E.2, the climate module in Table E.3, and policy institutions in Table E.4

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Key References	Lamperti et al. (2018a, 2020)	Lamperti et al. (2019a, 2021)	Safarzynska and van den Bergh (2022)	Czupryna et al. (2020)	Gerdes et al. (2022)
Macroeconomic Framework	K+S Model	Lamperti et al. (2018a)	Safarzynska and van den Bergh (2017b,a)	No prior framework	Influenced by Rengs and Scholz-Wackerle (2019)
General Properties					
Stock-flow Consistent	Yes	Yes	Unspecified	Unspecified	Unspecified
Timestep	Quarter	Quarter	Year	Year	Monthly with annual-only events

Table continues on next page

E.1. MACROECONOMIC OVERVIEW

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Regions	Single	Single	Single	10 regions: Africa, Japan, China, India, Rest of Asia, Europe, North America, Central and South America, Commonwealth of Independent States, Middle East	Two: Global North (with capital goods), Global South (with mines)
Consumption					
Reference	Lamperti et al. (2018a)	Lamperti et al. (2018a)	Safarzynska and van den Bergh (2017a)	Czupryna et al. (2020)	Gerdes et al. (2022)
Goods	Single C-Good	Single C-Good	Multiple goods. Each firm offers good with differentiated quality depending on maximum attainable quality and duration of producing a given good	Multiple goods: agriculture, textiles, chemicals, other manufacturing, transport, and other services	Single C-Good
Production	Constant returns to scale with labour and capital vintages	Constant returns to scale with labour and capital vintages	CES production function with labour, capital and energy	CES with capital, labour and energy	Leontief function with capital and labour
Factor Demand	Proportional to target production	Proportional to target production	Cost minimization of CES given production target	Cost minimization of CES given production target	Proportional to target production

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Supplied Quantity	Expected demand with adjustment for inventory	Expected demand with adjustment for inventory	Weighted average of current sales and actual demand	Tatonnement: increase if excess demand and price above market (Assenza et al., 2015)	Target production depends on expected demand with adjustment for target inventory
Pricing	Variable markup based on market share	Variable markup based on market share	Variable markup on cost based on past market power. Costs include variable cost and fixed cost of producing quality good	Tatonnement: increase if below market and excess demand, decrease if reverse (Assenza et al., 2015)	Tatonnement: Price is adjusted downward if strongly lower sales, and upward if strong excess demand. Bounded from below by production costs. Adoption of a new price is successful with a fixed probability
Investment	Based on expected demand, with scrapping based on expected gains.	Based on expected demand, with scrapping based on expected gains.	Based on expected demand	Based on excess of capital factor demand over existing capital	Target based on expected demand.
Household Demand	All income consumed	All income consumed	Consume based on target wealth-to-permanent-income ratio	Part of income (fixed propensity) plus any savings	Budget is wage income plus proportion of savings

Table continues on next page

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Household Choice	N/A	N/A	Intensity of choice function with utility per product. Product-utility is a Cobb-Douglas type aggregation of quality, price and number of other buyers	Stone-Geary utility function over n goods, with a required minimum consumption. Energy consumption is split into electricity and fuels based on CES, then fuels is split into coal, gas and oil based on CES	Household has shortlist of preferred firms, with 25% change household might change firm to a lower price one. With 25% chance household replaces firm that didn't satisfy demand with random alternative
Market Protocol	Not explicit	Not explicit	Not explicit	Regional clearing except transport and other services. Consumers visit producers (unspecified order) adjusting demands based on prices as clearing proceeds	Two purchasing rounds: a fraction of desired demand bought in round one, the remainder in round two. If a firm cannot satisfy demand, a household tries another firm until satisfied or no supply left.
Capital					
Reference	Lamperti et al. (2018a)	Lamperti et al. (2018a)	Safarzynska and van den Bergh (2017a)	Czupryna et al. (2020)	Gerdes et al. (2022)
Capital Goods	Multiple vintages: labour-productivity, energy-efficiency, emission intensity	Multiple vintages: labour-productivity, energy-efficiency, emission intensity	N/A no explicit sector	Single type of capital	Single type of machine

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Production Technology	Constant returns to scale with labour and energy	Constant returns to scale with labour and energy	N/A	CES with capital, labour and energy	Leontief with capital, labour and resources
Technological Change	Nelson and Winter (1982) imitation-innovation process based on R&D investment (See Dosi et al., 2010)	Nelson and Winter (1982) imitation-innovation process based on R&D investment (See Dosi et al., 2010)	Nelson and Winter (1982) two-step process for improving CES technical coefficients. Maximum product quality and technical coefficients increase exogenously over time	Exogenous noisy growth in CES technical coefficients	None. Fixed technical production coefficients
Prices	fixed markup	fixed markup	N/A since not a sector	Tatonnement: increase if below market average and excess demand (Assenza et al., 2015)	Tatonnement: Price is adjusted downward if strongly lower sales, and upward if strong excess demand. Bounded from below by production costs. Adoption of a new price is successful with a fixed probability
Quantity	Proportional to expected demand	Proportional to expected demand	N/A	Based on expected demand, subtracting own requirements first	Target depends on expected demand with an inventory adjustment

Table continues on next page

E.1. MACROECONOMIC OVERVIEW

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Market Protocol	Choice by price, productivity, energy-efficiency. Can consume from a subset of producers. Delivery at end of period	Choice by price, productivity, energy-efficiency. Can consume from a subset of producers. Delivery at end of period	N/A since not a sector	Assenza et al. (2015): each consumer connected to fixed number of producers, purchasing from lowest producers in order of lowest price first. Excess demand is partially satisfied. Once fixed network is completed, additional demand is fulfilled by excess supply or proportionally allocated. Adjustment for import/export: if a region has an export surplus, imports are preferred	Consumption good firms have shortlist of capital firms, updated every 3 months with 25% change household might change firm to a lower price one. With 25% chance household replaces firm that didn't satisfy demand with random alternative. Firm attempts to buy equal fraction of demand from all sellers, if unsatisfied continues until at least 95% satisfied or no more stock
Labour & Household Income					
Reference	Lamperti et al. (2018a)	Lamperti et al. (2018a)	Safarzynska and van den Bergh (2022)	Czupryna et al. (2020)	Gerdes et al. (2022)
Supply	Fixed	Fixed	Fixed	Fixed Regionally	Fixed Regionally
Differentiation	None	None	None	None	Global south adjusts labour productivity based on local pollution and duration of employment

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Wages	Single market wage based on productivity, price, unemployment	Single market wage based on productivity, price, unemployment	Wages are fixed, with a proportional increase when demand exceeds supply of labour	Unspecified	All firms adjust wage based on change in price of their good over prior 12 months, with a maximum adjustment rate. Downward rigid. Mining employee wages are adjusted for productivity half of firm profits are distributed regionally in proportion to their existing savings. Excess R&D funds are distributed analogously
Other income	None	None	A share of energy of energy and capital rents. Interest on deposits.	Dividends from shareholdings in companies	Regional markets with random matching. Capital and consumption goods have a maximum amount of new hires, and maximum proportion of labour that can be fired. Mines
Market Protocol	Not explicit	Not explicit	Randomly allocated workers	Randomly allocated within regions	
Credit & Financing					
Reference	Lamperti et al. (2018a)	Lamperti et al. (2021) based on Dosi et al. (2015)	Safarzynska and van den Bergh (2017a)	Czupryna et al. (2020)	Gerdes et al. (2022)

Table continues on next page

E.1. MACROECONOMIC OVERVIEW

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Credit Demand	Desired investment net of cash	Desired investment net of cash	Desired investment net of cash	Desired investment	Desired investment net of cash
External Finance	Bank loans	Bank loans, and government bonds for government purchased by banks	Bank loans	Household investment	Bank loans
Firm Bankruptcy	Negative liquid assets or zero market share. Replaced by firm's representing industry averages	Negative liquid assets or zero market share	Inability to pay back loans, high market share firms can extend loan payback by some periods. Capital goes to the bank, which resells to new entrants with fixed probability	Consumer & capital firms go bankrupt if there is a lack of demand or lack of production factors. Fuel extraction companies go bankrupt if all resources are exhausted. Ownership covers the losses	No production for 12 months
Firm Entry	Replace bankrupt firm by industry average firm	Replace bankrupt firm by industry average	new firm enters with fixed probability, offering random product quality with production technology greater than best current technology. Firm demands start-up loan for initial investments, successfully granted with a fixed probability. First period has a zero markup price	Unspecified	Unspecified

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Credit Supply	Pecking-order based on net-worth-to-sales. Upper bound to credit based on debt-to-sales	Bounded by each banks' equity, which is subject to a Basel-II capital rule. Pecking order based on firm's credit worthiness	Firm receives loan if below exogenous debt-to-equity ratio.	Household decide on overall planned investment (income net of consumption)	Unbounded
Interest rate	Markup on central bank rate	Risk premium based on client position in credit ranking (markup is based on quartile within the banks clientele)	The same offer from all banks, with electricity having a lower interest rate	No interest. Firms pay out profits	Fixed rate
Regulation	Maximum credit set by credit multiplier rule	Time-varying capital adequacy ratio	Minimum level of reserves held at central bank, with fixed minimum and higher reserves depending on deposits	None	None
Bank Bankruptcy	None	Equity (net worth) is negative	Equity or reserves are negative. Lending banks write off the loans to defaulted banks	None	None
Bank Entry	None	Bankrupt bank is bailed out by the government. Bailout up to a (fraction of the equity of the smallest incumbent)	None	None	None

Table continues on next page

E.2. ENERGY AND RESOURCE MODULES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Market Protocol	Single bank, pecking-order credit	Each bank has a pecking order process for its list of clients	Loans granted based on debt-to-equity ratio. If a bank has insufficient funds to grant a loan, it asks other banks for loans (order of liquidity) until liquidity requirements are satisfied	Households distribute planned investment among companies based on their current ownership shares and value of companies planned capital increase. If planned investment exceeds demand, money is distributed back to owners	All credit granted if profit rate is larger than interest rate adjusted for credit-lenience

E.2 Energy and Resource Modules

Table E.2 – Overview of the Energy and Resource Module of four Agent-based Integrated Assessment Models

	DSK & DSK-FIN	ABMIAM	CFHS	GRSW
Reference	Lamperti et al. (2018a)	Safarzynska and van den Bergh (2022) based on Safarzynska and van den Bergh (2011); Safarzynska (2012)	Czupryna et al. (2020)	Gerdes et al. (2022)
Energy Markets				

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK & DSK-FIN	ABMIAM	CFHS	GRSW
Primary Energy Types	representative fossil fuel, representative renewable flow	three fuels: (1) increasing unit cost over time, (2) decreasing unit costs over time by Brownian motion, (3) constant price. Associated with gas, coal and renewable energy	Coal, gas, oil, nuclear, hydro, wind and solar	N/A
Producer Heterogeneity	green (renewable) and brown (fossil) power plants with capital vintage dependent cost-structure, thermal efficiency, environmental impact. Production requires only capital, and fossil fuel for brown power plants (one unit per unit electricity)	heterogeneous plants differentiated by: age, productivity, energy source, installed capacity, maximum lifespan, capacity factor	Heterogeneous plants: fuel type, production capacity, storage capacity (solar and wind), lifetime, capacity factor, operation cost factor, electricity transmission loss factor, thermal efficiency factor	N/A

