



HAL
open science

Asset pricing with heterogeneous agents and non-normal return distributions

Arthur Beddock

► **To cite this version:**

Arthur Beddock. Asset pricing with heterogeneous agents and non-normal return distributions. Business administration. Université Paris sciences et lettres; Tilburg School of Economics and Management (TISEM), 2021. English. NNT : 2021UPSLD003 . tel-03545639

HAL Id: tel-03545639

<https://theses.hal.science/tel-03545639v1>

Submitted on 27 Jan 2022

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.



THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à l'Université Paris-Dauphine
Dans le cadre d'une cotutelle avec Tilburg University

Evaluation d'Actifs avec Agents Hétérogènes et Rendements Non-normaux

Asset Pricing with Heterogeneous Agents and Non-normal Return Distributions

Soutenue par

Arthur BEDDOCK

Le 2 septembre 2021

Ecole doctorale n° ED 543

Ecole doctorale SDOSE

Spécialité

Gestion, Finance

Composition du jury :

Serge DAROLLES Professeur, Université Paris-Dauphine	<i>Président</i>
Milo BIANCHI Professeur, Toulouse School of Economics	<i>Rapporteur</i>
Esther EILING Professeur, University of Amsterdam	<i>Rapporteur</i>
Bas WERKER Professeur, Tilburg University	<i>Examineur</i>
Elyès JOUINI Professeur, Université Paris-Dauphine	<i>Directeur de thèse</i>
Frans de ROON Professeur, Tilburg University	<i>Directeur de thèse</i>
Clotilde NAPP Directrice de recherche, Université Paris-Dauphine	<i>Directrice de thèse</i>

Asset Pricing with Heterogeneous Agents and Non-normal Return Distributions

Proefschrift ter verkrijging van de graad van doctor aan Tilburg University op gezag van de rector magnificus, prof. dr. W.B.H.J. van de Donk, en de Université Paris Dauphine op gezag van de president, prof. E.M. Mouhoud, in het openbaar te verdedigen ten overstaan van een door het college voor promoties aangewezen commissie in de Aula van Tilburg University op donderdag 2 september 2021 om 16.00 uur door

ARTHUR NICOLAS GEORGES MICHEL BEDDOCK

geboren te Paris, Frankrijk

Promotores

Prof. dr. Elyès Jouini (Université Paris-Dauphine)

Prof. dr. Frans de Roon (Tilburg University)

Copromotor

Dr. Clotilde Napp (Université Paris-Dauphine)

Leden Promotiecommissie

Prof. dr. Milo Bianchi (Toulouse School of Economics)

Prof. dr. Serge Darolles (Université Paris-Dauphine)

Dr. Esther Eiling (Universiteit van Amsterdam)

Prof. dr. Bas Werker (Tilburg University)

Acknowledgements

I have spent the last five years writing this dissertation, and it goes without saying that none of it could have been possible without the assistance and support of the numerous people who accompanied me through the many ups and downs along the way. I would like to take the time now to thank them all for their participation in this unique adventure!

First, my deepest gratitude goes to my supervisors, Elyès Jouini, Clotilde Napp, and Frans de Roon. It has been a privilege to work with you and under your supervision. I am thankful for your constant guidance, your invaluable advice, and, especially, for everything I had the chance to learn by your side. This will follow me for the rest of my career, and I am thus forever indebted to you. Elyès, thank you also for having been an outstanding co-author and for your faith in me.

I am also grateful to Milo Bianchi, Serge Darolles, Esther Eiling, and Bas Werker for having accepted to be part of my PhD committee. I really appreciate the time and effort you have spent reviewing my chapters. The comments you gave me have already greatly improved my work, and I am sure that they will also be helpful for future revision steps.

Next, I sincerely thank Paul Karehnke, who has not only been a great co-author but who has also been a tremendous support over the last five years. From our first meeting in Sydney to today, your help has always been precious to me, and I truly hope that our collaboration will continue in the coming years.

ACKNOWLEDGMENTS

Moreover, I have had the chance to interact with professors, from Dauphine, Tilburg, and elsewhere, who have guided me and given me useful advice. In particular, I thank Julio Crego, Marie-Pierre Dargnies, Jérôme Dugast, Sebastian Ebert (with whom I already look forward to working next year at the Frankfurt School of Finance & Management), Edith Ginglinger, Carole Gresse, Evan Jo, Bradley Paye, Fabrice Riva, Mattias Villani, and Ole Wilms. I am also thankful to all the participants of the Finance Theory Group Summer School 2019, the Financial Risk International Forum 2021, the AFFI Conference 2021, and of Dauphine and Tilburg brown bags for their helpful feedback.

On another note, I would like to express my gratitude to the French State for funding my doctoral studies, to the American Finance Association for offering me a grant to attend its 2019 conference, and to the Université Paris Dauphine - PSL for additional financial assistance. On the administrative side, thank you to Françoise Carbon, Valérie Fleurette, Loes de Groot, Ank Habraken, and Stéphanie Salon, who have made sure that my co-supervision was carried out in the most appropriate way and who have always been there to answer my questions. I also had a great time teaching corporate finance, microeconomics, and statistics to Dauphine students, and I thus would like to acknowledge that the weekly break from research that they offered me was a key ingredient in the completion of my thesis.

Additionally, let me emphasize that it would be wrong to say that the daily life of a PhD student is a lonely journey. From my research master—where I became friends with Augustin, Benjamin, Dup, Etienne, Félix, Louis, and Pierre, who have continuously supported me ever since and who even tried to read my papers—to my fifth year of joint PhD, I have encountered a lot of people who have made my research life very pleasant. I thus thank all past and present finance PhD students that I met. In particular, thank you to Béatrice, Bouchra, Camille, Caroline, Hugo, John, Louis, Saad, Théo, and Thomas.

As research is not limited to finance, I would also like to mention Alexis and Armand, two prominent young mathematicians and friends, and my former roommate Morgan, who is currently completing his PhD in economics, whose advice has always been beneficial.

Furthermore, I warmly thank my non-academic friends—especially Louis, Nicolas, Presnel, Samy, and Valentin—for distracting me and for all the good times we spent together. I also have a special thought for Daniel.

Last but not least, I will never forget the unconditional support and love I continuously received from my parents, Pascale and Richard, my uncle, Jean-François, my siblings, Alice, Clara, Eléonore, Nicolas, Théo, and Victor, and my girlfriend, Laureen. Because it would take too much time and place to enumerate everything you brought to me, I will keep it short: thank you for being you!

Paris, June 2021.

Contents

Introduction	1
1 Live fast, die young: equilibrium and survival in large economies	5
1.1 Introduction	8
1.2 The model	14
1.3 Representative agent and equilibrium	18
1.4 Market characteristics	20
1.4.1 Risk-free rate and market price of risk	20
1.4.2 Market volatility	23
1.4.3 Trading volume	25
1.5 Consumption shares and utility-maximizing agents	29
1.5.1 Surviving and market-dominating agents	29
1.5.2 Utility-maximizing agents	32
1.6 Conclusion	36
1.A Proofs	38
1.B Useful computations and additional results	48
1.B.1 Moments with respect to the density $\tilde{V}_{\delta,\rho,k,t}$	48
1.B.2 Consumption shares and utility-maximizing agents in the correlated case	49
2 Disagreeing forever: a testable model with non-vanishing belief heterogeneity	57

CONTENTS

2.1	Introduction	60
2.2	Theoretical part	66
2.2.1	An overlapping heterogeneous generations model	66
2.2.2	A dual approach of the model	71
2.2.3	The stock price and its dynamics	74
2.3	Empirical test of the model	78
2.3.1	Market belief dispersion data	78
2.3.2	Predicting market returns	80
2.3.3	Predicting market volatility	82
2.4	Conclusion	85
2.A	Proofs	87
2.B	Additional empirical results	95
2.B.1	Predicting market excess returns	95
2.B.2	Rolling window analysis of the raw simple market index returns at the quarterly horizon	95
3	Two skewed risks	109
3.1	Introduction	112
3.2	A simple skewed bivariate distribution	116
3.2.1	Definition of the SBN distribution	116
3.2.2	Empirical fit	119
3.3	Portfolio choice and asset pricing	122
3.3.1	Optimal choice	122
3.3.2	Equilibrium returns	127
3.4	Conditional risk measures	131
3.4.1	ΔCoVaR	132
3.4.2	CoES	135
3.5	Conclusion	137

3.A	Details on the distribution	139
3.A.1	Calculating the scaling factors	139
3.A.2	Moments of the distribution	140
3.A.3	Marginal distribution	143
3.A.4	Conditional distribution	144
3.B	Proofs	145
3.C	Additional results	153
3.C.1	Sharpe ratio in the portfolio choice setting	153
3.C.2	Sample variance of co-skewness	153
	Substantial summary in French	167

CONTENTS

Introduction

This thesis is constituted of three chapters that represent individual papers in the area of asset pricing and behavioral finance. Its overall goal is to study the impacts of releasing simplifying unrealistic assumptions frequently made in asset pricing models. Specifically, while agents are often assumed to be identical (which eases their aggregation into one representative agent) and returns to be normally distributed (which has many computational advantages), I focus on agent heterogeneity and on non-normal asset return distributions. The first chapter—published in *Economic Theory* in April 2021—is entitled *Live fast, die young: equilibrium and survival in large economies* and is jointly written with Elyès Jouini. The second one is named *Disagreeing forever: a testable model with non-vanishing belief heterogeneity* and is my job market paper. Lastly, the article presented in Chapter 3 is a joint work with Paul Karehnke and is called *Two skewed risks*.

The first two chapters deal with agent heterogeneity and consider models that incorporate a continuum of heterogeneous agents who agree to disagree. They echo recent works showing that heterogeneity has large impacts on the market characteristics and should therefore be taken into account.

More precisely, Chapter 1 is a theoretical paper where Elyès Jouini and I study a continuous-time economy populated by agents who have heterogeneous beliefs about what the future holds and heterogeneous time preferences. One of our contributions is to further allow for some correlation between these two types of heterogeneity. Empirical evidence indeed shows that they are negatively correlated: the most optimistic agents are also the most

patient ones. We therefore study the theoretical impacts of such correlated heterogeneities on the behavior of financial markets. We fully characterize the risk-free rate which is procyclical and the market price of risk which is countercyclical, and we show that a negative correlation between the two types of heterogeneity reduces the former and enhances the latter. In addition, we assume that an asset, whose dividend process is given by the total endowment of the economy, is available for trading. Importantly, we derive that a higher belief dispersion increases the overall trading volume, and that the case where the most optimistic investors are also the most patient ones induces some excess volatility in the market. In the last part of the paper, assuming that the two types of heterogeneity are uncorrelated, we finally study the characteristics of some specific agents, namely the surviving agent, the ex-post utility-maximizing agent, and the ex-ante utility-maximizing agent. We thereby contribute to the literature by showing that a shorter life might be more rewarding than a longer one, as the surviving agent of the economy is not necessarily the one who maximizes her utility over her lifetime.

The second chapter *Disagreeing forever: a testable model with non-vanishing belief heterogeneity* similarly deals with a continuum of heterogeneous investors but focuses only on the belief heterogeneity. Another difference with Chapter 1 is that I study more general utility functions as I assume that the investors are endowed with constant relative risk aversion utility functions rather than logarithmic ones. Furthermore, the paper contains both a theoretical and an empirical part. In fact, my main contribution is to derive a tractable overlapping heterogeneous generations model where the belief heterogeneity does not vanish as time goes by, and I thus obtain a stationary model that I can test empirically over long periods. In the theoretical part, I first compute the model equilibrium and determine the planned actions of the successive generations of investors, who keep the same belief from a generation to another and maximize the utility they derive from either consumption or bequest. Considering a similar stock as the one in Chapter 1, I further look at the implications of belief heterogeneity on the characteristics of this stock. In particular, I find that

the stock mean return and volatility both increase when the belief dispersion increases, and I derive non-vanishing belief dispersion effects, meaning that these positive relations should remain no matter the horizon considered. Thus, I specify four hypotheses that are tested in the empirical part of the paper. More precisely, using the Institutional Brokers Estimate System Unadjusted Summary database to construct the market belief dispersion variables, the empirical analysis shows that the positive relation between the market returns and the market belief dispersion is verified in the data for all horizons, and empirical evidence further points more towards the approval of the positive model-implied relation between the market volatility and the market belief dispersion than towards its rejection.

Finally, Chapter 3 departs from the assumption that asset returns are normally distributed (implying that they have a null skewness). Defining and using a skewed distribution in a simple two-asset framework, Paul Karehnke and I analyze how skewness and its interaction with correlation affect portfolio choice, asset prices, and popular risk metrics.

More precisely, we assume that returns follow the split bivariate normal distribution, which is a mixture of four scaled bivariate normal distributions. Skewness is thereby introduced by allowing the standard deviation parameters to differ in each of the four quadrants. Bivariate Kolmogorov-Smirnov tests and a comparison of implied and empirical co-skewnesses show that this distribution provides a good empirical fit, and is therefore well-suited for our study. Importantly, our focus is on simplicity and intuition due to the similarities of the split bivariate normal with the bivariate normal distribution. In a pure exchange economy, we find that capital asset pricing model alphas are non-zero when the assets have different levels of skewness. Additionally, one of our contributions is to show theoretically that the asset that underperforms also has high systematic skewness, high idiosyncratic skewness, and high maximum returns, as empirical studies suggest. Moreover, the main result of our portfolio choice analysis is to show that underdiversification becomes optimal when the correlation is high and skewnesses differ. Therefore, we highlight the role of correlation in the diversification-skewness trade-off, and we thus extend Mitton and Vorkink

(2007) who consider only one skewed asset. Lastly, we provide simple expressions for some bivariate risk metrics, namely ΔCoVaR (defined in Adrian and Brunnermeier, 2016) and the conditional expected shortfall (Acharya et al., 2017), when risks follow the split bivariate normal distribution, and thereby show that skewness has a large quantitative impact on them.

Chapter 1

Live fast, die young: equilibrium and survival in large economies

Joint work with Elyès Jouini.

Published in *Economic Theory*, Volume 71, Issue 3, April 2021.

Note that the version presented in the thesis slightly differs from the one published because it takes into account some minor comments of the committee members.

Abstract

We model a continuous-time economy with a continuum of investors who differ both in belief and time preference rate, and analyze the impacts of these heterogeneities on the behavior of financial markets. In particular, we allow the two types of heterogeneity to be correlated: a negative correlation means that the most optimistic agents are also the most patient ones. We fully characterize the risk-free rate, which is procyclical, and the market price of risk, which is countercyclical. When the two types of heterogeneity are negatively correlated, the former is lower and the latter higher compared to the standard case. A negative correlation also leads to a higher market volatility. Moreover, we find that the trading volume increases with the variance of the belief heterogeneity distribution. Finally, the surviving agent of this economy is not necessarily the one who maximizes her utility over her lifetime: a shorter life might be more rewarding than a longer one.

Keywords: Heterogeneous beliefs, Heterogeneous time preference rates, Continuum of agents, Asset pricing, Market elimination, Surviving agent

JEL classification: D53 - D90 - G02 - G11 - G12

1.1 Introduction

A common assumption made in asset pricing models is that all investors have a rational belief. Most models also rely on the assumption that all investors have the same time preferences. Although these assumptions are useful as they permit to aggregate all the agents into one representative agent, empirical evidence indicates that some investors are more optimistic than others and some more patient than others, questioning the pertinence of such hypotheses.

In spite of this evidence, there have been several arguments to support the first assumption. Following Friedman (1953), it has been argued that, although some investors might have a biased belief towards optimism or pessimism, they should not be of interest as this should lead them to make wrong choices and to go extinct (see, e.g., Sandroni, 2000). However, as Kogan et al. (2006) point out, survival and market impact are different concepts and they need to be studied separately. In particular, Yan (2008) shows that the elimination process takes a long time and that biased investors should therefore not be neglected. A recent study by Bottazzi et al. (2018) also states that heterogeneous agents are not necessarily eliminated in the long run and that the non-optimality of an agent's portfolio can correct for the inaccuracy of her belief, leading to her survival. A second important argument in favor of the belief homogeneity assumption argues that, as there is no reason for an average bias to exist in the economy, agents should be rational on average and the effects induced by biased investors should cancel out. Jouini and Napp (2011) find that this is not the case and that unbiased disagreement can not be considered as agreement. Finally, similarly to the pragmatic beliefs concept of Hvide (2002), there is an argument supporting the belief homogeneity assumption which states that irrational agents, observing that rational agents are being more successful, should adopt the same belief as the most successful ones. Questioning this third argument, Jouini and Napp (2016) show that irrational agents might do better than rational ones.

The time preference rate homogeneity assumption has also been examined. Following

the seminal work of Samuelson (1937), it has been widely accepted that a unique discount rate can be used to condense intertemporal choices. However, empirical studies (see, e.g., Frederick et al., 2002) show that this assumption does not hold in the real world and that there exists a great variety of time preference rates across investors. At a country level, Wang et al. (2016) highlight this heterogeneity and show that, in addition to economic factors, it can also be explained by cultural ones.

This suggests that these two assumptions might be unreasonable and that heterogeneous agents, both in belief and time preference rate, could have an impact on financial markets. Empirical works support this rationale and show how important are the belief and time preferences heterogeneities on various markets (see, e.g., Buraschi and Jiltsov, 2006 (option markets), Beber et al., 2010 (currency markets), Buraschi et al., 2014 (credit markets)). Investors' heterogeneity also explains, at least partly, empirical facts, as the implied risk aversion smile (Ziegler, 2007).

In this paper, we further investigate these impacts.¹ More precisely, we study their joint impact and allow the two types of heterogeneity to be correlated. To the best of our knowledge, we are the first to theoretically consider such a correlation, which seems to exist empirically. In fact, survey evidence suggests that we can assume a negative correlation, or, stated differently, that the most optimistic agents are also the most patient ones.² We therefore consider a continuous-time equilibrium model with a continuum of heterogeneous agents who do not share the same belief about the future nor the same time preference rate, and we allow some correlation to exist between the characteristics of each agent. Hence, our economy is characterized by the presence of optimistic and pessimistic agents and, for an identical belief, patient and impatient ones. Unlike other studies (see, e.g., Li, 2007, Berrada et al., 2018), we assume that the agents do not learn. We make this assumption because

¹We choose to focus on belief and time preference rate heterogeneity for modeling convenience. Those are also two of the most popular types of heterogeneity considered in the literature. We leave the study of other types of heterogeneity, e.g., the heterogeneity in risk aversions, for future work—note however that such a type of heterogeneity leads to less tractable results (see, e.g., Cvitanic et al., 2012).

²We thank Luc Arrondel who checked this fact on the PATER survey and obtained a negative correlation in a qualitative regression framework. The result is significant at a 5% level.

we want to focus on market elimination of given beliefs rather than on their elimination through possible learning. Moreover, it allows us to be consistent with the view that the agents' heterogeneity comes from psychological biases. If this is the case, heterogeneity can be explained by behavioral distortions and there is therefore no added value to assume learning. The fact that we find that some agents with a wrong belief might be better off than some rational agents further adds to the debate initiated by Grossman and Stiglitz (1980) on the economic rationale for learning which suggests that agents might have no incentive to learn.

We first compute the equilibrium of our model. We then study the impact investors' heterogeneity has on some equilibrium characteristics and, in particular, we determine how the correlation between the two types of heterogeneity affects them. Focusing on an economy with uncorrelated heterogeneities, we also determine which agent survives in the long run and which agent maximizes her expected utility (ex-post and ex-ante).

Belief and time preference rate heterogeneities impact the risk-free rate and the market price of risk (see, e.g., David, 2008). In fact, in our economy, we observe countercyclical market prices of risk and procyclical risk-free rates. These results, obtained considering a continuum of heterogeneous agents, are in line with, and complement, previous studies, such as Jouini and Napp (2011) or Bhamra and Uppal (2014), who consider models with only two groups of investors. They are also consistent with empirical observations that "equity risk premia seem to be higher at business cycles troughs than they are at peaks" (Campbell and Cochrane, 1999) and that the short term rate is a procyclical indicator of economic activity (see, e.g., Friedman, 1986). These effects are amplified when there is an average optimistic bias in the economy and lowered when investors are pessimistic on average. Moreover, unlike previous studies, we find that the higher the variance of the belief distribution, the higher these effects. We also find that the asymptotic market price of risk equals the one obtained in a homogeneous framework, whereas the asymptotic risk-free rate is smaller than the standard one. This is due to the fact that only the most patient investors, characterized by an almost

null time preference rate, survive in the long run. The study of the correlation between the two types of heterogeneity leads to an additional conclusion: a negative correlation between beliefs and time preference rates reduces the risk-free rate and enhances the market price of risk. This new result is interesting in light of the risk premium and the risk-free rate puzzles. We further consider an asset whose dividend process is given by the total endowment of the economy and study the volatility of the stochastic differential equation satisfied by this asset, namely the market volatility. Unlike the case where the beliefs and the time preferences are uncorrelated, we find that the two types of heterogeneity do impact the market volatility when they are correlated. In particular, if we assume a negative correlation between them, we observe some excess market volatility. The more impatient the agents are on average, the less important this effect is. Considering the time preference rate heterogeneity and allowing some correlation between beliefs and time preferences, we therefore complement Atmaz and Basak (2018), who derive a similar result of excess volatility in presence of a continuum of agents who are only heterogeneous in beliefs. In their settings, the main driver of the excess volatility is the belief dispersion whereas it does not play a role in our analysis. This is due to the fact that we consider agents endowed with logarithmic utility functions, which are known to mitigate the belief dispersion effect. The effect we establish is then an additional effect due to the correlation between the two types of heterogeneity. In a finite time economy, populated by a patient and an impatient agent, Li (2007) derives a similar result: he finds that the market volatility is higher than the aggregate endowment volatility if the patient agent is more optimistic than the impatient one. Hence, our result extends his finding to the case of a continuum of agents, which allows us to study in more depth the role of correlation. The last market characteristic that we study is the trading volume. Unlike the homogeneous agents case, where no one trades as all investors act identically, we observe some trading in the economy. At the investor level, we see that the more patient an agent is, the more she trades. This result is due to a wealth effect and is strengthened when the correlation is higher in absolute value. Looking at the impact of beliefs on the trading volume for agents

sharing the same time preference rate, we see that there is no clear relation between trading volume and belief bias. At the global level, as some agents are progressively driven out of the market, we find that the trading volume decreases with time in the uncorrelated case. This overall trading volume also depends positively on the variance parameter of the distribution of beliefs. In fact, more heterogeneous agents imply more trading possibilities.

Finally, in the last part of the paper, we study some characteristics of the agents. Consistent with Yan (2008), we find that the surviving agent of the economy is the most patient of the unbiased agents. Similarly, forming groups of investors sharing the same time preference rate, we show that the surviving agent of each group is the unbiased one. However, we show that this agent does not necessarily have the highest ex-post or ex-ante utility of her group. In particular, if there is an average optimistic bias in the economy, the ex-post utility-maximizing agent is an optimistic agent whose bias is between the rational one and the average one. This result is in line with Jouini and Napp (2016) who show that “even if they are eliminated in the long run, irrational agents might rationally stay irrational,” in the sense that, for given time preferences, having a biased belief might allow an agent (thus defined as an irrational one) to have a higher level of utility over her whole life than the utility level of an unbiased (rational) agent sharing the same time preferences. In terms of utility, such an agent therefore does not have any incentive to change her belief to become a rational agent. Stated differently, this is consistent with the fact that, under certain circumstances, a shorter life (i.e., associated with a low survival rate) might be more rewarding than a longer one. Our result complements Jouini and Napp (2016)’s findings as we deal with a continuum of agents who are therefore price takers. Indeed, in their two-agent model, each agent has an impact on the equilibrium price, and it is therefore natural that the deviation from rationality of an agent could benefit her through the impact this distortion has on the price. We find that this result still holds in an economy populated by an infinity of investors where the price impact of each agent is null. Our result also echoes De Castro and Yannelis (2013) who, looking at investors with asymmetric information, show that it might

be beneficial for investors to manipulate the information they process in order to maximize their utility. Hence, they find that being biased and deviating from rationality might be rewarding in terms of utility. Lastly, similarly to Detemple and Murthy (1994), our economy is characterized by waves of optimism and pessimism, as we observe that the aggregate consumption share of optimistic (resp. pessimistic) agents increases in good (resp. bad) states of the world.

Related literature A growing number of papers has been interested in the study of the different types of heterogeneity. More specifically, an important stream of the literature focuses on belief heterogeneity. Most of these papers consider a model populated by two or a finite number of investors who differ only in their belief and analyze the equilibrium properties of such economies (see, e.g., Basak, 2005, Jouini and Napp, 2007, Won et al., 2008). The time preference rate heterogeneity has received less attention. Becker and Mulligan (1997) explain how one investor's time preferences are endogenously determined. Gollier and Zeckhauser (2005) study an economy whose consumers have different constant discount rates and derive implications considering optimal allocations. They show that the representative agent of this economy has a decreasing discount rate. Finally, other papers study simultaneously several types of heterogeneity without considering a potential correlation between them (see, e.g., Cvitanic et al., 2012, Bhamra and Uppal, 2014).

Our paper adds to this literature. While most models deal with a finite number of agents, we consider a continuum of investors. This allows us to consider all possible beliefs and time preference rates. On the technical side, this also allows us to use statistical distributions to describe heterogeneity and, therefore, to characterize it with a limited number of parameters. This methodology derives from Cvitanic et al. (2012) and is similar to the one of Atmaz and Basak (2018), who are the closest to our work.³ They consider a continuum of heterogeneous investors, who differ only in belief, and express the main characteristics of the economy

³For other studies looking at a continuum of agents see, e.g., Chan and Kogan (2002) and Xiouros and Zapatero (2010) who look at economies with a continuum of agents with heterogeneous risk aversions.

using the parameters of the agents' distribution. We differ in three main ways. First, our model has intermediate consumption, which allows us to address interest rates issues. Second, we consider an economy with two types of potentially correlated heterogeneities. The distribution we use to characterize the continuum of investors is therefore a bivariate one, and we are able to study the combined effects of these two heterogeneities. Third, we choose to use logarithmic utility functions, as we know that this type of utility function enables to separate the role of time preference rate heterogeneity and of belief heterogeneity as long as they are independent.

The paper is organized as follows. Section 1.2 presents the model. We determine the equilibrium of our economy and the characteristics of the representative agent in Section 1.3. Section 1.4 presents some market characteristics. Section 1.5 reviews the survival and utility-maximizing issues and Section 1.6 concludes. All proofs are reported in Appendix 1.A and Appendix 1.B contains useful computations and some additional results.

1.2 The model

In a continuous-time framework, we consider a pure exchange Arrow-Debreu economy with a single non-storable consumption good—which we use as numeraire—and a continuum of risk-averse agents who maximize their expected utility for future consumption.

Uncertainty is modeled as usual by a filtered probability space $(\Omega, F, (F_t), P)$, where Ω is the set of states of nature, F is the σ -algebra of observable events, (F_t) describes how information is revealed through time, and P is the (objective) probability measure giving the likelihood of occurrence of the different events in F .

The aggregate endowment process in the economy is denoted by e^* , and we assume that it follows the following stochastic differential equation

$$de_t^* = \mu e_t^* dt + \sigma e_t^* dW_t,$$

where W is a standard unidimensional $((F_t), P)$ -Brownian motion and $(\mu, \sigma) \in \mathbb{R} \times \mathbb{R}_+$ are given constants. Stated differently, we make the assumption that e^* is a geometric Brownian motion with a drift coefficient μ and a volatility σ .

We consider infinitely-lived agents who consume at each date and who all have the same logarithmic instantaneous utility function u , such that $u(x) = \ln(x)$. However, we make the assumption that each agent is characterized by both a subjective belief Q_δ —associated to δ —and a time preference rate ρ .⁴ We call Agent (δ, ρ) the agent endowed with the subjective belief Q_δ , which is assumed to be equivalent to P and which gives the subjective likelihood of occurrence of the different events perceived by this agent, and the time preference rate ρ .⁵ We call Group $(\delta, .)$ the group of agents who share the same belief Q_δ but differ in their time preference rates and, similarly, we call Group $(., \rho)$ the group of agents who share the same time preference rate ρ but differ in their beliefs.

There are therefore two types of heterogeneity in the economy we study.