Table continues on next page

E.2. ENERGY AND RESOURCE MODULES

	DSK & DSK-FIN	ABMIAM	CFHS	GRSW
Production Technology	Production requires only capital, and fossil fuel for brown power plants (one unit per unit electricity)	Cobb-Douglas function of capital, labour and fuel with plant-specific TFP inversely related to thermal efficiency increased by Gaussian every period and fixed exponents (substitution factors).	Fixed capacities. Planned electricity fixed for nuclear, beta distribution for all other types (maintenance and weather) with a shape proportional to capacity factors. Solar has double capacity in daytime, and zero at night. Maximum supply augmented by storage and reserve rates for combustion plants	N/A
Production Capacity	Unitary per plant	Plant-specific chosen at creation of plant	Plant-specific	N/A
Production Cost	price of fossil fuels in relation to thermal efficiency for brown plants, zero for green plants	cost of labour, fixed operating costs, fuel costs	operation cost proportional to production, plus labour cost	N/A
Pricing	fixed markup on the marginal producers cost (supplier of last unit)	Inverse demand function with markup	Market clearing in each sub-period, with stress factors if demand ζ maximal provision	N/A
Energy Storage	No	No	Yes, for solar and wind plants	N/A
Final Energy Types	Electricity	Electricity	Electricity (Fuel considered a direct primary purchase)	N/A
Market Structure	Central authority activating plants until demand is met	Cournot game: each plant produces to maximize profits given	Time split into sub-steps (stochastic generation profiles)	N/A

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK & DSK-FIN	ABMIAM	CFHS	GRSW
Depreciation	All plants have a fixed lifetime	Plants have a fixed lifetime	Plants have a fixed lifetime, each period oldest part of capacity for a plant is depreciated.	N/A
Physical Investment	Cost of new brown plant is zero, while new green plant has vintage-dependent fixed cost. Expansion is done when demand exceeds maximum production. Green plants are preferred as long as their fixed cost is less than present value of the cost of the most efficient brown plant	Once a fixed lifetime is reached, a plant exists and the owner invests in a new plant. The type of fuel is chosen based on expected profits from that fuel. A new plant receives a loan to construct, this has to be paid back at the end of its lifetime	Each period a share of old capacity is replaced by investment if maximum expected demand cannot be met post-depreciation. Maximal capacity increases are bounded from above. Solar and wind also invest in storage if electricity prices are volatile.	N/A
Technological Innovation	Fraction of total past sales invested in R&D, allocated in proportion to revenue generated. Nelson and Winter (1982) two-step process: (1) draw for successful innovation, (2) beta-distribution draw or proportional improvement	None - fixed Cobb-Douglas exponents	Overnight investment costs for solar and wind decrease to a specified floor. Substitution of electricity for fuels increases to match intra-fuel substitutability	N/A
Final Demand	C-Good and K-good sectors	C-good market	All sectors except electricity producers	N/A

Table continues on next page

E.2. ENERGY AND RESOURCE MODULES

	DSK & DSK-FIN	ABMIAM	CFHS	GRSW
Market Protocol	Plants activated in order of cost (green first) until demand is met	Plants choose production to maximize profit based on Cournot game with linear inverse demand function (coefficient adapte to guarantee i). Factor demands are based on marginal productivity	In each region, time is in sub-steps with half as night (no solar, low demand) and half as day (higher demand) and cleared each time to develop an average price for the period.	N/A
Resource Markets				
Types of Resources	Representative fossil fuel	Three stylized fuels	Coal, Crude Oil, Natural Gas	Representative extracted resource
Stock and Extraction	Infinite stock	Infinite stock	Finite, regional stocks. Marginal costs of extraction follow Rogner curve. Planned production is adjusted for regional depletion rates (easy to extract first).	Infinite stock
Production Technology	N/A	N/A	Capital gives capacity, fixed total labour. Capital adjusted based on expected vs. realized demand. No innovation.	Leontief dependent on labour and capital

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK & DSK-FIN	ABMIAM	CFHS	GRSW
Price	Exogenous variable	Geometric Brownian Motions	Intersection of supply and vertical demand (short-term inelasticity)	Tatonnement: Price is adjusted downward if strongly lower sales, and upward if strong excess demand. Bounded from below by production costs. Adoption of a new price is successful with a fixed probability
Market Protocol	N/A	N/A	Centrally cleared per fuel (coal, crude, gas). Supplied at the marginal extraction cost, consumers bid quantities and accept market price. If demand _i > planned supply, the highest price is taken and excess demand distributed over suppliers (assuming below maximum extraction). If demand _i > max.extraction _i , proportional rationing.	Capital good firms have shortlist of capital firms, updated every 3 months with 25% change household might change firm to a lower price one. With 25% chance household replaces firm that didn't satisfy demand with random alternative. Firm attempts to buy equal fraction of demand from all sellers, if unsatisfied continues until at least 95% satisfied or no more stock

E.3 Climate Modules

Table E.3 – Overview of the Implementation of Climate Modules in four Agent-based Integrated Assessment Models

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Reference Paper	Lamperti et al. (2020)	Lamperti et al. (2019a)	Safarzynska and van den Bergh (2022)	Czupryna et al. (2020)	Gerdes et al. (2022)
Climate System					
Reference Framework	C-ROADS model of Sterman et al. (2012, 2013). Similar to Nordhaus (1992)	Single-equation framework similar to Matthews et al. (2009, 2012)	Nordhaus (2017)	Petschel-Held et al. (1999)	Single equation framework
Timescale	Annual	Quarterly (like model)	Annual	Annual	Monthly
Pollution Sources	Consumption Goods Sector, Capital Goods Sector, and Energy Sector emit CO2	Consumption Goods Sector, Capital Goods Sector, and Energy Sector emit CO2	Energy sector (fuel-specific) emits CO2	Energy sector (fuel-specific). Unclear Emits CO2	Capital firms in the global North emit CO2. Mines in the global South cause local Pollution

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Description	Two-layer model with two loops: (1) increased "natural" primary CO2 production with CO2 levels, (2) oceans' capacity to uptake carbon falls with CO2 concentration increases. Radiative forcing determines mean temperature increase.	Fixed ratio for change in temperature following change in cumulative emissions	Changes in atmospheric carbon depend on existing stock and past emissions with a permanent and transient part. Stock of carbon affects global mean temperature anomaly $T = \omega \frac{\ln \frac{E_t}{E_{pre}}}{\ln 2}$	Three-equation carbon cycle. Cumulative emissions increase by economy, leading to increase in carbon concentration (cumulative + annual emission - difference to prehistoric values). Temperature change is proportional to carbon intensity vs. prehistoric and temperature vs. prehistoric. Regional temperatures differ, but increase by the same as global mean temperature	Emission increases monthly, and decreases as a fixed proportion of initial concentration. Local pollution of the global south depends on mine production rate and their pollution coefficients
Damage Feedback					
Types of Damage	Capital stocks, labour productivity, inventories, energy-efficiency	Labor productivity, capital stocks	Consumer budgets	Agricultural, Labour, Natural Disasters	Capital stocks, Labor productivity of workers in mines due to local pollution

Table continues on next page

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Damage Distribution	Beta distribution of damage proportion (uni-modal right-skewed), with a mean increasing in temperature, and the right tail increasing with temperature variability (higher risk of extreme events)	Damage follows Nordhaus (2017) deterministic damage function: a quadratic function of temperature levels with two response parameters	Total damage proportional to increase in temperature $(1 + \zeta_1 T_t^{\zeta_2})^{-1}$	Deterministic functions. Agriculture: literature estimates for Europe/Africa/Rest-of-World. Labour: quadratic function of temperature difference to 13°C. Natural disasters: regional regression estimates of USD bn. loss per degree increase in temperature	Damage magnitude is a proportion of total capital level, λ monthly number of disasters (hits) is a nonlinear function of current emissions compared to the baseline, parameterized by an acceleration rate.
Damage Allocation	Each firm is hit with a random draw affecting capital or labour productivity	All firms are hit by the aggregate shock modified by a small Gaussian noise parameter	Based on modification of Dennig et al. (2015): consumption post damage is reduced by consumer-specific share of total damage, dependent on the share of wealth (through wealth-elasticity of damage)	Agricultural damages are reductions in all CES coefficients for Agriculture sector. Labour damages are reductions to labour efficiency of capital/consumer good companies. Natural disasters: reduce all firms' capital, reduce output of consumer and capital good firms	Each hit hits a random firm (capital, consumption or mine) allowing for multiple hits per month. Firms lose capital equivalent to the damage value. Local pollution per mine decreases the mine's employees productivity (from 1 down to a minimum of 1/3). Workers are replaced every 30 years (same employer and savings)

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Tipping Points	Implicit due to higher probability of extreme events	None	None	Non-linearity in data, but no tipping points	None

E.4 Policy Studies

Table E.4 – Overview of the Policy Experiments of four Agent-based Integrated Assessment Models

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Reference	Lamperti et al. (2018a, 2020)	Lamperti et al. (2019a, 2021)	Safarzynska and van den Bergh (2022)	Czupryna et al. (2020)	Gerdes et al. (2022)
Government					
Fiscal Policy (Outflow)	Unemployment benefits as a fixed fraction of current market wage	Unemployment benefits as a fixed fraction of market wage. Bank bailout costs	None	None	Unemployment benefits as 80% of mean regional wage
Budget Deficits	Taxes on firm profits	Taxes on firm incomes and worker incomes. Government pays interest on its bonds to banks and the central bank, which are issued if deficits are positive.	None	None	Taxes on firm profits
Central Bank					
Interest rate	fixed baseline rate	Taylor type rule with target inflation and unemployment rate	Fixed rates	None	Fixed baseline rate

Table continues on next page

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Credit regulation	sets credit multiplier (on deposits) to limit total debt	Bank specific credit multiplier depending on risk weighted assets and bank equity	Minimum reserves and fraction of deposits. Critical Debt-to-equity value for loan-granting	None	None
Policy Experiments					
Fiscal	No Experiments	No experiments	Carbon Tax determined from the social cost of carbon in DICE models, the formula proposed by Rezai and Van der Ploeg (2016). Revenues are distributed as equal lump-sum to citizens	No Experiments	Global civilized market institution: sets base fine for emissions (capital firms) and local pollution (mines). Subsidises via: (a) no subsidy, (b) fines collected are redistributed in region of collection based on firms' fraction of total capacity, (c) Same as b, but a share of global north funds is reallocated to the south, (d) same as b but funds are multiplied by a government grant scheme
Monetary	No Experiments	No Experiments	No Experiments	No Experiments	No Experiments

Table continues on next page

APPENDIX E. ABIAM MODEL STRUCTURES

	DSK	DSK-FIN	ABMIAM	CFHS	GRSW
Regulation	No Experiments	(1) Carbon-risk adjustment: firms rank in banks' pecking order becomes the average of their credit rank and their emissions rank. (2) Green public guarantees: government backs loans to green firms completely. (3) Green Basel-II: exclude loans to green firms from the credit multiplier regulation, thus increasing total credit supply	No Experiments	No Experiments	No Experiments
Other	Exogenous increases to fossil fuel prices mimicking a carbon tax / green investment subsidy, and decreases mimicking current fossil subsidies (Done in Lamperti et al., 2020)	No experiments	Test carbon tax when probability of a electricity plant being renewable is 50% exogenously (e.g. regulation driven)	No Experiments	No Experiments

Appendix **F**

The Dynamic Solow Model

F.1 Parameterisation of the Dynamic Solow Model

In this appendix, we examine the model's parameters and discuss how they affect the behavior of the dynamical system (7.24) in the phase space.