First, as all probabilities are equivalent, the agents agree on the volatility of the aggregate endowment (for a study of such type of disagreement see, e.g., Duchin and Levy, 2010) but disagree on their estimation of its drift.⁶ All the agents of Group $(\delta, .)$ believe that the aggregate endowment growth rate is given by $\mu_\delta = \mu + \sigma\delta$, and their bias towards optimism or pessimism is thus given by $\sigma\delta$. Hence, if $\delta > 0$, we have $\mu_\delta > \mu$, and they are therefore considered as optimistic agents. Conversely, if $\delta < 0$, they are considered as pessimistic ones. In the case where $\delta = 0$, we consider them as rational agents. We denote by M_δ the density of Q_δ with respect to P , i.e., $\frac{dQ_\delta}{dP} = M_\delta$. From their point of view, the aggregate endowment

⁴In the remainder of the analysis, we interchangeably use Q_δ or δ to refer to the belief of a given agent.

⁵If several agents have the same belief about the future and the same time preference rate, it is easy to check that their aggregate behavior is the same as the behavior of a single agent who has the same characteristics and whose initial endowment is equal to their aggregate endowment. We may then consider them just as one agent.

⁶The equivalence of all probabilities is a necessary condition for the existence of equilibrium. Otherwise, agents who do not believe in the possibility of a given state would want to sell an infinite quantity of the Arrow-Debreu asset associated to that state if its price is positive, and agents who do believe in the possibility of that state would want to buy an infinite quantity of that asset if its price is zero. This is in line with the result of Merton (1980) that the expected return is harder to estimate than the variance.

process follows the following stochastic differential equation

$$de_t^* = (\mu + \sigma\delta) e_t^* dt + \sigma e_t^* dW_{\delta,t},$$

where W_δ is a standard unidimensional $((\mathcal{F}_t), Q_\delta)$ -Brownian motion, such that $W_{\delta,t} = W_t - \delta t$, and, by Girsanov, we have $dM_{\delta,t} = \delta M_{\delta,t} dW_t$.

Second, we allow some agents to be more patient than others: the higher her time preference rate ρ is, the more impatient is the agent. Indeed, for a high ρ , an agent discounts more her future consumption and is therefore more willing to consume quickly. Conversely, for a small ρ , she is more willing to save for future consumption as she does not discount it much.

Hence, Agent (δ, ρ) aims at maximizing her von Neumann Morgenstern utility for future consumption of the form

$$\mathbb{E}^{Q_\delta} \left[\int_0^\infty \exp(-\rho t) \ln(c_{\delta,\rho,t}^*) dt \right] = \mathbb{E} \left[\int_0^\infty \exp(-\rho t) M_{\delta,t} \ln(c_{\delta,\rho,t}^*) dt \right],$$

where $c_{\delta,\rho}^*$ is her consumption stream.

Finally, as we consider a continuum of agents, we use a probability density function to describe their initial wealth share distribution.⁷ This distribution—given exogenously at $t = 0$ —depends on a parameter $k \in \mathbb{R}$ that allows the two types of heterogeneity to be correlated. For instance, if we assume a negative correlation between them, we have that the most optimistic agents are also the most patient ones. We define $\nu_{\delta,\rho,k}$ as the share of total initial endowment owned by Agent (δ, ρ) at $t = 0$. Formally, we assume

$$\nu_{\delta,\rho,k} = \frac{1}{\sqrt{2\pi\bar{\omega}}} \exp\left(-\frac{(\delta + k\rho - \delta_0)^2}{2\bar{\omega}^2}\right) \frac{\vartheta^l}{\Gamma(l)} \rho^{l-1} \exp(-\rho\vartheta).$$

⁷As already underlined, we may aggregate the agents who have the same characteristics (δ, ρ) . Therefore, the distribution of the individuals across the different belief and time preference characteristics is not relevant by itself, and the relevant information is rather the distribution of initial endowment among the set of individual characteristics.

For a given time preference rate, the beliefs are distributed according to a Gaussian distribution. This is as in Atmaz and Basak (2018). Moreover, even though we take the heterogeneity in beliefs as given ex-ante and assume no learning, this Gaussian assumption is consistent with models where investors have different private information coming from white noises (see, e.g., Kyle, 1985). Similarly, for a given belief, we use a Gamma distribution to describe the heterogeneity in time preference rates. Weitzman (2001) uses a similar distribution and provides empirical evidence that supports this choice.

For computational reasons, we assume throughout the paper that l equals two. The other parameters are chosen exogenously at time $t = 0$ before equilibrium is reached.

When $k = 0$, there is no correlation between the two types of heterogeneity, and the density can therefore be decomposed into two independent components: a Gaussian density with a mean δ_0 and a standard deviation $\bar{\omega}$, that describes the initial belief heterogeneity, and a Gamma density with a shape parameter $l = 2$ and a rate parameter ϑ , that describes the initial heterogeneity in time preference rate. In this situation, the initial average bias towards optimism or pessimism of the agents is given by $\sigma\delta_0$ and, if $\delta_0 > 0$, we say that the economy has an (exogenous initial) optimistic bias.⁸ The belief dispersion is given by the standard deviation $\bar{\omega}$, the average time preference rate is given by $\frac{2}{\vartheta}$, and the dispersion of the rates of the agents is given by the standard deviation $\frac{\sqrt{2}}{\vartheta}$.

In the general case, the correlation between the belief and the time preference rate is given by a function of k . In particular, at $t = 0$, the correlation function is given by

$$\frac{-\sqrt{2}k}{\sqrt{\bar{\omega}^2\vartheta^2 + 2k^2}}.$$

In order to express this function and the main results of our model as functions of the central (co)moments of δ and ρ at each date, we define the following time-dependent

⁸Note that not imposing $\delta_0 = 0$ implies that we do not necessarily require that the agents are symmetric around rationality (ex-ante).

stochastic density function

$$\tilde{\nu}_{\delta,\rho,k,t} = \frac{\nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t}}{\int \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho},$$

and we denote by $\mathbb{E}_t(\cdot)$ the time-dependent mean with weights given by $\tilde{\nu}_{\delta,\rho,k,t}$. Using some algebra, we derive that the correlation function is an inverted S-shaped function such that the correlation is null when k equals zero and such that, for all time t , its sign is opposite to the sign of k .⁹

Finally, let us introduce the following notations that will be useful in the next. We denote respectively by φ and Φ the density and cumulative distribution functions of the standard normal distribution, i.e., $\varphi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right)$ and $\Phi(x) = \int_{-\infty}^x \varphi(s) ds$. We denote by sgn the sign function, i.e., $sgn(x) = 1$ if $x > 0$, $sgn(x) = -1$ if $x < 0$, and $sgn(0) = 0$.

1.3 Representative agent and equilibrium

In this section, we study the Arrow-Debreu equilibrium of this economy and the characteristics of the representative agent.

In such a model, an Arrow-Debreu equilibrium is defined by a positive density price p^* and a continuum of consumption plans $(c_{\delta,\rho}^*)_{\delta \in \mathbb{R}, \rho \in \mathbb{R}_+^*}$, each one maximizing the von Neumann Morgenstern utility for future consumption of the corresponding agent under her budget constraint and such that the market clears.

The representative agent of this economy is a fictitious agent who, if endowed with the total wealth of the economy, would have a marginal utility equal to the equilibrium price. This agent is therefore obtained by construction from the economy characteristics. In the proof of Proposition 1.1, we construct such an agent, and we denote by \bar{Q} her belief—associated to $\bar{\delta}_t$ —, by \bar{M} the density of \bar{Q} with respect to P , and by $\bar{\rho}_t$ her time preference

⁹Explicit computations of the correlation function and of the other central (co)moments are reported in Appendix 1.B.1.

rate. Let emphasize that, unlike common investors in the economy, the representative agent's characteristics $\bar{\delta}_t$ and $\bar{\rho}_t$ can be time- and state-dependent. This is because, by construction, the representative agent results from the aggregation of common investors and because it appears that her characteristics are given by consumption-weighted averages of the individual agents' characteristics (with time- and state-dependent consumptions). This result is also supported by prior works (see, e.g., Gollier and Zeckhauser, 2005).

Proposition 1.1. *1. At the equilibrium, the state price density and the consumption plans are given by*

$$p_t^* = (e_t^*)^{-1} \int \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho,$$

$$c_{\delta,\rho,t}^* = (p_t^*)^{-1} \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t},$$

where $\lambda_{\delta,\rho,k}$ is defined such that $\lambda_{\delta,\rho,k} = \rho \nu_{\delta,\rho,k} \left(\int \frac{\lambda_{\delta,\rho,k}}{\rho} d\delta d\rho \right)$.

2. The representative agent's time preference rate is a time-dependent consumption-weighted average of the individual time preference rates and is given by $\bar{\rho}_t = \frac{\mathbb{E}_t(\rho^2)}{\mathbb{E}_t(\rho)}$, and the associated variance is given by $\sigma_{\bar{\rho}} = \frac{\mathbb{E}_t(\rho^3)}{\mathbb{E}_t(\rho)} - \bar{\rho}_t^2$.
3. The representative agent's belief is a time-dependent consumption-weighted average of the individual beliefs and is given by $\bar{\delta}_t = \frac{\mathbb{E}_t(\delta\rho)}{\mathbb{E}_t(\rho)}$, and the associated variance is given by $\sigma_{\bar{\delta}} = \frac{\mathbb{E}_t(\delta^2\rho)}{\mathbb{E}_t(\rho)} - \bar{\delta}_t^2$.

When $k \neq 0$, explicit computations of $\bar{\rho}_t$ and $\bar{\delta}_t$ lead to

$$\bar{\rho}_t = \frac{\sqrt{1 + \bar{\omega}^2}}{|k|\sqrt{t}} \Psi(X_t), \quad \bar{\delta}_t = \frac{\delta_0 + \bar{\omega}^2 W_t}{1 + t\bar{\omega}^2} - \frac{\text{sgn}(k)}{\sqrt{t}\sqrt{1 + t\bar{\omega}^2}} \Psi(X_t),$$

with

$$X_t = \frac{\sqrt{1 + t\bar{\omega}^2}}{|k|\sqrt{t}} \left(t + \vartheta + k \frac{W_t - \delta_0 t}{1 + t\bar{\omega}^2} \right), \quad (1.1)$$

and

$$\Psi(x) = \frac{2 + x^2 - \frac{1-\Phi(x)}{\varphi(x)}(3x + x^3)}{-x + \frac{1-\Phi(x)}{\varphi(x)}(1 + x^2)}.$$

We easily get that $\Psi(X_t)$ is non-negative and converges to zero when t goes to infinity. Hence, we see that the belief of the representative agent is more optimistic when the correlation is positive and more pessimistic when the correlation is negative. In other words, if the most optimistic agents are also the most impatient (resp. patient) ones, the representative agent is more optimistic (resp. pessimistic). Looking at the asymptotic behavior of her belief, we also derive that this agent tends to be the rational one. Similarly, we have that, unlike the other agents, the time preference rate of the representative agent is not a constant and goes to zero asymptotically. Hence, in the long run, the representative agent of the economy is rational and more patient than all the agents, which is consistent with the survival implications derived in Section 1.5.

In the uncorrelated case, the characteristics of the representative agent are given by $\bar{\rho}_t = \frac{l+1}{\vartheta+t}$ and $\bar{\delta}_t = \frac{\delta_0 + \bar{\omega}^2 W_t}{1 + t\bar{\omega}^2}$, with the associated variances respectively given by $\sigma_{\bar{\rho}} = \frac{l+1}{(\vartheta+t)^2}$ and $\sigma_{\bar{\delta}} = \frac{\bar{\omega}^2}{1 + t\bar{\omega}^2}$. In this case, while $\frac{l}{\vartheta}$ and δ_0 (resp. $\frac{l}{\vartheta^2}$ and $\bar{\omega}^2$) measure the time preference rate and belief averages (resp. variances) with weights given by the agents' initial total endowment, $\bar{\rho}_t$ and $\bar{\delta}_t$ (resp. $\sigma_{\bar{\rho}}$ and $\sigma_{\bar{\delta}}$) measure their averages (resp. variances) with weights given by agents' current consumption. Consistent with Gollier and Zeckhauser (2005), who study an economy where agents differ only in time preferences, we derive that the time preference rate of the representative agent decreases with time. Moreover, we observe that, for a given t , both variances are non-negative constants and do not depend on W . As underlined by Atmaz and Basak (2018), this is due to the assumption of an unbounded investor type space. Indeed, in the case of bounded beliefs, they argue that the wealth transfer accumulates to one type of investor and that the belief dispersion therefore goes to zero in extreme states. This is not the case in the presence of a continuum of heterogeneous investors. Adding heterogeneity in time preference rates, we see that a similar reasoning holds for $\sigma_{\bar{\rho}}$.

1.4 Market characteristics

In this section, we derive several characteristics of the market and study how the correlated heterogeneities impact them.

1.4.1 Risk-free rate and market price of risk

Let first recall that, in the standard homogeneous case, the risk-free rate and the market price of risk are both time- and state-independent and given by

$$\begin{aligned} r^f(\text{std}) &= \mu - \sigma^2 + \hat{\rho}, \\ MPR(\text{std}) &= \sigma, \end{aligned}$$

where $\hat{\rho} > 0$ stands for the homogeneous time preference rate that all the agents agree on.

We have the following result.

Proposition 1.2. *In our economy, the risk-free rate and the market price of risk are given by*

$$\begin{aligned} r_t^f &= \mu - \sigma^2 + \bar{\rho}_t + \sigma \bar{\delta}_t, \\ MPR_t &= \sigma - \bar{\delta}_t. \end{aligned}$$

As already shown in the literature (see, e.g., Bhamra and Uppal, 2014), we observe that, unlike the standard case, they are not time- and state-independent. We complement the previous literature by extending the result to the case of a continuum of agents, which allows us to express the risk-free rate and the market price of risk as functions of the moments of the heterogeneity distribution. Moreover, we are able to consider the impact of the correlation on such quantities.