We begin with equation (7.24b) that describes sentiment dynamics. Parameter β_1 defines the relative importance of the herding and random behaviors of firms. In an unforced situation ($\beta_2 = 0$), the number of stable equilibrium points, to which the firms' sentiment s converges, doubles at $\beta_1 = 1$ from one to two. For $\beta_1 < 1$, random behavior prevails since there is a single equilibrium at $s = 0$, meaning firms fail to reach a consensus opinion. Conversely, for $\beta_1 > 1$, herding behavior rules as equation (7.24b) generates a polarized, bi-stable environment with one pessimistic ($s < 0$) and one optimistic ($s > 0$) equilibrium states. It is sensible to assume $\beta_1 \sim 1$, otherwise firms would unrealistically behave either randomly or in perfect synchronicity. We set $\beta_1 = 1.1$, implying a slight prevalence of herding over randomness. In addition, we set $\beta_2 = 1$ to ensure that analysts' influence on firms' managers likewise appears in the leading order.

We now consider the information dynamics in (7.24c). The terms under the hyperbolic tangent describe the impacts of economic growth and exogenous news on the collective opinion of analysts h . We assume these two sources of information are of equal importance. Thus, we expect that $\gamma\omega_y = O(1)$ in the feedback term and we model ξ_t as an Ornstein-Uhlenbeck process with an $O(1)$ standard deviation and short decorrelation timescale τ_ξ . Note that $\omega_y \ll 1$ and accordingly $\gamma \gg 1$.

Finally, we inspect the economic dynamics in (7.24a). In this equation, different terms determine leading behaviors on separate timescales. We show in F.4 that the last three terms (with technology growth rate ε estimated on the basis of observed total factor productivity) are in balance in the long run. However, if we consider short timescales, the change in sentiment \dot{s} becomes dominant. Thus, equation (7.24a) can be approximated

in the short run as $\dot{z} \sim \rho c_1 \dot{s}$ and we set $\rho c_1 = 1$. We also note that by construction $c_2 \ll c_1$ to ensure that the term $c_2 s$ does not contribute to capital demand dynamics on short timescales. Hence we expect $c_2 \ll 1$.

As highlighted in Section 7.1, there is a segregation of characteristic timescales that emerges naturally from the types of decisions faced by the different agents in the model: $\tau_\xi \ll \tau_h \ll \tau_s \ll \tau_y \ll 1/\varepsilon$. This segregation facilitates the transfer of the impact of instantaneous news shocks ξ_t across multiple timescales. The estimates for the timescales are discussed in Section 7.1.

The parameters c_2 and γ are central to the system's behavior in the phase space. Increasing c_2 stabilizes the system, strengthening convergence towards the stable equilibria and creating a higher barrier between attracting regions. The role of γ is twofold. As γ grows from zero, its immediate effect is to destabilize the system due to growing economic feedback. However, as γ continues to increase, it exerts a stabilizing effect similar to that of c_2 because of the term γc_2 in the equilibrium condition:

$$\operatorname{arctanh}(s) - \beta_1 s = \beta_2 \tanh(\gamma c_2 s + \gamma \varepsilon), \quad (\text{F.1})$$

which follows from equations (7.24) for $\dot{h} = \dot{s} = \dot{z} = \xi_t = 0$. Consequently, the potential to generate autonomous economic instability is limited. In particular, there exists a critical value¹ of $c_2 \sim 10^{-4}$ below which feedback may generate a limit cycle and above which it does not. Figure F.1 depicts the formation and subsequent destruction, for $c_2 = 10^{-4}$, of the limit cycle as γ increases.

In this paper, we argue that realistic economic behaviors cannot be explained by a stochastic limit cycle. Therefore, we proceed to study the system for $c_2 \gtrsim 10^{-4}$, which ensures a bi-stable configuration without a limit cycle. Figure F.2 illustrates that as c_2 increases, the barrier between attracting regions grows stronger, resulting in less frequent crossings from one region to the other (i.e. cycle duration increases). We seek c_2 at the lower end of this range to reduce cycle duration.

Similarly, the barrier between attracting regions grows stronger as γ increases, resulting in infrequent transitions between the attractors. Relatively small values of γ , however, dampen feedback, leading to weak dynamics and stochastic-like behavior. Accordingly, we focus on values of γ between these two extremes. Figure F.3 depicts the dynamics under different values of γ , with balanced dynamic behaviors for a reasonably wide range thereof.

We select $c_2 = 7 \times 10^{-4}$ and $\gamma = 2000$ for the base cases studied in Section 7.2. Note that $c_2 \ll 1$, $\gamma \gg 1$ and $\gamma \omega_y = O(1)$, as expected. Figure F.4 provides the base case phase portrait ($\xi_t = 0$) projected on the (s, h) -plane, showing attracting regions around the two stable equilibria as well as long trajectories passing near each attractor and ending at the opposite equilibrium. These trajectories, which emerge due to strong feedback ($\gamma \gg 1$), allow the economy to transition quickly between expansions and contractions.

¹Subject to the values set for the other parameters.




Figure F.1 – Development of a stable limit cycle with increasing γ for $\xi_t = 0$, $c_2 = 1 \times 10^{-4}$ and all other parameter values from the base case (F.2). The left panels show the phase portraits projected on the (s, z) -plane and the right panels plot $s(t)$. Classification of equilibrium points is provided in footnote 19 in Section 7.2. As γ increases, a large stable limit cycle emerges and then vanishes, demonstrating the destabilizing effect of γ at low values and its stabilizing effect at high values. (i) Stable dynamics for $\gamma = 350$: red trajectories converge to the left focus and blue trajectories converge to the right focus. (ii) The equilibria become unstable and a large stable limit cycle emerges for $\gamma = 1000$. (iii) The left node and the saddle point vanish while the limit cycle persists at $\gamma = 4000$. (iv) The dynamics are again stable at $\gamma = 15000$: trajectories converge to the stable focus and the limit cycle disappears.




Figure F.2 – The effect of c_2 on the dynamics of sentiment $s(t)$ for $\xi_t \neq 0$. As c_2 increases from the base case value $c_2 = 7 \times 10^{-4}$ (left) to $c_2 = 9.5 \times 10^{-4}$ (right), the barrier separating the two attracting regions grows stronger. The system spends more time captive to the attractors, reducing the frequency of the crossings between them and lengthening the duration of fluctuations. Note that the system tends to stay longer at the expansion attractor (where $s > 0$) owing to the asymmetry induced by technological growth $\varepsilon > 0$. All other parameters are from the base case (Table F.1).




Figure F.3 – The effect of γ on the dynamics of sentiment $s(t)$ for $\xi_t \neq 0$. At low γ , the system's behavior is dominated by noise as the barrier between the two attracting regions is weak. As γ increases, the barrier grows stronger and the system becomes extremely bi-stable. Reasonably balanced dynamics emerge in the range from $\gamma = 1500$ to $\gamma = 2500$. Note the asymmetry caused by technological growth becomes exacerbated as γ increases in accordance with equation (F.1). All other parameters are from the base case (Table F.1).




Figure F.4 – System dynamics in the base case (F.2). Left: A phase portrait ($\xi_t = 0$) projected on the (s, h) -plane. The portrait depicts stable foci, separated by a saddle point, and the large trajectories relevant for regime transitions. Right: A trajectory ($\xi_t \neq 0$) projected on the (s, h) -plane. The stable foci are at the center of the two attracting regions, within which the trajectory is dense. The transit of the economy between these regions corresponds to regime transitions between contractions and expansions, occurring at much shorter intervals than the periods during which the economy is captive to an attractor. The trajectory was smoothed by a Fourier filter to remove harmonics with periods less than 500 business days for clean visualization.




Figure F.5 – The effect of ϵ on the dynamics of sentiment $s(t)$ for $\xi_t \neq 0$. As ϵ increases from the base case value $\epsilon = 2.5 \times 10^{-5}$ (left) to $\epsilon = 7.5 \times 10^{-5}$ (right), the system behavior begins to exhibit a stronger asymmetry between the contraction and expansion attractors. All other parameters are from the base case (Table F.1).

The attractors are not connected and the economy cannot cross the boundary separating them in the absence of exogenous news shocks ξ_t . It takes a random news event to force the economy, entrapped by one attractor, across the boundary. Once it crosses the boundary, the economy finds itself on the long trajectory that takes it swiftly to the other attractor, where the economy remains captive until another news event instigates the next regime transition by again forcing the economy across the boundary. At this point the economy is carried back to the entrapment region where it started. This is the coherence resonance mechanism that is at the heart of the economic fluctuations in our model.

As a final comment, we note that if $\varepsilon = 0$, the equilibrium condition (F.1) is symmetric to $s \rightarrow -s$. Technology growth, $\varepsilon > 0$, causes an asymmetry² wherein the equilibrium at $s > 0$ becomes stronger than that at $s < 0$ (to the extent that the latter vanishes above a certain threshold). As a result, the system tends to stay longer in the region where economic sentiment is positive, accelerating the economy's long-term growth. The asymmetry, however, vanishes in the limit cycle regime, whether periodic or stochastic (Section 7.3). Figure F.5 illustrates this asymmetric behavior.

F.2 Model Variables and Parameters

Table F.1 outlines the notation used in the Dynamic Solow Model, highlighting the set of variables (V), constants (C) and noise (N), as well as the respective constants' baseline values.

²Note that γ amplifies the asymmetry via the term $\gamma\varepsilon$ in equation (F.1).

Table F.1 – Parameters and Notation of the Dynamic Solow Model

Definition	Type	Base Case
Production (Section 7.1)		
Y Production	V	-
K Capital in production	V	-
y Log production	V	-
k Log capital in production	V	-
τ_Y Production characteristic timescale (business days)	C	1000
A_0 Initial technology level	C	1
ε Daily technology growth rate	C	2.5×10^{-5}
ρ Capital share in production	C	1/3
Capital Supply (Section 7.1)		
K_s Capital supply	V	-
k_s Log capital supply	V	-
λ Proportion of income saved	C	0.15
δ Daily depreciation rate	C	2×10^{-4}
Capital Demand (Section 7.1)		
K_d Capital demand	V	-
k_d Log capital demand	V	-
z Sentiment level	V	-
h Information level	V	-
c_1 Capital demand sensitivity to \dot{s}	C	3
c_2 Capital demand sensitivity to s	C	7×10^{-4}
β_1 Sentiment herding factor	C	1.1
β_2 Sentiment sensitivity to h	C	1.0
γ Feedback strength factor	C	2000
ξ_t News shocks	E	-
τ_s Sentiment characteristic timescale (business days)	C	250
τ_h Information characteristic timescale (business days)	C	25
τ_ξ News shocks characteristic timescale (business days)	C	5
Limiting Cases (Section 7.2)		
z Production growth indicator	V	-
ω_y $\omega_y = \tau_y^{-1}$	C	0.001
General Case (Section 7.3)		
y_0 Asymptotic growth rates of y	V	-
k_{s0} Asymptotic growth rates of d_s	V	-
k_{d0} Asymptotic growth rates of k_d	V	-
R Classic Solow growth rate $R = \varepsilon/(1 - \rho)$ (daily)	C	3.75×10^{-5}
$1/\varepsilon$ Technology growth timescale (business days)	C	4×10^4

F.3 Approximate Solution to the Supply-Driven Regime

In this appendix, we solve equation (7.22) approximately through use of the boundary layer technique and obtain the economy's path in analytic form in the intermediate and long run under the supply-driven regime ($K = K_s$).

The starting point of our derivation is equation (7.22), for convenience repeated here:

$$\tau_Y \ddot{K} + (1 + \tau_Y \delta) \dot{K} + \delta K = \lambda K^\rho e^{\varepsilon t}. \quad (\text{F.2})$$

Recall that $1 \ll \tau_Y \ll 1/\varepsilon$, where τ_Y is the timescale in which output adjusts to changes in the level of capital and $1/\varepsilon$ is the timescale of output growth in the long run. We aim to capture the dynamics on these two timescales by solving equation (F.2) on the interval $t \geq O(\tau_Y)$. For simplicity, we assume that $\tau_Y \delta \gg 1$, which implies that $\tau_Y \delta \dot{K}$ is much larger than \dot{K} and $\tau_Y \ddot{K}$ on the interval $t \geq O(\tau_Y)$, allowing us to derive a more compact solution.

First, we consider equation (F.2) for $t \gg \tau_Y$. In this outer region, $\tau_Y \dot{K} \ll K$ and we can approximate the solution to (F.2) by the solution to equation:

$$\delta K_o = \lambda K_o^\rho e^{\varepsilon t}, \quad (\text{F.3})$$

which is given by

$$K_o = \left(\frac{\lambda}{\delta} e^{\varepsilon t} \right)^{\frac{1}{1-\rho}}. \quad (\text{F.4})$$

Next, we consider equation (F.2) on the interval $O(\tau_Y) \leq t \ll 1/\varepsilon$, where $e^{\varepsilon t} \rightarrow 1$ and $\tau_Y \dot{K}$ is not necessarily substantially smaller than K . In this inner region, we can approximate the solution to (F.2) by the solution to

$$\tau_Y \dot{K}_i + K_i = \frac{\lambda}{\delta} K_i^\rho. \quad (\text{F.5})$$

This is the Bernoulli equation and its solution is given by

$$K_i = \left(B e^{-(1-\rho)\frac{t}{\tau_Y}} + \frac{\lambda}{\delta} \right)^{\frac{1}{1-\rho}}, \quad (\text{F.6})$$

where B is the constant of integration.

Solutions K_o and K_i must match in the overlapping interval $\tau_Y \ll t \ll 1/\varepsilon$. This is satisfied for any value of B since $K_o \rightarrow \left(\frac{\lambda}{\delta}\right)^{\frac{1}{1-\rho}}$ and $K_i \rightarrow \left(\frac{\lambda}{\delta}\right)^{\frac{1}{1-\rho}}$, as follows from (F.4) and (F.6). Thus, we approximate the solution to equation (F.2) in this region by

$$K_m = \left(\frac{\lambda}{\delta} \right)^{\frac{1}{1-\rho}}. \quad (\text{F.7})$$

The approximate solution to equation (F.2) that is uniformly valid for all $t \geq O(\tau_y)$ is given by

$$K = K_i + K_o - K_m = \left(\frac{\lambda}{\delta}\right)^{\frac{1}{1-\rho}} \left(\left(B e^{-\left(\frac{1-\rho}{\tau_y}\right)t} + 1 \right)^{\frac{1}{1-\rho}} + e^{\left(\frac{\varepsilon}{1-\rho}\right)t} - 1 \right), \quad (\text{F.8})$$

where B has been rescaled for convenience.

As a final step, we obtain the solution for output Y by inverting the equation of capital motion (7.4):

$$Y = \frac{1}{\lambda} (\dot{K} + \delta K). \quad (\text{F.9})$$

Note that $\dot{K} \ll \delta K$ on the interval $t \geq O(\tau_Y)$ due to the simplifying assumption $\tau_Y \delta \gg 1$. Therefore, the corresponding uniform approximation for output Y , valid for all $t \geq O(\tau_Y)$, is given by

$$Y = \left(\frac{\lambda}{\delta}\right)^{\frac{\rho}{1-\rho}} \left(\left(B e^{-\left(\frac{1-\rho}{\tau_y}\right)t} + 1 \right)^{\frac{1}{1-\rho}} + e^{\left(\frac{\varepsilon}{1-\rho}\right)t} - 1 \right). \quad (\text{F.10})$$

F.4 Asymptotic Analysis of Long-Term Growth

In this appendix, we study the behavior of the Dynamic Solow model in the long run by seeking $y \sim y_0 t$, $k_d \sim k_{d0} t$ and $k_s \sim k_{s0} t$ in equations (7.27)-(7.33) at large values of t .

Asymptotic Behavior in the Supply-Driven Regime ($k_d > k_s$)

We first consider the situation where capital demand exceeds supply, which entails $k = k_s$ under the market clearing condition (7.32), and obtain the resulting growth rates.

For $t \gg 1$, the production equation (7.27) becomes

$$e^{(\rho k_{s0} + \varepsilon - y_0)t} - 1 = \tau_y y_0. \quad (\text{F.11})$$

Consequently, $(\rho k_{s0} + \varepsilon - y_0)t$ must be constant, which in turn implies that

$$y_0 = \rho k_{s0} + \varepsilon, \quad (\text{F.12})$$

with a precision of up to $O(1/t)$. Similarly, capital supply equation (7.28) yields

$$k_{s0} = \lambda e^{(y_0 - k_{s0})t} - \delta, \quad (\text{F.13})$$

so that $(y_0 - k_{s0})t$ is constant and, therefore, with a precision of up to $O(1/t)$:

$$k_{s0} = y_0. \quad (\text{F.14})$$

It follows from equations F.12 and F.14 that

$$y_0 = k_{s0} = \frac{\varepsilon}{1 - \rho} \equiv R, \quad (\text{F.15})$$

where R denotes the classic Solow growth rate.³

To determine the growth rate of capital demand k_{d0} , we average equation (7.29) with respect to time, noting that $\bar{s} = 0$ since s is bounded:

$$k_{d0} = c_2 \bar{s}, \quad (\text{F.16})$$

where the bar denotes the time average.

Then we average equation (7.31) while noting that $\bar{h} = 0$ since h is bounded and that $H(k_s, k_d) = 0$ from (7.26) (no feedback) to obtain

$$\bar{h} = \overline{\tanh(\xi_t)} = \tanh(\overline{\xi_t}) = 0, \quad (\text{F.17})$$

where we have assumed that fluctuations are small to allow us to take averages under the hyperbolic tangent⁴. Similarly averaging equation (7.30) leads to

$$\bar{s} = \overline{\tanh(\beta_1 s + \beta_2 h)} = \tanh(\beta_1 \bar{s}). \quad (\text{F.18})$$

Equation (F.18) has three solutions for $\beta_1 > 1$: $s = 0$, $s_- < 0$, and $s_+ > 0$, where $s_- = -s_+$. Our focus is on s_- and s_+ as they correspond to the stable equilibrium points⁵. The system spends most of its time in the attracting regions that surround each of these two equilibria and transits rapidly between them when forced by exogenous noise. In the long run, the time spent in transit is negligible relative to the length of time during which the system is entrapped by the attractors. The attractors have the same strength and are located symmetrically in s , thus the system tends to spend an equal amount of time at each of them at large t . Therefore, its average position with respect to sentiment s must be zero. More formally, taking s_- and s_+ as the attractors' proxies, we estimate the long-term average sentiment as

$$\bar{s} = \frac{1}{2} (s_- + s_+) = 0. \quad (\text{F.19})$$

Hence equation (F.16) yields

$$k_{d0} = 0. \quad (\text{F.20})$$

This result is intuitively clear: the growth of demand is driven in the long run by average sentiment, which converges to zero because its dynamics are symmetric in the absence of feedback. We conclude that in the supply-driven regime the economy's growth is, as expected, independent of capital demand and matches the classic Solow growth, $y_0 = k_{s0} = R$, while capital demand is stagnating ($k_{d0} = 0$). We verify these results via numerical simulations in Section 7.3.

³This same result also follows from equation (7.23) for $t \geq O(1/\varepsilon)$.

⁴This simplifying assumption does not severely restrict applicability as $\tanh x$ is approximated reasonably well by a linear function for $-1 \leq x \leq 1$ and the noise amplitude is $O(1)$ in (F.17).

⁵For the base case value $\beta_1 = 1.1$, we have $s_{\pm} \approx \pm 0.5$ from (F.18).

Asymptotic Behavior in the Demand-Driven Regime ($k_d < k_s$)

In the demand-driven regime, the market clearing condition (7.32) yields $k = k_d$, so that equation (7.27) becomes

$$e^{(\rho k_{d0} + \varepsilon - y_0)t} - 1 = \tau_y y_0. \quad (\text{F.21})$$

Consequently,

$$y_0 = \rho k_{d0} + \varepsilon, \quad (\text{F.22})$$

with a precision of up to $O(1/t)$. Similarly, equation (7.28) takes the form:

$$k_{s0} = \lambda e^{(y_0 - k_{s0})t} - \delta e^{k_d - k_s}. \quad (\text{F.23})$$

The term $\delta e^{k_d - k_s}$ can be neglected as it is exponentially small for $k_d < k_s$; therefore, with a precision of up to $O(1/t)$:

$$y_0 = k_{s0}. \quad (\text{F.24})$$

And finally, averaging equation (7.29) leads to

$$k_{d0} = c_2 \bar{s}. \quad (\text{F.25})$$

We can rewrite equations (F.22) and (F.24) as

$$y_0 = k_{s0} = R + \rho(k_{d0} - R). \quad (\text{F.26})$$

It follows that if $\bar{s} > \frac{R}{c_2}$, then the economy's long-term growth exceeds the classic Solow growth rate R . For the base case values of c_2 , ε and ρ in our model, we find $\bar{s} > 0.05$.