Let first study the uncorrelated case. We have

$$r_t^f = \mu - \sigma^2 + \frac{l+1}{\vartheta+t} + \sigma \frac{\delta_0 + \bar{\omega}^2 W_t}{1 + t\bar{\omega}^2},$$

$$MPR_t = \sigma - \frac{\delta_0 + \bar{\omega}^2 W_t}{1 + t\bar{\omega}^2}.$$

First, we notice that the belief heterogeneity affects both market characteristics. If on average the agents have a pessimistic bias ($\delta_0 < 0$), the risk-free rate is lower and the market price of risk is higher. This echoes the result of Bhamra and Uppal (2014), who show, in an economy populated by two agents, that the former (resp. latter) depends positively (resp. negatively) on the weighted arithmetic mean of the beliefs of individuals agents. This is intuitive as a pessimistic economy on average rewards agents who take risks more than an economy where the agents are optimistic on average and therefore willing to take more risks. Second, we see that the time preference rate heterogeneity only impacts the risk-free rate: the higher is the average time preference rate, or, stated differently, the more impatient the agents are on average, the higher the risk-free rate of the economy is.

Concerning the state dependency, we see that, for a given t , as r^f is a linear increasing function of W and MPR a linear decreasing one, the risk-free rate is procyclical and the market price of risk countercyclical. This latter result is consistent with the empirical observation of Campbell and Cochrane (1999), who say that the equity premium seems to be smaller when the economy is doing well, and with the theoretical implications of Jouini and Napp (2011). We complement these findings as we observe that the higher the variance of the belief distribution, the higher the procyclical effect on the risk-free rate and the countercyclical effect on the market price of risk.

In the general case, we notice that, as t tends to infinity, the median risk-free rate tends to $\mu - \sigma^2$. Hence, asymptotically, the risk-free rate of the economy tends to the one of an economy populated by homogeneous investors having a zero time preference rate. This result is consistent with the survival implications of the model that are presented in the next

section: when time goes to infinity, as the agents who survive are the most patient ones, the share of patient agents in the economy becomes larger. As those agents are characterized by almost null time preference rates, it leads to a lower risk-free rate. Similarly, we observe that MPR tends to σ . We derive that asymptotically the effects of the belief heterogeneity on the market price of risk vanish.

Finally, we study the additional impact of the correlation between beliefs and time preference rates on the risk-free rate and the market price of risk, and we see that this impact depends on the sign of the correlation between the two types of heterogeneity. In particular, a negative correlation ($k > 0$) leads to a lower risk-free rate and a higher market price of risk. Thereby, this novel effect, induced by the correlated heterogeneities, helps to solve, at least partly, the risk-free rate and the market price of risk puzzles.

Insert Figure 1.1 here.

We illustrate this correlation effect for the median risk-free rate in Figure 1.1 in which we plot its time evolution, using standard parameters to define the initial wealth distribution of the investors and the economy process. Formally, we set $\mu = 14.23\%$, $\sigma = 8.25\%$, $\delta_0 = 0$, $\bar{w} = 3.39\%$, $l = 2$, and $\vartheta = 50$.¹⁰ The parameters of the economy process and of the belief heterogeneity distribution are the same as in Atmaz and Basak (2018). In particular, we assume that there is no aggregate belief bias. In the uncorrelated case, the parameters of the time preferences distribution implies an initial average time preference rate of 4% and a standard deviation of 2.83%. These are similar values to those used in Weitzman (2001).¹¹ We also set k to match given levels of median correlation at $t = 1$. We study five economies: a strongly negatively correlated one (with a median correlation of -0.75 at $t = 1$), a negatively correlated one (-0.25), an uncorrelated one, a positively correlated one (0.25), and a strongly positively correlated one (0.75). We observe that, for all t , the higher

¹⁰We use the same set of parameters throughout the paper.

¹¹Asking 2160 economists about the right time preference rate to use in practice, he concludes that the average time preference rate is 4% with a standard deviation of 3%.

the heterogeneity correlation is, the higher is the median risk-free rate, and that this effect decreases with time.

1.4.2 Market volatility

We now consider the asset S , whose dividend process is given by the total endowment of the economy. The value of S is known and given by

$$S_t = \frac{\mathbb{E}_t \left[\int_t^\infty p_s^* e_s^* ds \right]}{p_t^*} = e_t^* \frac{\int \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}{\int \rho \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}.$$

Using Ito's Lemma, we identify the volatility parameter σ_S of the stochastic differential equation followed by S , which corresponds to the market volatility.

Proposition 1.3. *In our economy, the market volatility is given by*

$$\sigma_{S,t} = \sigma - \frac{\text{cov}_t(\delta, \rho)}{\mathbb{E}_t(\rho)} = \sigma + \frac{\text{sgn}(k)}{\sqrt{t}\sqrt{1+t\bar{\omega}^2}} \Upsilon(X_t),$$

with X_t given in Equation (1.1) and

$$\Upsilon(x) = \frac{2 - 3\frac{1-\Phi(x)}{\varphi(x)}x + \left(\frac{1-\Phi(x)}{\varphi(x)}\right)^2(-1+x^2)}{-x + \frac{1-\Phi(x)}{\varphi(x)}(1+2x^2) - \left(\frac{1-\Phi(x)}{\varphi(x)}\right)^2(x+x^3)}.$$

A first important point to notice is that, for $k = 0$, we obtain $\sigma_{S,t} = \sigma$. Hence, we derive that when the two types of heterogeneity are uncorrelated, there is no heterogeneity impact on the market volatility.

When we allow some correlation to exist between the belief heterogeneity and the time preference one, we observe that the market volatility becomes time- and state-dependent. The initial impact of the two correlated types of heterogeneity is easy to derive and given

by

$$\sigma_{S,0} = \sigma + \frac{k}{\vartheta}.$$

For general t , as the mean of the time preference rates is positive, there is some excess volatility induced by the two correlated heterogeneities when the covariance between them is negative. Concretely, we derive that, when the most optimistic agents are also the most patient ones, the market volatility increases. Conversely, when the covariance is positive, the presence of heterogeneous investors decreases the market volatility. This is consistent with Li (2007), who derives a similar result in an economy populated by a patient and an impatient agent. More precisely, he finds that $\sigma_S > \sigma$ (resp. $<$, $=$) when the patient agent is more (resp. less, as) optimistic than the impatient agent. We therefore complement this result by showing that it still holds when there is an infinite number of agents, and we relate this effect to the parameters of the statistical distribution of the agents' characteristics.

Insert Figure 1.2 here.

Figure 1.2 shows the evolution of the median market volatility over time and illustrates this volatility effect. We notice that this effect is stronger when the absolute value of the correlation between the two types of heterogeneity is higher.

Looking at the asymptotic behavior of Υ , we find that the function converges to zero, and we thus derive that the market volatility tends to its standard value when time goes to infinity. Hence, we conclude that the volatility effect tends to vanish. However, we note that it takes a substantial amount of time for this effect to disappear, as we see that, even after 100 years, there is still a large effect for the most correlated economies considered.

1.4.3 Trading volume

We now analyze how the trading volume is impacted by the presence of heterogeneous investors. Note that, in a standard homogeneous economy, the agents make their decisions based on the same belief and the same time preference rate. Hence, they all act identically

and there is no trading as they can not find other agents who are willing to trade with them. This is not the case when there is some heterogeneity among the agents.

To compute the trading volume of Agent (δ, ρ) , we first compute explicitly her total wealth, given by

$$V_{\delta,\rho,t} = \frac{\mathbb{E}_t \left[\int_t^\infty p_s^* c_{\delta,\rho,s}^* ds \right]}{p_t^*}.$$

Using Ito's Lemma, we then derive the stochastic differential equation this total wealth follows. We find $dV_{\delta,\rho,t} = \mu_{V_{\delta,\rho,t}} V_{\delta,\rho,t} dt + \sigma_{V_{\delta,\rho,t}} V_{\delta,\rho,t} dW_t$, with

$$\begin{aligned} \mu_{V_{\delta,\rho,t}} &= \mu - \rho + \sigma\delta + \bar{\rho}_t - (\sigma + \delta)\bar{\delta}_t + \bar{\delta}_t^2, \\ \sigma_{V_{\delta,\rho,t}} &= \sigma + \delta - \bar{\delta}_t. \end{aligned}$$

We also know that the total wealth of Agent (δ, ρ) can be decomposed into three components: the number of shares α she invests in the risky asset S described in the previous subsection, the number of shares β she invests in the risk-free asset B whose drift is given by the risk-free rate, and what she consumes.

Hence, we have

$$\begin{aligned} dV_{\delta,\rho,t} &= \alpha_{\delta,\rho,t} dS_t + \beta_{\delta,\rho,t} dB_t - c_{\delta,\rho,t}^* dt \\ &= \left(\alpha_{\delta,\rho,t} \mu_{S,t} S_t + \beta_{\delta,\rho,t} r_t^f B_t - c_{\delta,\rho,t}^* \right) dt + \alpha_{\delta,\rho,t} \sigma_{S,t} S_t dW_t. \end{aligned}$$

By identification, we derive the number of shares Agent (δ, ρ) should optimally invest in the risky asset and in the risk-free asset. In particular, we obtain that the optimal number of shares Agent (δ, ρ) should invest in the risky asset is given by $\alpha_{\delta,\rho,t} = \frac{V_{\delta,\rho,t} \sigma_{V_{\delta,\rho,t}}}{S_t \sigma_{S,t}}$.

Using a similar approach as in recent studies in continuous time (see, e.g., Xiong and Yan, 2010, Longstaff and Wang, 2012), we define the trading volume of Agent (δ, ρ) as the absolute value of the volatility in the stochastic differential equation satisfied by $\alpha_{\delta,\rho}$.

Proposition 1.4. *The trading volume of Agent (δ, ρ) at time t in state of the world W_t is given by*

$$\frac{V_{\delta, \rho, t}}{S_t} \frac{1}{\sigma_{S, t}} \left| \left((\sigma + \delta - \bar{\delta}_t) \left(\delta - \mathbb{E}_t(\delta) + \frac{1}{\sigma_{S, t}} \left(\frac{\text{cov}_t(\delta^2, \rho)}{\mathbb{E}_t(\rho)} - \frac{\text{cov}_t(\delta, \rho) (\mathbb{E}_t(\delta) + \bar{\delta}_t)}{\mathbb{E}_t(\rho)} \right) \right) - \sigma_{\bar{\delta}} \right) \right|.$$

We see that the trading volume of Agent (δ, ρ) is time- and state-dependent. It also depends on the belief and the time preference rate of the agent being considered. In unreported graphs, looking at the trading volumes of some agents sharing the median time preference rate but having different beliefs, we do not find any clear link between the trading volume and the belief bias. However, we observe a negative relation between the time preference rate of an investor and her trading volume, and we illustrate this result in Figure 1.3.

Insert Figure 1.3 here.

We plot the evolution over time of the median trading volume of three unbiased investors with different time preferences. The different time preference rates are chosen to partition the Gamma distribution used to describe the initial time preferences heterogeneity when there is no correlation between the two types of heterogeneity. The time preference rate of the first (resp. second, third) investor is the 10th (resp. 50th, 90th) percentile of this distribution. In other words, this agent is more patient than 90% (resp. 50%, 10%) of the population and is therefore called the patient (resp. neutral, impatient) agent. Panel A shows the uncorrelated case, while Panel B (resp. Panel C) shows the evolution of their trading volume in an economy where the median correlation at $t = 1$ is -0.75 (resp. 0.75).

We observe the same pattern in all panels: the more patient an agent is, the more she trades. Observing that the trading volume of the patient agent is higher in Panel B and Panel C than in Panel A, we deduce that a higher absolute value of the correlation between beliefs and time preference rates seems to strengthen this effect. Similarly, we observe that the trading volume of the impatient agent is low in the uncorrelated economy and almost equals zero in the correlated ones. Even if trades implying this agent may occur, we derive

that, being impatient, she wants to sell her shares independently of the state of the world. Hence, as her trading volume is defined as the volatility of her optimal portfolio, it decreases quickly.

To get a deeper insight of the mechanism behind this result linking trading volume and time preference rate, let us focus on the uncorrelated case. When $k = 0$, the formula of Proposition 1.4 simplifies to

$$|\sigma_{\alpha_{\delta,\rho,t}}| = \frac{V_{\delta,\rho,t}}{S_t} \frac{(\sigma + \delta - \bar{\delta}_t)(\delta - \bar{\delta}_t) - \sigma\bar{\delta}}{\sigma}.$$

Hence, we derive that more patient agents trade more due to a wealth effect. As suggested by the survival implications of the model—derived in the next section—, it takes more time for more patient agents to be driven out of the market, and they therefore benefit more from the economy growth, which allows them to trade more. Consistent with this intuition, we observe that the trading volume is procyclical.

The overall trading volume of the economy can further be obtained by summing over the agents' trading volumes and dividing this quantity by two to prevent double summation of the shares traded across investors. Formally, we have

$$Vol_t = \frac{1}{2} \int |\sigma_{\alpha_{\delta,\rho,t}}| d\delta d\rho.$$

After some algebra, we derive that, when there is no correlation between the belief and the time preferences heterogeneities, it is given by the following formula.

Proposition 1.5. *In the uncorrelated economy, the trading volume at time t in state of the world W_t is given by*

$$Vol_t = \frac{\sqrt{\sigma\bar{\delta}}}{2\sigma} \left(\left(\sigma + \sqrt{\sigma^2 + 4\sigma\bar{\delta}} \right) \varphi \left(\frac{-\sigma + \sqrt{\sigma^2 + 4\sigma\bar{\delta}}}{2\sqrt{\sigma\bar{\delta}}} \right) - \left(\sigma - \sqrt{\sigma^2 + 4\sigma\bar{\delta}} \right) \varphi \left(\frac{-\sigma - \sqrt{\sigma^2 + 4\sigma\bar{\delta}}}{2\sqrt{\sigma\bar{\delta}}} \right) \right).$$

The trading volume is state-independent and, because some agents are progressively driven out of the market, decreases with time. This result extends Atmaz and Basak (2018) to heterogeneous time preference rates. In fact, as explained in Section 1.3, the unboundedness of the investor type space implies that $\sigma_{\bar{\delta}}$ is constant with respect to W and, therefore, trades occur independently of the state of the world. As in Atmaz and Basak (2018), we also observe that, for t being fixed, the trading volume increases with the variance associated to the belief bias of the representative agent and, consequently, with the dispersion coefficient of the belief heterogeneity distribution. This is because a higher belief dispersion in the economy means that the investors are more heterogeneous and that there is therefore more agents willing to trade.

1.5 Consumption shares and utility-maximizing agents

In this section, we only consider the case where the two types of heterogeneity are uncorrelated. Looking at the consumption shares, we characterize the agents who survive in the long run and those who dominate the market depending on the state of the world. We then look at the agents who maximize their ex-post utility and their ex-ante one.¹²

1.5.1 Surviving and market-dominating agents

In this subsection, we focus on the consumption shares of the agents. Letting $\tau_{\delta,\rho}$ denote the consumption share of Agent (δ, ρ) , we have

$$\tau_{\delta,\rho,t} = \frac{c_{\delta,\rho,t}^*}{e_t^*} = \frac{\lambda_{\delta,\rho,0} \exp(-\rho t) M_{\delta,t}}{\int \lambda_{\delta,\rho,0} \exp(-\rho t) M_{\delta,t} d\delta d\rho}.$$

First, characterizing the surviving agent, we study the survival properties of our economy.

¹²In Appendix 1.B.2, we briefly discuss the impact of correlation on the characteristics of the surviving agent and on the belief of the ex-post utility-maximizing agent. In particular, we find that it does not affect the survival implications of the model and that, in case of a negative correlation, the heterogeneity effect on the ex-post utility-maximizing agent is mitigated.

Formally, we say that an agent survives if her consumption share does not approach zero almost surely when time goes to infinity. We define the survival index of Agent (δ, ρ) as in Yan (2008).¹³ Similarly to his finding, we obtain the following result that is stated only for the sake of analysis completeness.

Proposition 1.6 (Yan). *In the long run, the only surviving agent is the one with the smallest survival index.*

We notice that the survival index depends positively on ρ . Hence, we derive that having a high time preference rate is a disadvantage for survival. Stated differently, the more impatient an agent is, the less likely she is to survive in the long run. This result is intuitive as an impatient agent, discounting her future utility more, prefers to consume today than to save for future consumption. Conversely, by saving more for future consumption, a more patient agent enhances her chances to survive. Equivalently, observing that the survival index is an increasing function of the absolute value of δ , we obtain that the lower the belief bias of an agent is, the better it is for her survival. Finally, as Agent (δ, ρ) and Agent $(-\delta, \rho)$ have the same survival index, having a bias towards optimism or pessimism is equally disadvantageous regarding survival issues.

Let us also study how fast an agent disappears from the market. To do so, we compute the average half-life of several agents and compare them. We define the half-life of an agent as the time taken for her current endowment to fall to half of her initial endowment. Formally, the half-life $t_{\delta, \rho}^{HL}$ of Agent (δ, ρ) is given by $t_{\delta, \rho}^{HL} = \{\inf t \text{ such that } \tau_{\delta, \rho, t} = \frac{\tau_{\delta, \rho, 0}}{2}\}$.

Insert Table 1.1 here.

We notice that the half-life is stochastic. Hence, we report the average half-lives in Table 1.1 and consider three different time preference rates and five different beliefs that

¹³Formally, it is given by $\rho + \frac{\delta^2}{2}$. We do not include the third component of Yan (2008)'s survival index relative to the risk aversion coefficient heterogeneity as we assume that all the agents have the same logarithmic utility function.

partition the initial wealth share distribution of the continuum of agents. The time preference rates are defined as in Section 1.4.3 and we therefore compare patient, neutral and impatient agents. The first (resp. second, third, fourth, and fifth) belief bias is such that 10% (resp. 25%, 50%, 75%, and 90%) of the agents are more pessimistic than the agents endowed with this belief bias. Using the same set of parameters as in the previous section, we therefore compute the average half-life of 15 specific agents and see that, without being the smallest, the smaller the survival index of an agent, the longer she survives.¹⁴

Finally, aggregating the agents into groups of agents sharing the same belief, we define formally the surviving group of this economy and its associated belief. To do so, we integrate the consumption shares of the agents who share the same belief with respect to the time preference rate and study their evolution. Formally, we compare for different δ the consumption shares of the Groups (δ, \cdot) —that we denote τ_δ —and study their limit when time goes to infinity. Easy computations give us that $\forall \delta \in \mathbb{R}^*$, $\lim_{t \rightarrow +\infty} \tau_{\delta,t} = 0$. This result means that the only group of agents who survive is Group $(0, \cdot)$, or, in other words, that the surviving group is the rational one, and that the surviving agent of the economy is the most patient agent in this group.

We now turn to the study of the market-dominating agents. More precisely, for a given t , we now analyze how the consumption shares of the agents evolve given the states of the world and study, for very good and very bad ones, which category of agents dominates the market, in the sense that their aggregate consumption share approaches one.

Proposition 1.7. *The aggregate consumption share of the optimistic (resp. pessimistic) agents τ_{opt} (resp. τ_{pes}) are given by*

$$\tau_{opt,t} = 1 - \Phi\left(-\frac{\bar{\delta}_t}{\sqrt{\sigma_\delta}}\right), \quad \tau_{pes,t} = \Phi\left(-\frac{\bar{\delta}_t}{\sqrt{\sigma_\delta}}\right),$$

¹⁴Note that the average half-lives are relatively long. Hence, we derive that the asymptotic results of the paper are true asymptotic results and not only approximations of a stationary state. Note also that these relatively long half-lives imply that non-surviving agents can potentially have a significant impact on the market over time.

and the optimistic (resp. pessimistic) agents dominate the market (resp. become extinct) for very good states of the world, i.e., when $W_t \rightarrow +\infty$, and become extinct (resp. dominate the market) for very bad ones, i.e., when $W_t \rightarrow -\infty$.

Notice that τ_{opt} is a monotonically increasing function of W and τ_{pes} a monotonically decreasing one. Hence, as in Jouini and Napp (2011), we derive that our economy is characterized by waves of optimism and waves of pessimism, in the sense that the consumption shares are biased in favor of the optimistic agents in the good states of the world and in favor of the pessimistic agents in the bad states of the world. Note that these waves of optimism and pessimism also depend on the initial average economy bias. For instance, when the agents are optimistic on average ($\delta_0 > 0$), the bias towards the consumption shares of the optimistic agents is even more important in good states of the world. Conversely, this wave of optimism is smaller if, on average, the agents are pessimistic ($\delta_0 < 0$).

Insert Figure 1.4 here.

Another way of showing the existence of such waves of optimism and pessimism is to show how the aggregate consumption share of all the optimistic agents evolves over time depending on the state of the economy (or, equivalently, how the aggregate consumption share of all the pessimistic agents evolves, as their sum adds to one). To do so, we consider three different trajectories and study the time evolution of τ_{opt} in each of them in Figure 1.4. The first trajectory is characterized by a series of positive events. Formally, we assume that $W_t = \sqrt{t}$ at each date. Similarly, we study a trajectory where bad events happen consistently ($W_t = -\sqrt{t}$ at each date) and one where neutral events happen consistently ($W_t = 0$ at each date). We see that, in the good trajectory, the proportion of optimistic investors tends to one when time goes to infinity. Similarly, in the bad trajectory, the economy tends to be populated by pessimistic investors only. We conclude that consistent good events result in an economy populated by investors who think that a good event is more likely to happen and vice versa. As the agents considered do not learn, this does not result from learning but from market elimination.

1.5.2 Utility-maximizing agents

We now look at the characteristics of the utility-maximizing agents of the economy. More precisely, for a given Group $(., \rho)$, we determine the belief an agent should have in order to have the highest ex-post (and ex-ante) utility level (with respect to the other agents).

Note that, ex-post, an agent knows which states of the worlds occurred. Hence, since the objective probability P governs the states of the world the agents face during their life, the ex-post utility of Agent (δ, ρ) is given on average by

$$U_{\delta, \rho}^{\text{ex-post}}(c_{\delta, \rho}^*) = \mathbb{E} \left[\int_0^{\infty} \exp(-\rho t) u(c_{\delta, \rho, t}^*) dt \right].$$

Conversely, Agent (δ, ρ) does not know ex-ante which states will be realized in the future. She therefore uses the subjective probability Q_{δ} to compute her ex-ante utility, given on average by

$$U_{\delta, \rho}^{\text{ex-ante}}(c_{\delta, \rho}^*) = \mathbb{E} \left[\int_0^{\infty} \exp(-\rho t) M_{\delta, t} u(c_{\delta, \rho, t}^*) dt \right].$$

We derive the following result.

Proposition 1.8. *1. Agent $(\delta^{\text{ex-post}}(\rho), \rho)$, the ex-post utility-maximizing agent of Group*

$$(\cdot, \rho), \text{ is characterized by } \delta^{\text{ex-post}}(\rho) = \frac{\delta_0}{1 + \frac{\omega^2}{\rho}}.$$

2. Agent $(\delta^{\text{ex-ante}}(\rho), \rho)$, the ex-ante utility-maximizing agent of Group (\cdot, ρ) , has a more optimistic (or less pessimistic) belief than the average agent.

3. In particular, when $\delta_0 \geq 0$, Agent $(\delta^{\text{ex-ante}}(\rho), \rho)$ has a more optimistic belief than the rational agent. When $\delta_0 < 0$, she has a more optimistic belief than the rational agent when σ is high.

4. Unless when $\delta_0 < 0$ and σ is high, there exists biased agents in Group (\cdot, ρ) whose ex-post and ex-ante utilities are higher than the one of the rational agent of this group.

5. In particular, when $\delta_0 \geq 0$, Agent $(\delta^{\text{ex-post}}(\rho), \rho)$ has a higher ex-post and ex-ante utility than Agent $(0, \rho)$.

From the first point of Proposition 1.8, we have that $\delta^{\text{ex-post}}(\rho) = 0$ when $\delta_0 = 0$. Hence, looking at a given group of agents having the same time preference rate, we derive that, if the economy has no aggregate belief bias, the ex-post utility-maximizing agent and the rational one (i.e., the surviving agent of the group) share the same belief and are therefore identical. However, if there exists a bias towards optimism or pessimism in the economy, the two beliefs differ and the agent who maximizes her ex-post utility is not the one who survives in the long run. Hence, as in Jouini and Napp (2016), ex-post, a shorter life might be more rewarding than a longer one. We also notice that the belief bias of the ex-post utility-maximizing agent tends to zero when $\bar{\omega}^2$ tends to infinity. As it denotes the variance of the initial belief distribution, we conclude that the more heterogeneity in belief there is, the more the agent who maximizes her ex-post utility in her group tends to be the one who survives in this group. The intuition behind this result is that the wider the beliefs are spread, the more extremely optimistic or pessimistic agents there are in the economy and, therefore, the quicker the extremely biased decisions these agents take lead them to go extinct. In other words, the more the variance of the belief distribution increases, the less the trade-off between having a more rewarding life based on biased decisions and having a longer life by being rational becomes favorable. We also notice that the initial aggregate belief bias of the economy does not impact $\delta^{\text{ex-post}}(\rho)$ when $\bar{\omega}^2$ goes to infinity. Finally, as $1 + \frac{\bar{\omega}^2}{\rho} > 1$, we know that $\delta^{\text{ex-post}}(\rho)$ is smaller than δ_0 in absolute value and that both have the same sign. We derive that the agent who maximizes her ex-post utility in her group has the same type of bias towards optimism or pessimism as the economy, but she is more rational than the average: in a euphoric economy ($\delta_0 > 0$), the ex-post utility-maximizing agent of each group is a less euphoric, yet optimistic, one and she is a less depressed, yet pessimistic, agent in a depressed economy ($\delta_0 < 0$).

Our setting, that also considers heterogeneous time preference rates, allows us to derive

another new result from this first point. We see that as ρ increases, the ex-post utility-maximizing belief goes from zero to δ_0 . In fact, looking at patient groups of agents, we have that the ex-post utility-maximizing agent of the group tends to be the rational agent and that, for more impatient groups of agents, she tends to be the one endowed with the average belief.

From the second point, we derive that, when there is no aggregate belief bias in the economy, Agent $(\delta^{\text{ex-ante}}(\rho), \rho)$ is more optimistic than the rational agent, who is the one who both survives and maximizes her ex-post utility. The intuition behind this result is that, ex-ante, an agent does not know that her bias will lead her to take wrong decisions. Hence, idealizing the reality by being optimistic allows her to have a higher ex-ante utility. However, ex-post, the wrong decisions she will take leads her life to be shorter and her ex-post utility to be smaller. For similar reasons, when the aggregate economy has a positive belief bias, we have that the ex-ante utility-maximizing agent of a given group of agents is more optimistic than both the surviving one and the ex-post utility-maximizing one of the group. Finally, when the agents are pessimistic on average, the situation is slightly different. Even if she is less pessimistic than the average of the agents, we do not know if Agent $(\delta^{\text{ex-ante}}(\rho), \rho)$ is pessimistic or optimistic. Hence, we study the sign of the derivative of the ex-ante utility function with respect to δ at the point $\delta = 0$. We find that, for a given time preference rate, the sign of this derivative depends on the volatility of the economy. The more the economy is volatile, the more the ex-ante utility-maximizing agent tends to be an optimistic one. This result is intuitive as a more volatile economy allows an agent to dream bigger, and it therefore leads her to be more optimistic ex-ante. Similarly, for a given volatility, we observe that the sign of the derivative depends on the group of agents we study. For instance, for the most patient groups of agents, i.e., looking at the groups Group (\cdot, ρ) for ρ going to zero, we see that this derivative is positive and that the ex-ante utility-maximizing agents of these groups are therefore optimistic. Conversely, when we consider the most impatient groups of agents, the derivative is negative and the ex-ante utility-maximizing agents of these groups

are therefore more optimistic than the average of the agents but are still pessimistic.

From the fourth point of the proposition, we derive that, as in Jouini and Napp (2016), it is possible to have an economy where the threat of elimination is not sufficient to push some agents towards rationality, and that such an economy is characterized by agents who should rationally stay irrational (if their goal is to maximize their level of utility). Note that Jouini and Napp (2016)'s finding is obtained in a two-agent setting where each of the two agents considered takes profit from her impact on equilibrium prices. Hence, we complement this result as we deal with a continuum of agents whose individual price impact is null. In fact, the mechanism at play is quite simple: in a positive growth economy and whatever the asset price is, an optimistic agent extracts more utility from the market portfolio than the rational agent does. Furthermore, when there is a dose of optimism in the economy, an agent whose belief is between the current market (aggregate) belief and rationality is closer to the market belief than the rational agent, and her current decisions are closer to the market portfolio.¹⁵ However, agents' beliefs are constant while the market's belief converges to rationality. Hence, no one can remain indefinitely between the market's belief and rationality, and the time spent within this range is obviously larger when the agent is closer to rationality. We therefore show that there is an optimal position where the advantage obtained while being between these two limits exceeds the subsequent disadvantage.