To estimate \bar{s} , we must consider three types of characteristic behavior possible in the demand-driven regime: noise-driven, limit cycle and coherence resonance behavior. Noise-driven behavior prevails when feedback is weak. This situation is, in its limit, equivalent to that of the supply-driven regime in which sentiment behaves symmetrically with respect to the origin. Therefore, $\bar{s} \rightarrow 0$. Thus, the noise-driven mode generates growth $y_0 \rightarrow \varepsilon$, which is lower than R .

The growth in the two other modes is studied numerically in Section 7.3. For completeness, we briefly note, first, limit cycles (periodic or stochastic) lead to $\bar{s} \rightarrow 0$ and $y_0 \rightarrow \varepsilon$ (as the economy tends to spend a half of its time in the region where $s > 0$ and the other half where $s < 0$) and, second, coherence resonance yields $\bar{s} > 0.05$ and $y_0 > R$, owing to the attractors' asymmetry caused by technological growth ($\varepsilon > 0$) in the presence of economic feedback ($\gamma > 0$).

As a final remark, it follows from (F.22) that, asymptotically, $z \sim z_0 t \sim (\rho k_{d0} + \varepsilon - y_0)t \sim O(1)$. The system's motion is therefore bounded in z . Its motion is likewise bounded in s and h , which vary between -1 and 1, as, at the boundaries, \dot{s} and \dot{h} are directed into the domain of motion as follows, respectively, from equations (7.30) and (7.31). Thus, the system's phase trajectories are bounded in the (s, h, z) -space.

Appendix G

The Self-Reflexive DSGE Model

G.1 Table of Parameters

Description	Definition	Value(s)
Variables		
c_t	Real consumption	Eq. (8.1)
n_t	Labour	Eq. (8.1)
G_t	Propensity to consume from Income	Eq. (8.31) [G_{\min}, G_{\max}]
\mathcal{I}_t	Income in period t	Eq. (8.2)
w_t	Real wage	Eq. (8.12)
b_t	Real income from one-period bond purchased in t	Eq. (8.7)
k_t	Real capital	Eq. (8.6)
q_t	Realised yield on capital	Eq. (8.6) [$0, q_t^*$]
i_t	Real value of total investment	Eq. (8.5) [$(1 - G_t)\mathcal{I}_t$]
F_t	Allocation of investment to capital	Eq. (8.36) [F_{\min}, F_{\max}]
z_t	Total factor productivity (TFP)	Eq. (8.9) [$0, \infty$]
q_t^*	Ideal capital returns	Eq. (8.13)
ξ	Investment risk process	Eq. (8.27) [$0, 1$]
μ_t^q	Estimated expected return	Eq. (8.32)
σ_t^q	Estimated investment volatility	Eq. (8.33)
\mathcal{S}_t	Estimated Sharpe ratio	Eq. (8.34)
\mathcal{C}_t	Consumption confidence	Eq. (8.30) [$-1, 1$]
Σ_t	Unbounded sentiment	Eq. (8.35)
Studied Parameters		
δ	Depreciation rate	Eq. (8.6) 0.001, 0.005, 0.02
a	Investment risk multiplier	Eq. (8.28) 15
c_0	Consumption confidence threshold	Eq. (8.30) [$0, 0.025$]
λ	Memory kernel for the Sharpe ratio	Eq. (8.32) 0.607, 0.905, 0.951, 0.98
ν	Interpolation between \mathcal{S}_t and \mathcal{C}_t	Eq. (8.35) 0.75, 1.0

Fixed Parameters

γ	Disutility of labour	Eq. (8.1)	1
π_t	Inflation rate	Eq. (8.2)	0.1%
r_t	Real interest rate	Eq. (8.2)	0.15%
p_t	Price level	Eq. (8.2)	1
α	Capital share in production	Eq. (8.8)	1/3
$1/(1 + \rho)$	Elasticity of substitution k vs. n	Eq. (8.8)	-7
z_0	Baseline value of the total factor productivity	Eq. (8.9)	0.05
η	Autocorrelation of total factor productivity	Eq. (8.9)	0.5
θ_c	Consumption rate transition width	Eq. (8.30)	300
G_{\min}	Minimum consumption rate	Eq. (8.31)	0.05
G_{\max}	Maximum consumption rate	Eq. (8.31)	0.95
\mathcal{N}	Scaling factor for the Sharpe ratio \mathcal{S}_t	Eq. (8.34)	1/4
F_{\min}	Minimum capital allocation	Eq. (8.36)	0.0
F_{\max}	Maximum capital allocation	Eq. (8.36)	1.0
θ_k	Allocation transition width	Eq. (8.36)	15

Additional Parameters

Ξ_k	Weighted proportion of time when $k_t < n_t$	Eq. (8.38)	[0, 1]
Ξ_c	Weighted proportion of time when $c_t < c_0$	Eq. (8.37)	[0, 1]
$[\mathcal{L}_k, \mathcal{H}_k]$	Low/High frequency Ξ_k	Sec. 8.2	
$[\mathcal{L}_c, \mathcal{H}_c]$	Low/High frequency Ξ_c	Sec. 8.2	
\mathcal{T}_λ	Timescale of allocation \mathcal{S}_t	Sec. 8.1	$1/ \log(\lambda) $
\mathcal{T}_η	Timescale of TFP shock z_t	Sec. 8.1	$1/ \log(\eta) $
\mathcal{T}_δ	Timescale of capital k_t	Sec. 8.1	$1/ \log(1 - \delta) $
$\mathcal{T}_<$	Average duration of consumption crises	Sec. 8.2	
$\mathcal{T}_>$	Average duration of high output periods	Sec. 8.2	

Table G.1 – Parameters and Notation of the Adapted DSGE model

Appendix **H**

Post-COVID Inflation

H.1 Table of Parameters

Table H.1 – Inactive Central Bank parameter set for Mark-0 Model (see Section 9.3). Parameters, which change for the Reactive Central Bank and Anchored/Floating Trust scenarios, can be taken from sections 9.5 and 9.5 respectively and are indicated in the table with a star. All rates are given in monthly time scales.

	Parameter description	Value
R_0	Ratio of hiring-firing rate (η^+/η^-):	2.0
Θ	Maximum credit supply available to firms :	3.2
Γ_0	Financial Fragility sensitivity:	0.0
ρ^*	Baseline interest rate:	0.001*
α_c	Influence of deposit rates on consumption	12
ϕ_π	Intensity of interest rate policy of Central Bank:	0.0*
π^*	Central Bank inflation target:	0.0*
τ^T	Inflation target parameter:	0.0*
g_w	Factor to adjust wages to inflation expectations:	0.8
g_p	Factor to adjust prices to inflation expectations:	0.8
y_0	Initial production:	0.7
γ_p	Parameter to set adjustment of prices:	0.01
η_0^-	Firing propensity:	0.2
α_Γ	Influence of loans interest rate on hiring-firing policy:	450
c_0	Fraction of savings in consumption budget:	0.5
ϕ	Revival probability per unit time:	0.1
ω	Moving average parameter:	0.2
δ	Dividend rate:	0.02
δ_e	Fraction of energy sector's equity redistributed:	0.04
g_e	Share of Energy Price share in GDP:	0.0325
μ	Easy-Credit policy multiplier:	1.3

H.2 Counterfactual Simulations

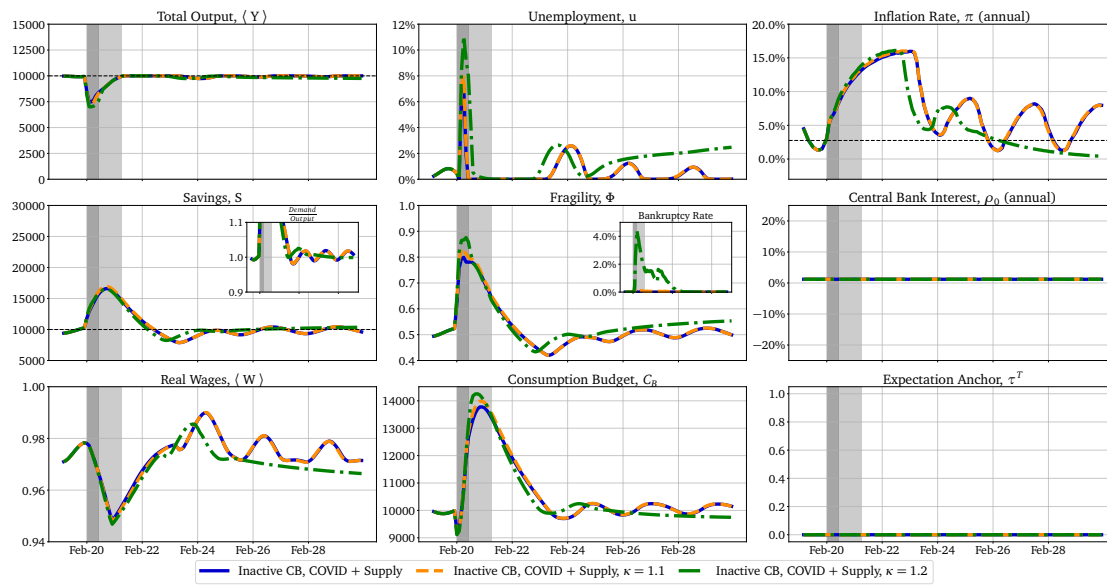


Figure H.1 – Economic dashboard for the Inactive Central Bank scenario without any policy and with a stronger COVID shock: The dynamics for a COVID shock strength of $\kappa = 1.0$ (blue), $\kappa = 1.1$ (orange) and $\kappa = 1.2$ (green) in the Inactive Central Bank scenario without any policy. Once the COVID shock gets too strong, the economy collapses in the long run.

H.3 Monetary Policy

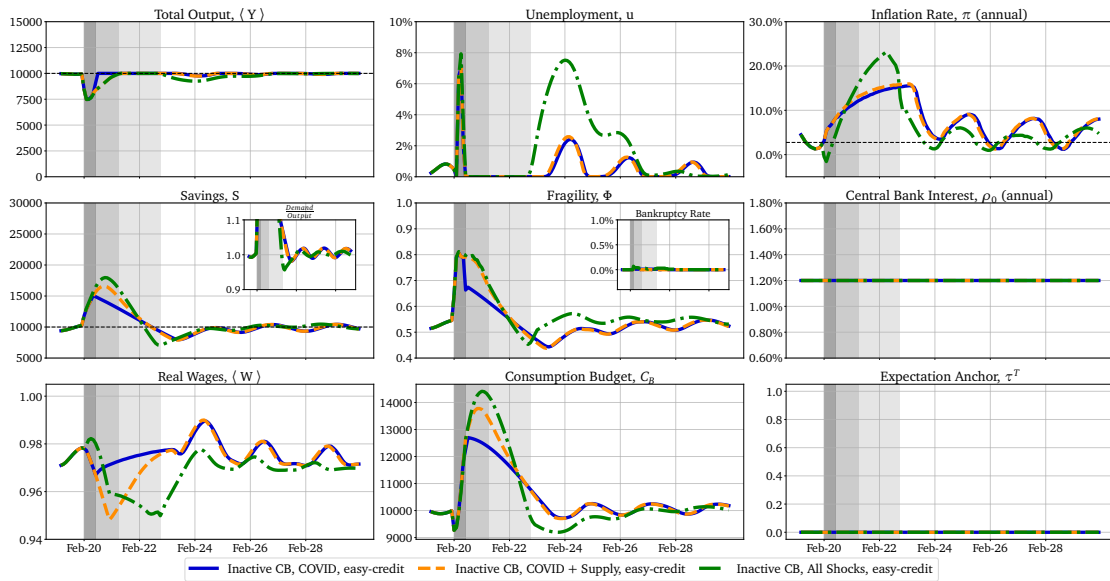


Figure H.2 – Dynamics for the three shocks for the Inactive Central Bank scenario: The dynamics for the three shocks, COVID only (blue), COVID and Supply Chain shock (orange) and all shocks (green) for the Inactive Central Bank scenario with Easy-Credit policy. The areas shaded in grey indicate the duration of the three shocks: the COVID shock lasting until the end of the dark grey area, the supply chain shock until the end of the grey area, and the price shock until the end of the light grey area.

APPENDIX H. POST-COVID INFLATION

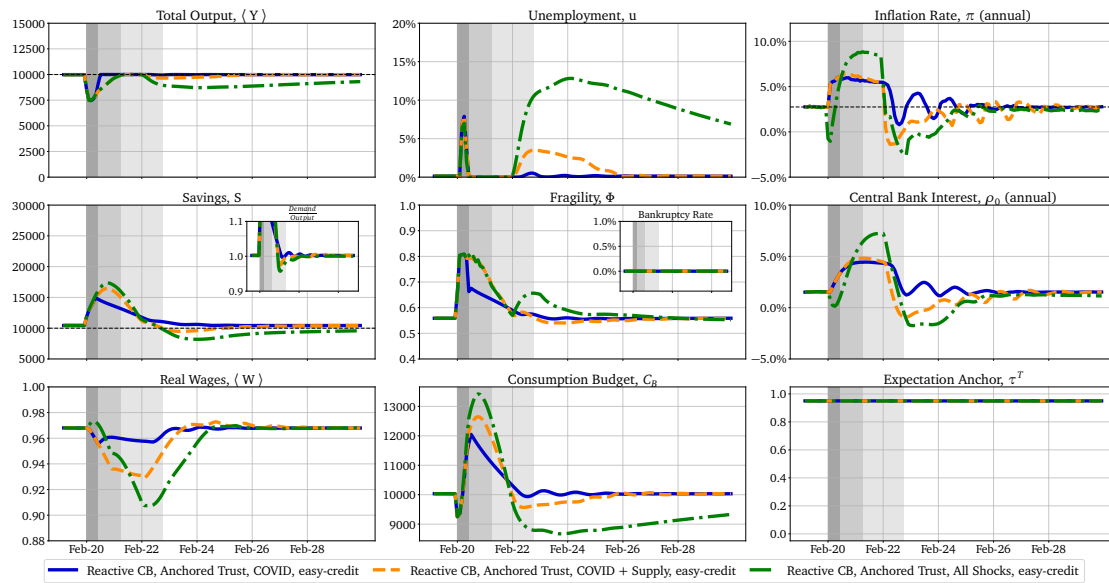


Figure H.3 – Dynamics for the three shocks for the Reactive Central Bank with Anchored Trust scenario: The dynamics for the three shocks, COVID only (blue), COVID and Supply Chain shock (orange) and all shocks (green) to the scenario with reactive Central Bank and Anchored Trust of economic agents with an Easy-Credit policy. The areas shaded in grey indicate the duration of the three shocks: the COVID shock lasting until the end of the dark grey area, the supply chain shock until the end of the grey area, and the price shock until the end of the light grey area.




Figure H.4 – Dynamics for the three shocks for the Reactive Central Bank with Floating Trust scenario: The dynamics for the three shocks, COVID only (blue), COVID and Supply Chain shock (orange) and all shocks (green) to the scenario with reactive Central Bank and Floating Trust of economic agents with an Easy-Credit policy. The areas shaded in grey indicate the duration of the three shocks: the COVID shock lasting until the end of the dark grey area, the supply chain shock until the end of the grey area, and the price shock until the end of the light grey area.




Figure H.5 – Dynamics for all scenarios and all shocks: The full dashboard for Figure 9.6 All dynamics are with Easy-Credit policy in the Inactive Central Bank scenario (blue), Reactive Central Bank with Anchored trust scenario (orange) and Reactive Central Bank with Floating Trust scenario (green). The areas shaded in grey indicate the duration of the three shocks: the COVID shock lasting until the end of the dark grey area, the supply chain shock until the end of the grey area, and the price shock until the end of the light grey area.

H.4 The Effect of Anchoring

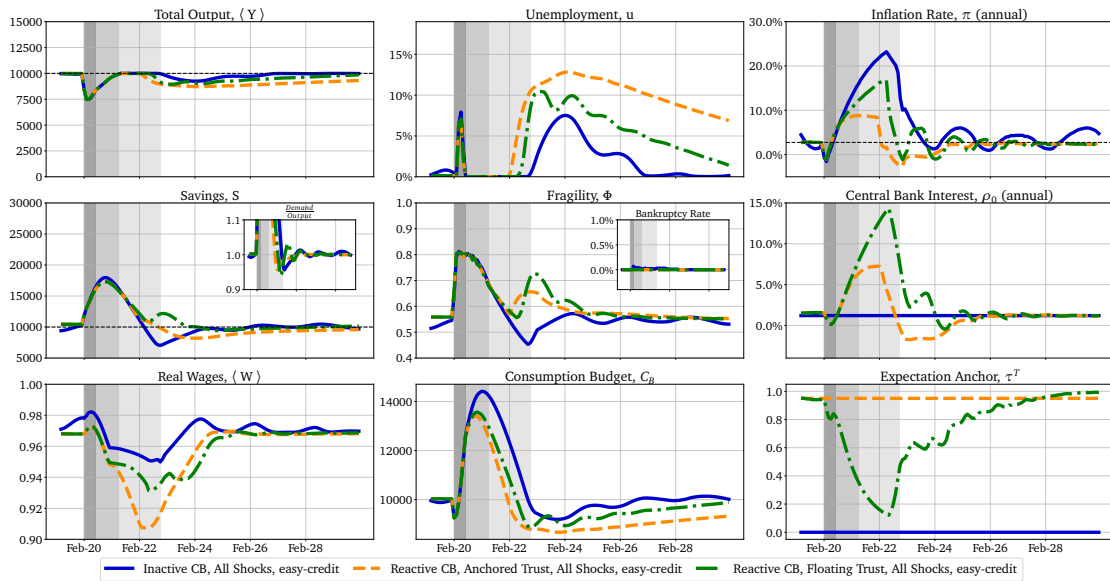


Figure H.6 – Dynamics for all shocks: All dynamics are with Easy-Credit policy in the Inactive Central Bank scenario (blue), Reactive Central Bank with Anchored trust scenario (orange dash) and Reactive Central Bank with no trust scenario (red dot dash). The areas shaded in grey indicate the duration of the three shocks: the COVID shock lasting until the end of the dark grey area, the supply chain shock until the end of the grey area, and the price shock until the end of the light grey area.

H.5 Monetary Policy with Stronger Central Bank Reaction

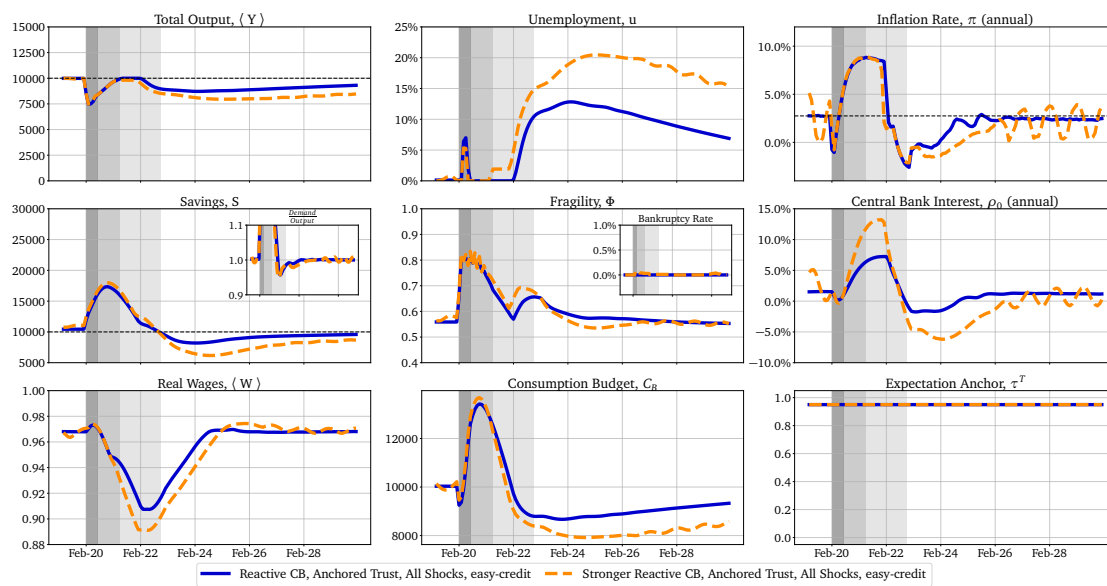


Figure H.7 – Economic dashboard for the stronger Reactive Central Bank with Anchored Trust scenario and all shocks: The dynamics for a Central Bank strength of $\phi_\pi = 1.0$ (blue) and $\phi_\pi = 2.0$ (orange). A stronger central bank is not able to decrease inflation due to an external price shock, but with stronger policies, the consumption of households reduces, which leads to an increase in unemployment.

H.6. SENSITIVITY OF MONETARY POLICY TO α_c AND α_Γ




Figure H.8 – Economic dashboard for the stronger Reactive Central Bank with Floating Trust scenario and all shocks: The dynamics for a Central Bank strength of $\phi_\pi = 1.0$ (blue) and $\phi_\pi = 2.0$ (orange) with Floating Trust of economic agents. A stronger Central bank can reduce excess inflation at the cost of increased unemployment. The central bank must maintain a balance between falling inflation and rising unemployment.

H.6 Sensitivity of Monetary Policy to α_c and α_Γ




Figure H.9 – Sensitivity of α_c for the Reactive Central Bank with Anchored Trust scenario and all shocks: As already described in Gualdi et al. (2017), larger α_c lead to a greater magnification of price trends (cet. par.). This amplification, in turn, causes the fluctuations in unemployment and inflation. Furthermore, larger α_c increases consumption which reduces unemployment due to higher demand.