1.6 Conclusion

Using a correlated statistical distribution to describe their belief and time preference rate heterogeneities, we develop a model with a continuum of heterogeneous investors and derive the implications these different types of heterogeneity have on the behavior of financial markets. We first define the representative agent and determine the equilibrium of this economy. We then look at some market properties (risk-free rate, market price of risk, market

¹⁵In Jouini and Napp (2016), a dose of optimism is obtained assuming that one agent is rational and the other one is optimistic. In the present model, it results from the condition $\delta_0 > 0$.

volatility, trading volume) and notice that the two types of heterogeneity affect them. More specifically, we find that the risk-free rate is procyclical and that the market price of risk is countercyclical. The economy trading volume is also impacted: we find that it increases with the variance of the belief heterogeneity distribution. Moreover, we show that a negative correlation between the two types of heterogeneity increases the volatility of the asset whose dividend process is given by the total endowment of the economy. The higher the correlation in absolute value is, the stronger this excess volatility effect is. A negative correlation between beliefs and time preference rates also decreases the risk-free rate and increases the market price of risk. Additionally, looking at the consumption shares of the agents, we derive that the economy is characterized by waves of optimism and pessimism. Lastly, we study the characteristics of some specific agents and find that the utility-maximizing agents (both ex-post and ex-ante) are different from the surviving one as long as there is an aggregate belief bias in the economy. When agents are optimistic on average, being an optimist reduces the lifetime but increases the utility compared to a rational agent. As in Jouini and Napp (2016), we therefore find that having a shorter life might be more rewarding than a longer one.

To the best of our knowledge, we are the first to consider a continuum of investors with two types of heterogeneity and to introduce some correlation between them. This allows us to derive new implications on the joint impacts of heterogeneities on the behavior of financial markets. Interesting extensions could be to add some risk aversion heterogeneity among the agents and to endogenize the aggregate endowment process in the economy. We leave this for future research.

1.A Proofs

Proof of Proposition 1.1

1. To find the equilibrium of this economy, we need to solve the following program

$$\begin{aligned} c_{\delta,\rho}^* &= c_{\delta,\rho}(p^*, M_\delta, e_{\delta,\rho}^*), \\ e_t^* &= \int c_{\delta,\rho,t}^* d\delta d\rho, \end{aligned}$$

where $e_{\delta,\rho}^* = \nu_{\delta,\rho,k} e^*$ is the initial endowment of Agent (δ, ρ) and

$$c_{\delta,\rho}(p, M, e) \equiv \underset{\mathbb{E} \left[\int_0^\infty p_t (c_{\delta,\rho,t} - e_t) dt \right] \leq 0}{\text{argmax}} \mathbb{E} \left[\int_0^\infty \exp(-\rho t) M_t u(c_{\delta,\rho,t}) dt \right].$$

The first order conditions give immediately

$$\begin{aligned} p_t^* &= (e_t^*)^{-1} \int \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho, \\ c_{\delta,\rho,t}^* &= (p_t^*)^{-1} \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t}, \end{aligned}$$

where $(\lambda_{\delta,\rho,k})_{\delta \in \mathbb{R}, \rho \in \mathbb{R}_+^*}$ are the inverse of the Lagrange multipliers, which satisfy $\lambda_{\delta,\rho,k} = \rho \nu_{\delta,\rho,k} \left(\int \frac{\lambda_{\delta',\rho',k}}{\rho'} d\delta' d\rho' \right)$.

Note that the equation giving the consumption of Agent (δ, ρ) helps to see how the representative agent—whose characteristics are derived in the remainder of the proof—is constructed. More formally

$$c_{\delta,\rho,t}^* = (p_t^*)^{-1} \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} \Leftrightarrow e_t^* = (p_t^*)^{-1} \int \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho.$$

The right hand-side of this equivalence can be interpreted as the equilibrium equation of a representative agent whose total consumption is e^* and whose characteristics are

such that $\exp(-\bar{\rho}_t) \bar{M}_t = \int \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho$.

Note also that we must have $\lambda_{\delta,\rho,k} = \rho \nu_{\delta,\rho,k} \left(\int \frac{\lambda_{\delta',\rho',k}}{\rho'} d\delta' d\rho' \right)$ to ensure that

$$\nu_{\delta,\rho,k} = \frac{\mathbb{E} \left[\int_0^\infty p_t^* c_{\delta,\rho,t}^* dt \right]}{\int \mathbb{E} \left[\int_0^\infty p_t^* c_{\delta',\rho',t}^* dt \right] d\delta' d\rho'}.$$

This last equality follows from the budget constraint of Agent (δ, ρ) which states that the value of the agent's consumption (i.e., $\mathbb{E} \left[\int_0^\infty p_t^* c_{\delta,\rho,t}^* dt \right]$) should (at most) equal the value of her endowment (given by the value of the fraction of the production process that she is endowed with, i.e., $\nu_{\delta,\rho,k} \times \mathbb{E} \left[\int_0^\infty p_t^* e_t^* dt \right] = \nu_{\delta,\rho,k} \times \int \mathbb{E} \left[\int_0^\infty p_t^* c_{\delta',\rho',t}^* dt \right] d\delta' d\rho'$).

Additionally, using Ito's lemma, we easily derive

$$\begin{aligned} dp_t^* &= \left(-\mu + \sigma^2 - \frac{\mathbb{E}_t(\rho^2)}{\mathbb{E}_t(\rho)} - \sigma \frac{\mathbb{E}_t(\delta\rho)}{\mathbb{E}_t(\rho)} \right) p_t^* dt + \left(-\sigma + \frac{\mathbb{E}_t(\delta\rho)}{\mathbb{E}_t(\rho)} \right) p_t^* dW_t \\ &= \mu_{p^*} p_t^* dt + \sigma_{p^*} p_t^* dW_t, \end{aligned}$$

with $\mathbb{E}_t(\cdot)$ the time-dependent mean with weights given by $\tilde{\nu}_{\delta,\rho,k,t}$ defined in Section 1.2.

2. 3. Let us denote by \bar{Q} the belief of the representative agent—associated to $\bar{\delta}_t$ —and by $\bar{\rho}_t$ her time preference rate. Let us also denote by \bar{M} the density of \bar{Q} with respect to P .

Recall that the representative agent of this economy is an agent who, if endowed with the total wealth of the economy, would have a marginal utility equal to the equilibrium price. Hence, we have $\bar{M}_t = \exp(\bar{\rho}_t t) p_t^* e_t^*$.

We derive

$$d\bar{M}_t = (\bar{\rho}_t + \mu_{p^*}(t) + \mu_{e^*}(t) + \sigma_{p^*}(t)\sigma_{e^*}(t)) \bar{M}_t dt + (\sigma_{p^*}(t) + \sigma_{e^*}(t)) \bar{M}_t dW_t$$

$$= \mu_{\bar{M}}(t)\bar{M}_t dt + \sigma_{\bar{M}}(t)\bar{M}_t dW_t.$$

Direct computations give $\mu_{\bar{M}}(t) = \bar{\rho}_t - \frac{\mathbb{E}_t(\rho^2)}{\mathbb{E}_t(\rho)}$ and $\sigma_{\bar{M}}(t) = \frac{\mathbb{E}_t(\delta\rho)}{\mathbb{E}_t(\rho)}$.

By definition, $\mu_{\bar{M}}(t) = 0$. The representative agent's time preference rate is therefore given by $\bar{\rho}_t = \frac{\mathbb{E}_t(\rho^2)}{\mathbb{E}_t(\rho)}$, and the associated variance is given by $\sigma_{\bar{\rho}} = \frac{\mathbb{E}_t(\rho^3)}{\mathbb{E}_t(\rho)} - \bar{\rho}_t^2$.

The representative agent's belief is given by $\sigma_{\bar{M}}$. Hence, we have $\bar{\delta}_t = \sigma_{\bar{M}} = \frac{\mathbb{E}_t(\delta\rho)}{\mathbb{E}_t(\rho)}$, and the associated variance is given by $\sigma_{\bar{\delta}} = \frac{\mathbb{E}_t(\delta^2\rho)}{\mathbb{E}_t(\rho)} - \sigma_{\bar{M}}(t)^2$.

By definition, finding the characteristics such that we have $\bar{M}_t = \exp(\bar{\rho}_t t) p_t^* e_t^*$ ensures the existence of the representative agent.

When $k \neq 0$, we use the computations of Appendix 1.B.1 to obtain the explicit computations of $\bar{\rho}_t$ and $\bar{\delta}_t$.

Let briefly study the function Ψ . We have $\Psi(X_t) = \frac{|k|\sqrt{t}}{\sqrt{1+t\bar{\omega}^2}} \frac{\mathbb{E}_t(\rho^2)}{\mathbb{E}_t(\rho)}$.

Knowing that the time preference rates are non-negative, we derive that $\Psi(X_t) \geq 0$.

Using Taylor expansions, we also easily derive that $\Psi(x) \underset{x \rightarrow +\infty}{\sim} \frac{3 + \mathcal{O}\left(\frac{1}{x^2}\right)}{x - \frac{6}{x} + \mathcal{O}\left(\frac{1}{x^3}\right)}$, and we conclude that $\Psi(X_t)$ converges to zero when t goes to infinity.

In the uncorrelated case ($k = 0$), the computations simplify to

$$\begin{aligned} \bar{\rho}_t &= \frac{\int_0^\infty \rho^2 \exp(-\rho t) \frac{\vartheta^l}{\Gamma(l)} \rho^{l-1} \exp(-\rho\vartheta) d\rho}{\int_0^\infty \rho \exp(-\rho t) \frac{\vartheta^l}{\Gamma(l)} \rho^{l-1} \exp(-\rho\vartheta) d\rho} = \frac{l+1}{\vartheta+t}, \\ \sigma_{\bar{\rho}} &= \frac{\int_0^\infty \rho^3 \exp(-\rho t) \frac{\vartheta^l}{\Gamma(l)} \rho^{l-1} \exp(-\rho\vartheta) d\rho}{\int_0^\infty \rho \exp(-\rho t) \frac{\vartheta^l}{\Gamma(l)} \rho^{l-1} \exp(-\rho\vartheta) d\rho} - \bar{\rho}_t^2 = \frac{l+1}{(\vartheta+t)^2}, \\ \bar{\delta}_t &= \frac{\int_{\mathbb{R}} \delta M_{\delta,t} \frac{1}{\sqrt{2\pi\bar{\omega}}} \exp\left(-\frac{(\delta-\delta_0)^2}{2\bar{\omega}^2}\right) d\delta}{\int_{\mathbb{R}} M_{\delta,t} \frac{1}{\sqrt{2\pi\bar{\omega}}} \exp\left(-\frac{(\delta-\delta_0)^2}{2\bar{\omega}^2}\right) d\delta} = \frac{\delta_0 + \bar{\omega}^2 W_t}{1+t\bar{\omega}^2}, \end{aligned}$$

$$\sigma_{\bar{\delta}} = \frac{\int_{\mathbb{R}} \delta^2 M_{\delta,t} \frac{1}{\sqrt{2\pi\bar{\omega}}} \exp\left(-\frac{(\delta - \delta_0)^2}{2\bar{\omega}^2}\right) d\delta}{\int_{\mathbb{R}} M_{\delta,t} \frac{1}{\sqrt{2\pi\bar{\omega}}} \exp\left(-\frac{(\delta - \delta_0)^2}{2\bar{\omega}^2}\right) d\delta} - \sigma_{\bar{M}}(t)^2 = \frac{\bar{\omega}^2}{1 + t\bar{\omega}^2}.$$

■

Proof of Proposition 1.2 Let consider an asset which does not pay dividends and let us denote by Z its price process. We have that p^*Z is a martingale. Hence, $\mu_Z + \mu_{p^*} + \sigma_Z \sigma_{p^*} = 0$. In the case of a riskless asset, we have that μ_Z is the risk-free rate and that $\sigma_Z = 0$. We obtain $r_t^f = -\mu_{p^*}(t)$.

In the case of a risky asset, we therefore have $\mu_Z - r^f + \sigma_Z \sigma_{p^*} = 0$, which leads to $\frac{\mu_Z - r^f}{\sigma_Z} = -\sigma_{p^*}$. Note that it does not depend on Z and that we therefore have $MPR_t = -\sigma_{p^*}(t)$.

Using Ito's Lemma on p^* we straightforwardly find the result (cf Proof of Proposition 1.1).

■

Proof of Proposition 1.3 We have $S_t = e_t^* \frac{\int \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}{\int \rho \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}$.

Using Ito's Lemma we get

$$\begin{aligned} \sigma_{S,t} &= \sigma + \frac{\int \delta \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}{\int \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho} - \frac{\int \delta \rho \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}{\int \rho \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho} \\ &= \sigma - \frac{\text{cov}_t(\delta, \rho)}{\mathbb{E}_t(\rho)}. \end{aligned}$$

Using the computations of Appendix 1.B.1, we obtain

$$\sigma_{S,t} = \sigma + \frac{\text{sgn}(k)}{\sqrt{t}\sqrt{1+t\bar{\omega}^2}} \Upsilon(X_t),$$

$$\text{with } \Upsilon(x) = \frac{2 - 3\frac{1-\Phi(x)}{\varphi(x)}x + \left(\frac{1-\Phi(x)}{\varphi(x)}\right)^2(-1+x^2)}{-x + \frac{1-\Phi(x)}{\varphi(x)}(1+2x^2) - \left(\frac{1-\Phi(x)}{\varphi(x)}\right)^2(x+x^3)}.$$

We can rewrite $\Upsilon(x)$ as

$$\Upsilon(x) = \frac{2 - 3\frac{1-\Phi(x)}{\varphi(x)}x + \left(\frac{1-\Phi(x)}{\varphi(x)}\right)^2(-1+x^2)}{\left(1 - \frac{1-\Phi(x)}{\varphi(x)}x\right)^2} \frac{1 - \frac{1-\Phi(x)}{\varphi(x)}x}{-x + \frac{1-\Phi(x)}{\varphi(x)}(1+x^2)}.$$

Moreover, as $\mathbb{V}_t(\rho) > 0$ and $\mathbb{E}_t(\rho) > 0$, we derive that

$$2 - 3\frac{1-\Phi(X_t)}{\varphi(X_t)}X_t + \left(\frac{1-\Phi(X_t)}{\varphi(X_t)}\right)^2(-1+X_t^2) > 0 \text{ and } \frac{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t}{-X_t + \frac{1-\Phi(X_t)}{\varphi(X_t)}(1+X_t^2)} > 0. \text{ Hence,}$$

we obtain that $\Upsilon(X_t) \geq 0$.