Figure H.10 – Sensitivity of α_Γ for the Reactive Central Bank with Anchored Trust scenario and all shocks: Larger values of α_Γ increases the influence of financial fragility on the hiring/firing policy of firms (cet. par.). This in turn leads to a larger downward adjustment of the workforce which increases unemployment.

H.7 Helicopter Money

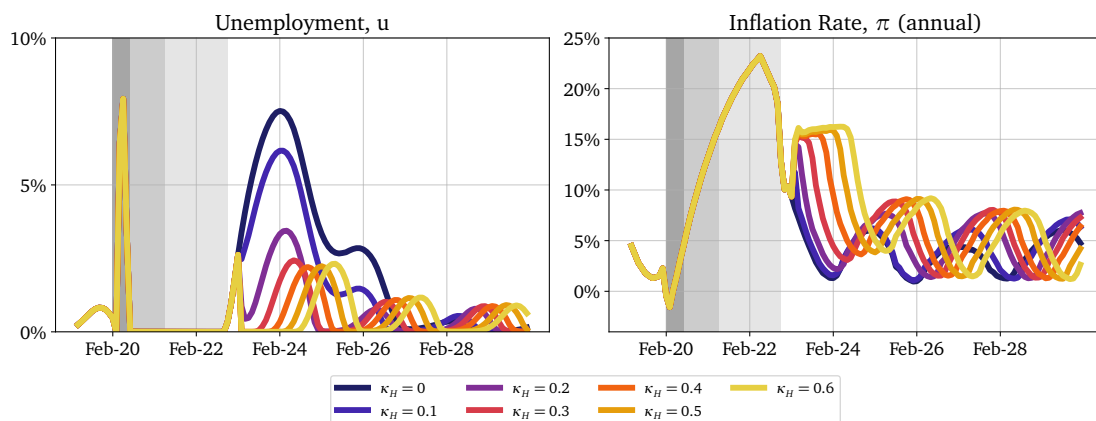


Figure H.11 – Helicopter Money in the Inactive Central Bank scenario and all shocks: Unemployment (left) and inflation (right) for a helicopter drop of size $\kappa_H S$ one month after the price shock. Already with $\kappa_H \geq 0.2$ unemployment is reduced almost to zero. A further increase of Helicopter Money only increases inflation without reducing unemployment more.




Figure H.12 – Helicopter Money in the Reactive Central Bank with Anchored Trust scenario: Unemployment (left) and inflation (right) for a helicopter drop of size $\kappa_H S$ one month after the price shock. With $\kappa_H \geq 0.3$ unemployment is reduced almost to zero. A further increase of Helicopter Money only increases inflation without reducing unemployment further. However, the Central Bank policy manages to keep inflation under control.

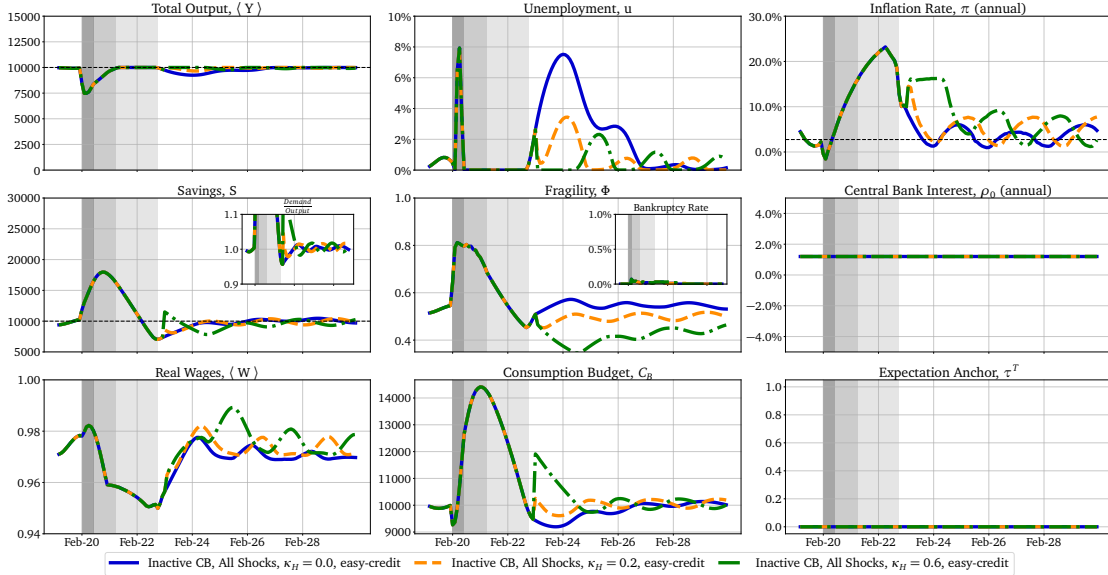


Figure H.13 – Economic dashboard for Helicopter Money in the Inactive Central Bank scenario and all shocks: Full dynamics for scenario without Helicopter Money (blue), for Helicopter Money with $\kappa_H = 0.2$ (orange) and $\kappa_H = 0.6$ (green) one month after the price shock and with Easy-Credit policy.

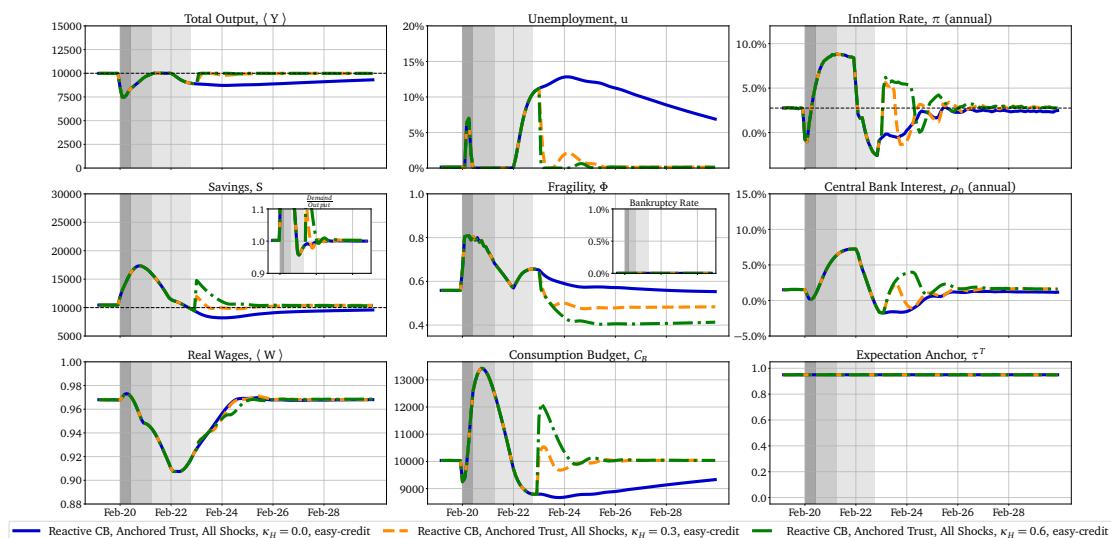


Figure H.14 – Economic dashboard for Helicopter Money in the Reactive Central Bank with Anchored Trust scenario and all shocks: Full dynamics for scenario without Helicopter Money (blue), for Helicopter Money with $\kappa_H = 0.3$ (orange) and $\kappa_H = 0.6$ (green) one month after the price shock and with Easy-Credit policy.




Figure H.15 – Economic dashboard for Helicopter Money in the Reactive Central Bank with Floating Trust and all shocks: Full dynamics for scenario without Helicopter Money (blue), for Helicopter Money with $\kappa_H = 0.2$ (orange) and $\kappa_H = 0.6$ (green) one month after the price shock and with Easy-Credit policy.

H.8 Windfall Tax




Figure H.16 – Windfall Tax in the Inactive Central Bank scenario: Unemployment (left) and inflation (right) for Windfall Tax of $\delta_e + \Delta\delta_e$ one year before the end of the price shock with a duration of two years. With $\Delta\delta_e \approx 4\%$ unemployment is reduced strongly. A further increase of tax does only increase unemployment again. For larger dividends, inflation increases because increased demand due to increased savings.




Figure H.17 – Windfall Tax in the Reactive CB with Anchored Trust scenario: Unemployment (left) and inflation (right) for Windfall Tax of $\Delta\delta_e$ one year before the end of the price shock with a duration of two years. With $\Delta\delta_e \approx 6\%$ unemployment is reduced strongly. A further increase of $\Delta\delta_e$ leads to a significant resurgence of unemployment.

APPENDIX H. POST-COVID INFLATION

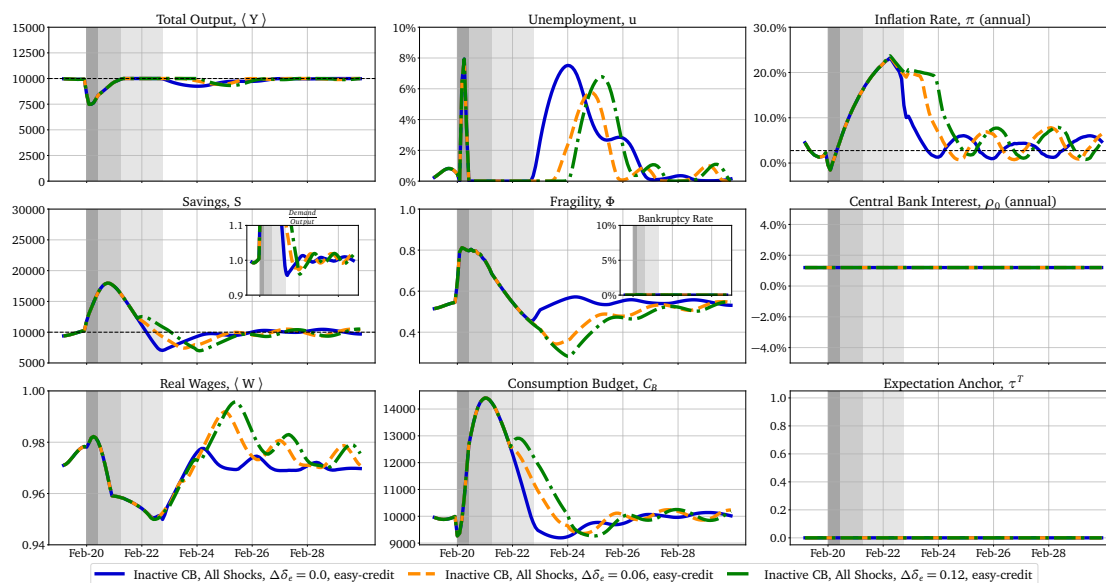


Figure H.18 – Economic dashboard for Windfall Tax in the Inactive Central Bank scenario, all shocks: Full dynamics for the Inactive Central Bank scenario without Windfall Tax (blue), for tax of $\Delta\delta_e = 6\%$ (orange) and $\Delta\delta_e = 12\%$ (green) one year before the end of the price shock with a duration of two years.




Figure H.19 – Economic dashboard for Windfall Tax in the Reactive Central Bank with Anchored Trust scenario and all shocks: Full dynamics for scenario without Windfall Tax (blue), for tax of $\Delta\delta_e = 6\%$ (orange) and $\Delta\delta_e = 12\%$ (green) one year before the end of the price shock with a duration of two years.

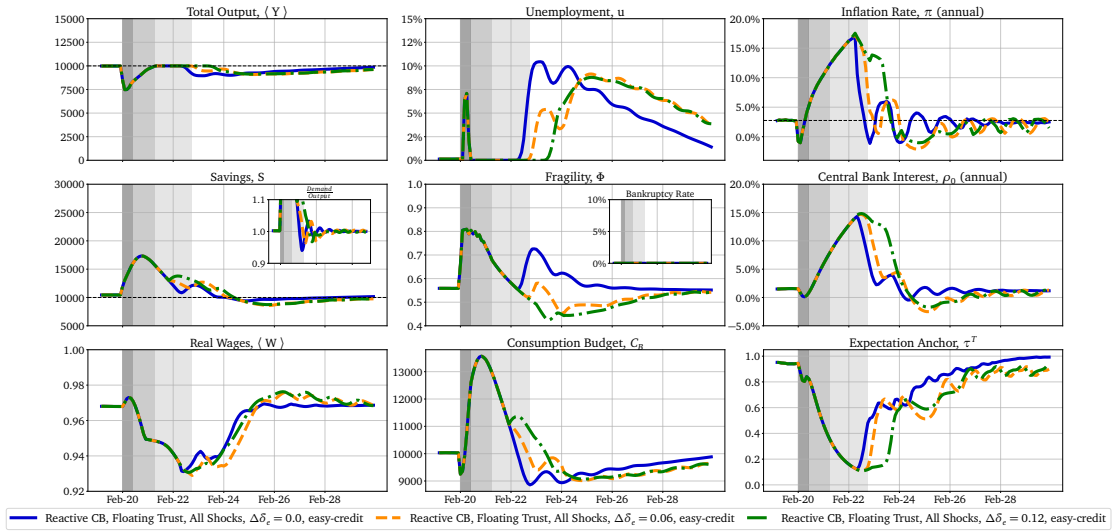


Figure H.20 – Economic dashboard for Windfall Tax in the Reactive Central Bank with Floating Trust scenario and all shocks: Full dynamics for scenario without Windfall Tax (blue), for tax of $\Delta\delta_e = 6\%$ (orange) and $\Delta\delta_e = 12\%$ (green) one year before the end of the price shock with a duration of two years.

H.9 Sloppiness

To assess the sensitivity of the model’s realisation we here follow the approach introduced in Part I of this thesis, with the MSE loss function. The eigenvalue spectrum spans several decades (Figure H.21), which indicates the sloppiness of the model at that set of parameters. Based on the spectrum, we select the four largest eigenvectors to consider for our analysis (see Figures H.22, H.23, H.24).

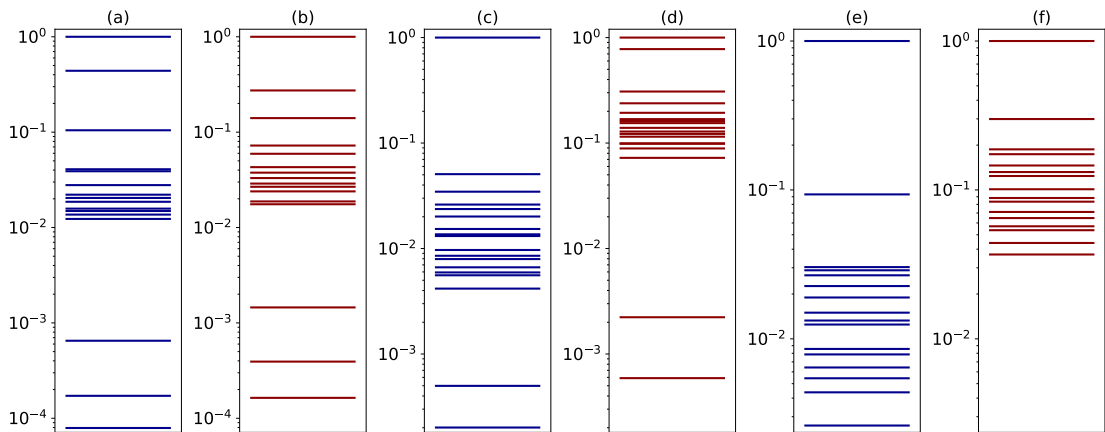


Figure H.21 – Eigenvalue Spectrum for three scenarios and all shocks. Sloppiness parameter $\epsilon = 0.01$, $n_{seeds} = 50$ target variable inflation (red) unemployment (blue), (a,b) Inactive Central Bank, (c,d) Reactive Central Bank with Anchored Trust, (e,f) Reactive Central Bank with Floating Trust.

APPENDIX H. POST-COVID INFLATION




Figure H.22 – Eigenvectors of the Hessian matrix for the Inactive Central Bank Scenario and all shocks: The length of the eigenvectors is normed to 1 and the colors indicate the target variables, unemployment (blue) and inflation (red).




Figure H.23 – Eigenvectors of the Hessian matrix for the Reactive Central Bank with Anchored Trust scenario and all shocks: The length of the eigenvectors is normed to 1 and the colors indicate the target variables, unemployment (blue) and inflation (red).

H.10. SENSITIVITY TO THE INDEXATION PARAMETERS g_p AND g_w




Figure H.24 – Eigenvectors of the Hessian matrix for the Reactive Central Bank with Floating Trust scenario and all shocks: The length of the eigenvectors is normed to 1 and the colors indicate the target variables, unemployment (blue) and inflation (red).

H.10 Sensitivity to the Indexation Parameters g_p and g_w




Figure H.25 – Sensitivity of $g_p = g_w = const$ for the Reactive Central Bank with Anchored Trust scenario and all shocks: Increasing $g_p = g_w$ increases slightly inflation (cet. par.), causing slightly higher interest rates which leads to a reduction of consumption budget that reduces demand and therefore increases unemployment.




Figure H.26 – Sensitivity of $\frac{g_p + g_w}{2} = const$ for the reactive Central Bank with Anchored Trust scenario and all shocks: (blue) When $g_w > g_p$, there is a higher bargaining power of worker which increases wages (cet. par.), therefore there is a higher consumption budget and higher demand which decreases unemployment. (yellow) For larger market power ($g_p > g_w$), there is the reverse effect of decreased wages causing a decrease in consumption budget and demand which therefore increases unemployment.

Titre: Génération et Exploration de Scénarios dans des Modèles Économiques Multi-Agents

Mots clés: Macroéconomie; Agent-based Models; Économie de la complexité

Résumé: L'économie, caractérisée par la non-linéarité, l'adaptabilité et la dynamique hors équilibre, présente des phénomènes émergents tels que des crises et des inégalités, façonnés par les réactions des agents et les politiques publiques. L'utilisation de modèles multi-agents (ABM) est une approche récente en macroéconomie qui génère ces phénomènes en simulant une multiplicité d'agents hétérogènes en interaction. Bien que cette méthode puisse mener à des phénomènes émergents, elle a souvent été critiquée comme étant une boîte noire où les mécanismes de causalité ne sont pas clairs et où il existe un ensemble trop vaste de dynamiques générées. Cette thèse propose une méthode pour aborder la question fondamentale : Quel est l'ensemble de phénomènes qualitativement différents qu'un Modèle Multi-Agent Macroéconomique (MABM) peut générer, et qu'est-ce qui régit leurs transitions ?

S'inspirant de la recherche en biophysique, l'idée centrale postule qu'il n'y a que quelques combinaisons critiques de paramètres qui gouvernent un résultat spécifique. En exploitant ces combinaisons avec un algorithme de *montée* en gradient, on peut efficacement découvrir l'ensemble de phénomènes différents qu'un MABM peut présenter. La pertinence de cette approche réside dans le fait de révéler d'une structure plus simple sous la complexité du

MABM, ouvrant la voie à des politiques efficaces qui abordent les directions critiques des paramètres. Cela suggère également que malgré la complexité d'un MABM et le nombre élevé de paramètres, l'ajustement de ces modèles ne nécessite que l'ajustement des directions critiques pour avoir un pouvoir prédictif.

La première partie de cette thèse développe les méthodes sous-jacentes à l'algorithme, mettant en évidence son efficacité sur les "Fourmis de Kirman", un modèle simple de comportement des agents. L'algorithme est ensuite démontré sur le MABM stylisé Mark-0 qui présente une phénoménologie riche avec un ensemble connu de phénomènes. Je montre comment nous pouvons récupérer cet ensemble de phénomènes malgré la complexité de la dynamique du modèle. La dernière partie de cette thèse adopte en fait une approche inverse, intégrant des interactions inter-agents dans des modèles macroéconomiques d'équilibre, dévoilant des phases émergentes et des crises endogènes dans ces modèles. En bref, cette thèse navigue dans le terrain complexe des ABMs, dévoilant leur potentiel dans la génération de différents scénarios qui peuvent être utilisés pour éclairer les politiques publiques dans des systèmes dynamiquement complexes.

Title: Scenario Discovery in a Complex Economy: Exploring the Parameter space of Agent-based Models

Keywords: Macroeconomics; Agent-based Models; Complexity Economics

Abstract: The economy, characterized by non-linearity, adaptability, and non-equilibrium dynamics, exhibits emergent phenomena, such as crises and inequalities, shaped by agents' reactions and policy interventions. Agent-based Modeling (ABM) is a recent modeling approach in macroeconomics that generates these phenomena from the ground up by simulating a multiplicity of heterogeneous interacting agents. While this method can generate emergent phenomena, it has often been critiqued as a black-box where causal mechanisms are unclear and there too vast set of generated dynamics. This thesis proposes a method to approach the fundamental question: What is the set of qualitatively different phenomena can an Macroeconomic Agent-based Model (MABM) generate, and what governs their transitions?

Drawing on research in biophysics, the core idea posits that there are only a few critical parameter combinations that govern a specific outcome. Exploiting these with a gradient ascent algorithm, one can effectively uncover the set of different phenomena a MABM can recover. The significance of this approach lies in revealing a simpler struc-

ture beneath MABM complexity, paving the way for effective policies that address critical parameter directions. It also suggests that despite the complexity of an MABM and the high number of parameters, fitting these models requires only fitting critical directions to have predictive power.

The first part of this thesis develops the methods behind the algorithm, highlighting its power on Kirman's Ants, a simple model of agent-herding. The algorithm is then demonstrated on the stylized Mark-0 MABM that has a rich phenomenology with a known set of phenomena. I show how we can recover this set of phenomena despite the complexity of the model's dynamics. The final part of this thesis actually adopts a reverse approach, embedding intra-agent interactions in equilibrium macroeconomic models, unveiling emergent phases and endogenous crises in these models. In its essence, this thesis navigates the intricate terrain of ABMs, unraveling their potential in generating different scenarios that can be used to inform policy decisions in dynamically complex systems.