Finally, using Taylor expansions, we easily get $\Upsilon(x) \underset{x \rightarrow +\infty}{\sim} \frac{\frac{1}{x} + \mathcal{O}\left(\frac{1}{x^3}\right)}{1 + \mathcal{O}\left(\frac{1}{x^2}\right)}$, and we conclude that $\Upsilon(X_t)$ converges to zero when t goes to infinity. ■

Proof of Proposition 1.4 We have $V_{\delta,\rho,t} = e_t^* \frac{\int_t^\infty \lambda_{\delta,\rho,k} \exp(-\rho s) M_{\delta,t} ds}{\int \lambda_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}$.

Using Ito's Lemma and rearranging the terms we find

$$\begin{aligned} dV_{\delta,\rho,t} &= (\mu - \rho + \sigma\delta + \bar{\rho}_t - (\sigma + \delta)\bar{\delta}_t + \bar{\delta}_t^2) V_{\delta,\rho,t} dt + (\sigma + \delta - \bar{\delta}_t) V_{\delta,\rho,t} dW_t \\ &= \mu_{V_{\delta,\rho,t}} V_{\delta,\rho,t} dt + \sigma_{V_{\delta,\rho,t}} V_{\delta,\rho,t} dW_t. \end{aligned}$$

As reported in Section 1.4.3, the optimal number of shares Agent (δ, ρ) should invest in the risky asset is given by $\alpha_{\delta,\rho,t} = \frac{V_{\delta,\rho,t}}{S_t} \frac{\sigma_{V_{\delta,\rho,t}}}{\sigma_{S,t}}$. Isolating the volatility of the stochastic differential equation satisfied by $\alpha_{\delta,\rho}$, tedious computations lead to

$$\sigma_{\alpha_{\delta,\rho,t}} = \frac{V_{\delta,\rho,t}}{S_t} \frac{1}{\sigma_{S,t}} \left((\sigma + \delta - \bar{\delta}_t) \left(\delta - \mathbb{E}_t(\delta) + \frac{1}{\sigma_{S,t}} \left(\frac{\text{cov}_t(\delta^2, \rho)}{\mathbb{E}_t(\rho)} - \frac{\text{cov}_t(\delta, \rho) (\mathbb{E}_t(\delta) + \bar{\delta}_t)}{\mathbb{E}_t(\rho)} \right) \right) - \sigma_{\bar{\delta}} \right).$$

Taking the absolute value of $\sigma_{\alpha_{\delta,\rho}}$ and knowing that $V_{\delta,\rho}$, S , and σ_S are non-negative, we derive the result of Proposition 1.4.

In the uncorrelated case, we have $\sigma_{S,t} = \sigma$, $\text{cov}_t(\delta^2, \rho) = 0$, and $\text{cov}_t(\delta, \rho) = 0$. Moreover

$$\bar{\delta}_t = \frac{\mathbb{E}_t(\delta\rho)}{\mathbb{E}_t(\rho)} = \frac{\text{cov}_t(\delta, \rho) + \mathbb{E}_t(\delta)\mathbb{E}_t(\rho)}{\mathbb{E}_t(\rho)} = \mathbb{E}_t(\delta).$$

Hence, in this case, the trading volume of Agent (δ, ρ) is given by

$$|\sigma_{\alpha_{\delta,\rho,t}}| = \frac{V_{\delta,\rho,t}}{S_t} \frac{(\sigma + \delta - \bar{\delta}_t)(\delta - \bar{\delta}_t) - \sigma_{\bar{\delta}}}{\sigma}.$$

■

Proof of Proposition 1.5 We have $Vol_t = \frac{1}{2} \int |\sigma_{\alpha_{\delta,\rho,t}}| d\delta d\rho$.

To compute this integral, we need to know the sign of the term in absolute value. We have

$$\begin{aligned} \sigma_{\alpha_{\delta,\rho,t}} > 0 &\Leftrightarrow (\sigma + \delta - \bar{\delta}_t)(\delta - \bar{\delta}_t) - \sigma_{\bar{\delta}} > 0 \\ &\Leftrightarrow \delta^2 + (\sigma - 2\bar{\delta}_t)\delta + \bar{\delta}_t^2 - \bar{\delta}_t\sigma - \sigma_{\bar{\delta}} > 0. \end{aligned}$$

Solving this equation, we get $\sigma_{\alpha_{\delta,\rho,t}} > 0 \Leftrightarrow \delta \in]-\infty, \delta_2[\cup]\delta_1, +\infty[$ with

$$\delta_1 = \frac{-\sigma + 2\bar{\delta}_t + \sqrt{\sigma^2 + 4\sigma_{\bar{\delta}}}}{2} \text{ and } \delta_2 = \frac{-\sigma + 2\bar{\delta}_t - \sqrt{\sigma^2 + 4\sigma_{\bar{\delta}}}}{2}.$$

Hence, we have $Vol_t = \frac{1}{2} \left(\int_0^\infty \int_{-\infty}^{\delta_2} \sigma_{\alpha_{\delta,\rho,t}} d\delta d\rho - \int_0^\infty \int_{\delta_2}^{\delta_1} \sigma_{\alpha_{\delta,\rho,t}} d\delta d\rho + \int_0^\infty \int_{\delta_1}^\infty \sigma_{\alpha_{\delta,\rho,t}} d\delta d\rho \right)$.

Moreover, as $\int_0^\infty \int_{-\infty}^\infty \sigma_{\alpha_{\delta,\rho,t}} d\delta d\rho = 0$, we have $Vol_t = - \int_0^\infty \int_{\delta_2}^{\delta_1} \sigma_{\alpha_{\delta,\rho,t}} d\delta d\rho$.

Tedious computations lead to

$$\begin{aligned} Vol_t = \frac{\sqrt{\sigma_{\bar{\delta}}}}{2\sigma} &\left(\left(\sigma + \sqrt{\sigma^2 + 4\sigma_{\bar{\delta}}} \right) \varphi \left(\frac{-\sigma + \sqrt{\sigma^2 + 4\sigma_{\bar{\delta}}}}{2\sqrt{\sigma_{\bar{\delta}}}} \right) \right. \\ &\left. - \left(\sigma - \sqrt{\sigma^2 + 4\sigma_{\bar{\delta}}} \right) \varphi \left(\frac{-\sigma - \sqrt{\sigma^2 + 4\sigma_{\bar{\delta}}}}{2\sqrt{\sigma_{\bar{\delta}}}} \right) \right). \end{aligned}$$

■

Proof of Proposition 1.6 Explicit computations of the consumption share of Agent (δ, ρ) lead to

$$\tau_{\delta, \rho, t} = \frac{\sqrt{1 + t\bar{\omega}^2}}{\sqrt{2\pi\bar{\omega}}} \frac{(\vartheta + t)^{l+1} \rho^l}{\Gamma(l+1)} \exp\left(-\left(\rho + \frac{\delta^2}{2}\right)t - \frac{\delta^2}{2\bar{\omega}^2} + \frac{1}{\sigma_{\bar{\delta}}}\delta\bar{\delta}_t - \frac{1}{2\sigma_{\bar{\delta}}}\bar{\delta}_t^2 - \rho\vartheta\right).$$

Defining the survival index of Agent (δ, ρ) by $\rho + \frac{\delta^2}{2}$ and studying the limit of $\tau_{\delta, \rho}$ when t goes to infinity, we derive the result straightforwardly. ■

Proof of Proposition 1.7 The aggregate consumption share of the optimistic (resp. pessimistic) agents τ_{opt} (resp. τ_{pes}) is given by

$$\tau_{opt, t} = \int_0^\infty \int_0^\infty \tau_{\delta, \rho, t} d\delta d\rho \quad \left(\text{resp. } \tau_{pes, t} = \int_0^\infty \int_{-\infty}^0 \tau_{\delta, \rho, t} d\delta d\rho\right).$$

Direct computations give us

$$\tau_{opt, t} = 1 - \Phi\left(-\frac{W_t + \frac{\delta_0}{\bar{\omega}^2}}{\sqrt{t + \frac{1}{\bar{\omega}^2}}}\right) = 1 - \Phi\left(-\frac{\bar{\delta}_t}{\sqrt{\sigma_{\bar{\delta}}}}\right), \quad \tau_{pes, t} = \Phi\left(-\frac{W_t + \frac{\delta_0}{\bar{\omega}^2}}{\sqrt{t + \frac{1}{\bar{\omega}^2}}}\right) = \Phi\left(-\frac{\bar{\delta}_t}{\sqrt{\sigma_{\bar{\delta}}}}\right).$$

We easily see that $\lim_{W_t \rightarrow +\infty} \tau_{opt, t} = 1$ and $\lim_{W_t \rightarrow +\infty} \tau_{pes, t} = 0$. Similarly $\lim_{W_t \rightarrow -\infty} \tau_{opt, t} = 0$ and $\lim_{W_t \rightarrow -\infty} \tau_{pes, t} = 1$. ■

Proof of Proposition 1.8

1. We have

$$c_{\delta, \rho, t}^* = \frac{\sqrt{1 + t\bar{\omega}^2}}{\sqrt{2\pi\bar{\omega}}} \frac{(\vartheta + t)^{l+1} \rho^l}{\Gamma(l+1)} \exp\left(-\left(\rho + \frac{\delta^2}{2}\right)t - \frac{\delta^2}{2\bar{\omega}^2} + \frac{1}{\sigma_{\bar{\delta}}}\delta\bar{\delta}_t - \frac{1}{2\sigma_{\bar{\delta}}}\bar{\delta}_t^2 - \rho\vartheta\right) e_t^*.$$

Some algebra leads to

$$\begin{aligned}
 U_{\delta,\rho}^{\text{ex-post}}(c_{\delta,\rho}^*) &= \mathbb{E} \left[\int_0^\infty \exp(-\rho s) \ln(c_{\delta,\rho,s}^*) ds \right] \\
 &= \int_0^\infty \exp(-\rho s) \left(\frac{\ln(s + \frac{1}{\bar{\omega}^2})}{2} + (l+1) \ln(s + \vartheta) - \frac{\delta_0^2}{2(s\bar{\omega}^4 + \bar{\omega}^2)} - \frac{s}{2(s + \frac{1}{\bar{\omega}^2})} \right) ds \\
 &\quad - \vartheta + \frac{l \ln(\rho) - \ln(\sqrt{2\pi}\Gamma(l+1)) - 1}{\rho} + \frac{2\delta_0\delta - \delta^2}{2\bar{\omega}^2\rho} + \frac{2\mu + \sigma^2 - \delta^2}{2\rho^2}.
 \end{aligned}$$

Deriving this expression with respect to δ , we get $\frac{\partial U_{\delta,\rho}^{\text{ex-post}}}{\partial \delta}(c_{\delta,\rho}^*) = \frac{-(1 + \frac{\bar{\omega}^2}{\rho})\delta + \delta_0}{\bar{\omega}^2\rho}$.

Hence, we have $\frac{\partial U_{\delta,\rho}^{\text{ex-post}}}{\partial \delta}(c_{\delta,\rho}^*) \Big|_{\delta=\delta^{\text{ex-post}}} = 0 \Leftrightarrow \delta^{\text{ex-post}} = \frac{\delta_0}{1 + \frac{\bar{\omega}^2}{\rho}}$.

Note that $\frac{\partial^2 U_{\delta,\rho}^{\text{ex-post}}}{\partial \delta^2}(c_{\delta,\rho}^*) = \frac{-(1 + \frac{\bar{\omega}^2}{\rho})}{\bar{\omega}^2\rho} < 0$, which ensures that this is a maximum.

2. Direct computations give

$$\begin{aligned}
 U_{\delta,\rho}^{\text{ex-ante}}(c_{\delta,\rho}^*) &= \mathbb{E}^{Q_\delta} \left[\int_0^\infty \exp(-\rho s) \ln(c_{\delta,\rho,s}^*) ds \right] = \mathbb{E} \left[\int_0^\infty \exp(-\rho s) M_{\delta,t} \ln(c_{\delta,\rho,s}^*) ds \right] \\
 &= \int_0^\infty \exp(-\rho s) \left(\frac{\ln(s + \frac{1}{\bar{\omega}^2})}{2} + (l+1) \ln(s + \vartheta) \right. \\
 &\quad \left. - \frac{\delta_0^2}{2(s\bar{\omega}^4 + \bar{\omega}^2)} - \frac{\delta_0\delta s}{1 + s\bar{\omega}^2} - \frac{s}{2(s + \frac{1}{\bar{\omega}^2})} - \frac{\delta^2 s^2}{2(s + \frac{1}{\bar{\omega}^2})} \right) ds \\
 &\quad - \vartheta + \frac{l \ln(\rho) - \ln(\sqrt{2\pi}\Gamma(l+1)) - 1}{\rho} + \frac{2\delta_0\delta - \delta^2}{2\bar{\omega}^2\rho} + \frac{2\mu + (\sigma + \delta)^2}{2\rho^2}.
 \end{aligned}$$

Deriving this expression with respect to δ , we get

$$\frac{\partial U_{\delta,\rho}^{\text{ex-ante}}}{\partial \delta}(c_{\delta,\rho}^*) = \int_0^\infty \exp(-\rho s) \left(-\frac{\delta_0 s}{1 + s\bar{\omega}^2} - \frac{\delta \bar{\omega}^2 s^2}{1 + s\bar{\omega}^2} \right) ds + \frac{\delta_0 - \delta}{\bar{\omega}^2\rho} + \frac{\delta + \sigma}{\rho^2}.$$

We also have

$$\frac{\partial^2 U_{\delta,\rho}^{\text{ex-ante}}}{\partial \delta^2} (c_{\delta,\rho}^*) = - \int_0^\infty \exp(-\rho s) \frac{\bar{\omega}^2 s^2}{1 + s\bar{\omega}^2} ds - \frac{1}{\bar{\omega}^2 \rho} + \frac{1}{\rho^2} < 0.$$

We derive that the derivative of the ex-ante utility function is a monotonic decreasing function. Hence, if this derivative evaluated at a given point is positive, the maximum of the ex-ante utility function is reached for a more optimistic belief (as a positive sign for the derivative means that the utility function is increasing).

In particular, for the average agent, simple computations give

$$\left. \frac{\partial U_{\delta,\rho}^{\text{ex-ante}}}{\partial \delta} (c_{\delta,\rho}^*) \right|_{\delta=\delta_0} = \int_0^\infty \exp(-\rho s) \left(-\frac{\delta_0 s}{1 + s\bar{\omega}^2} - \frac{\delta_0 \bar{\omega}^2 s^2}{1 + s\bar{\omega}^2} \right) ds + \frac{\delta_0 + \sigma}{\rho^2} = \frac{\sigma}{\rho^2} > 0.$$

3. When $\delta_0 > 0$, we have $\left. \frac{\partial U_{\delta,\rho}^{\text{ex-ante}}}{\partial \delta} (c_{\delta,\rho}^*) \right|_{\delta=0} > 0$, as this derivative is a monotonic decreasing function which is positive for $\delta = \delta_0$.

When $\delta_0 < 0$, the sign of the derivative for $\delta = 0$ depends on σ . We have

$$\left. \frac{\partial U_{\delta,\rho}^{\text{ex-ante}}}{\partial \delta} (c_{\delta,\rho}^*) \right|_{\delta=0} = - \int_0^\infty \exp(-\rho s) \frac{\delta_0 s}{1 + s\bar{\omega}^2} ds + \frac{\delta_0}{\bar{\omega}^2 \rho} + \frac{\sigma}{\rho^2}.$$

Using basic algebra and the fact that, for all s , we have $\frac{s\bar{\omega}^2}{1 + s\bar{\omega}^2} \leq 1$, we obtain

$$\int_0^\infty \exp(-\rho s) \frac{s\bar{\omega}^2}{1 + s\bar{\omega}^2} \rho ds \leq \int_0^\infty \exp(-\rho s) \rho ds = 1$$

which leads to

$$- \int_0^\infty \exp(-\rho s) \frac{\delta_0 s}{1 + s\bar{\omega}^2} ds + \frac{\delta_0}{\bar{\omega}^2 \rho} \leq 0.$$

Hence, we derive that when σ is sufficiently high, the derivative of the ex-ante utility function evaluated at $\delta = 0$ is positive.

4. 5. By definition Agent $(\delta^{\text{ex-post}}(\rho), \rho)$ has a higher ex-post utility than the rational agent. Moreover, when $\delta_0 \geq 0$, we have $0 \leq \delta^{\text{ex-post}}(\rho) \leq \delta_0$. Using a similar argument as for the proof of point 3. we derive that she also has a higher ex-ante utility.

When $\delta_0 < 0$ and σ is low, the ex-ante utility-maximizing agent of Group $(., \rho)$ has a less optimistic belief than the rational agent (see proof of point 3.). We have $\delta_0 \leq \delta^{\text{ex-ante}}(\rho) \leq 0$ (see proof of point 2.). Using the fact that the utility functions are monotonic decreasing functions once they have reached their maximum, we conclude that Agent $(\delta^{\text{ex-post}}(\rho), \rho)$ (resp. Agent $(\delta^{\text{ex-ante}}(\rho), \rho)$) has also a higher ex-ante (ex-post) utility than the rational agent if $\delta^{\text{ex-ante}}(\rho) \leq \delta^{\text{ex-post}}(\rho) \leq 0$ (resp. $\delta^{\text{ex-post}}(\rho) \leq \delta^{\text{ex-ante}}(\rho) \leq 0$).

■

1.B Useful computations and additional results

1.B.1 Moments with respect to the density $\tilde{\nu}_{\delta,\rho,k,t}$

We define

$$X_t = \frac{\sqrt{1+t\bar{\omega}^2}}{|k|\sqrt{t}} \left(t + \vartheta + k \frac{W_t - \delta_0 t}{1+t\bar{\omega}^2} \right).$$

Tedious computations give

$$\begin{aligned} \mathbb{E}_t(\delta) &= \int \delta \tilde{\nu}_{\delta,\rho,k,t} d\delta d\rho = \frac{\int \delta \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho}{\int \nu_{\delta,\rho,k} \exp(-\rho t) M_{\delta,t} d\delta d\rho} \\ &= \frac{\delta_0 + \bar{\omega}^2 W_t}{1+t\bar{\omega}^2} - \frac{\operatorname{sgn}(k)}{\sqrt{t}\sqrt{1+t\bar{\omega}^2}} \frac{-X_t + \frac{1-\Phi(X_t)}{\varphi(X_t)}(1+X_t^2)}{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t}, \\ \mathbb{E}_t(\delta^2) &= \int \delta^2 \tilde{\nu}_{\delta,\rho,k,t} d\delta d\rho \\ &= \frac{\bar{\omega}^2}{1+t\bar{\omega}^2} + \left(\frac{\delta_0 + \bar{\omega}^2 W_t}{1+t\bar{\omega}^2} \right)^2 + \frac{2+X_t^2 - \frac{1-\Phi(X_t)}{\varphi(X_t)}(3X_t+X_t^3)}{t(1+t\bar{\omega}^2)} \\ &\quad - 2 \frac{\operatorname{sgn}(k)}{\sqrt{t}\sqrt{1+t\bar{\omega}^2}} \frac{-X_t + \frac{1-\Phi(X_t)}{\varphi(X_t)}(1+X_t^2)}{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t} \frac{\delta_0 + \bar{\omega}^2 W_t}{1+t\bar{\omega}^2}, \\ \mathbb{E}_t(\rho) &= \int \rho \tilde{\nu}_{\delta,\rho,k,t} d\delta d\rho = \frac{\sqrt{1+t\bar{\omega}^2}}{|k|\sqrt{t}} \frac{-X_t + \frac{1-\Phi(X_t)}{\varphi(X_t)}(1+X_t^2)}{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t}, \\ \mathbb{E}_t(\rho^2) &= \int \rho^2 \tilde{\nu}_{\delta,\rho,k,t} d\delta d\rho = \frac{1+t\bar{\omega}^2}{k^2 t} \frac{2+X_t^2 - \frac{1-\Phi(X_t)}{\varphi(X_t)}(3X_t+X_t^3)}{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t}, \\ \mathbb{E}_t(\delta\rho) &= \int \delta\rho \tilde{\nu}_{\delta,\rho,k,t} d\delta d\rho \\ &= \frac{\delta_0 + \bar{\omega}^2 W_t}{|k|\sqrt{t}\sqrt{1+t\bar{\omega}^2}} \frac{-X_t + \frac{1-\Phi(X_t)}{\varphi(X_t)}(1+X_t^2)}{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t} - \frac{1}{kt} \frac{2+X_t^2 - \frac{1-\Phi(X_t)}{\varphi(X_t)}(3X_t+X_t^3)}{1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t}. \end{aligned}$$

We derive

$$\begin{aligned}
 \mathbb{V}_t(\delta) &= \mathbb{E}_t(\delta^2) - \mathbb{E}_t(\delta)^2 \\
 &= \frac{\bar{\omega}^2}{1 + t\bar{\omega}^2} + \frac{2 - \frac{1-\Phi(X_t)}{\varphi(X_t)}(5X_t + X_t^3) + \left(\frac{1-\Phi(X_t)}{\varphi(X_t)}\right)^2(-1 + 6X_t^2 + 2X_t^4)}{t(1 + t\bar{\omega}^2)\left(1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t\right)^2} \\
 &\quad - \frac{\left(\frac{1-\Phi(X_t)}{\varphi(X_t)}\right)^3(3X_t^3 + X_t^5)}{t(1 + t\bar{\omega}^2)\left(1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t\right)^2}, \\
 \mathbb{V}_t(\rho) &= \mathbb{E}_t(\rho^2) - \mathbb{E}_t(\rho)^2 = \frac{1 + t\bar{\omega}^2}{k^2t} \frac{2 - 3\frac{1-\Phi(X_t)}{\varphi(X_t)}X_t + \left(\frac{1-\Phi(X_t)}{\varphi(X_t)}\right)^2(-1 + X_t^2)}{\left(1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t\right)^2}, \\
 \text{cov}_t(\delta, \rho) &= \mathbb{E}_t(\delta\rho) - \mathbb{E}_t(\delta)\mathbb{E}_t(\rho) = -\frac{1}{kt} \frac{2 - 3\frac{1-\Phi(X_t)}{\varphi(X_t)}X_t + \left(\frac{1-\Phi(X_t)}{\varphi(X_t)}\right)^2(-1 + X_t^2)}{\left(1 - \frac{1-\Phi(X_t)}{\varphi(X_t)}X_t\right)^2}.
 \end{aligned}$$

Finally, we have $\text{corr}_t(\delta, \rho) = \frac{\text{cov}_t(\delta, \rho)}{\sqrt{\mathbb{V}_t(\delta)\mathbb{V}_t(\rho)}}$.


As $\mathbb{V}_t(\rho) > 0$, we know that $2 - 3\frac{1-\Phi(X_t)}{\varphi(X_t)}X_t + \left(\frac{1-\Phi(X_t)}{\varphi(X_t)}\right)^2(-1 + X_t^2) > 0$. Hence, we conclude that $\text{sgn}(\text{corr}_t(\delta, \rho)) = \text{sgn}(\text{cov}_t(\delta, \rho)) = -\text{sgn}(k)$.

1.B.2 Consumption shares and utility-maximizing agents in the correlated case

In unreported computations that are available upon request, we find that the correlation between the two types of heterogeneity does not affect the survival implications of the model. As in the uncorrelated economy, the surviving agent is the most patient of the rational agents. However, the correlation has an impact on the utility-maximizing agents of the economy. The ex-post utility-maximizing agent of Group $(., \rho)$ is characterized by $\delta^{\text{ex-post}}(\rho) = \frac{\delta_0 - k\rho}{1 + \frac{\bar{\omega}^2}{\rho}}$. We notice that this ex-post utility-maximizing agent shares similarities with the ex-post utility-maximizing agent of the uncorrelated economy. In particular, we observe that this agent is biased and that she therefore differs from the surviving one. In fact, the only group of


agents for which the surviving agent is also the one maximizing her ex-post utility is the one characterized by $\rho = \frac{\delta_0}{k}$. We also observe that the correlation between the two types of heterogeneity impacts the characteristics of such an agent. We see that the parameter k impacts negatively $\delta^{\text{ex-post}}(\rho)$, meaning that a positive k , or, stated differently, a negative correlation between the belief and the time preference rate, reduces the optimistic optimal bias an agent should have to be the ex-post utility-maximizing one. In other words, if we assume that the more optimistic agents are also the most patient ones, the heterogeneity effect on the ex-post utility-maximizing agent is mitigated.

Figure 1.1: Time evolution of the median risk-free rate depending on the correlation



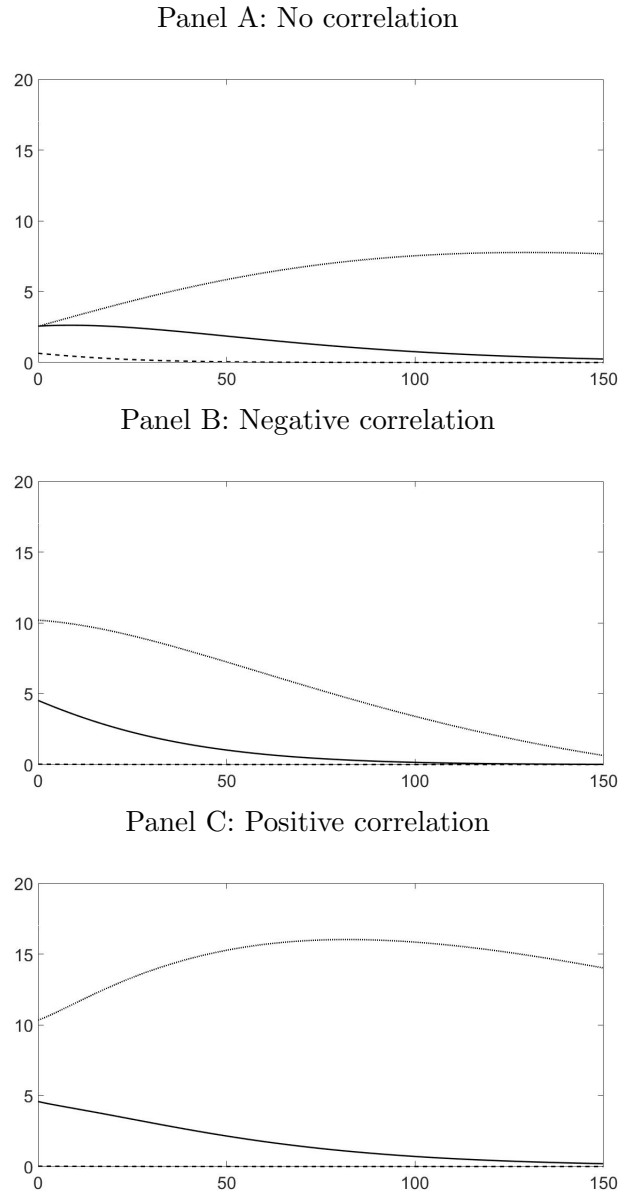
The figure presents the evolution over time of the median risk-free rate for different levels of correlation between the two types of heterogeneity. The dotted (resp. dash-dot, thin, dashed, solid) line corresponds to a median correlation at time $t = 1$ of -0.75 (resp. -0.25 , 0 , 0.25 , 0.75). The baseline parameter values are: $\mu = 14.23\%$, $\sigma = 8.25\%$, $\delta_0 = 0$, $\bar{w} = 3.39\%$, $l = 2$, and $\vartheta = 50$.

Figure 1.2: Time evolution of the median market volatility depending on the correlation




The figure presents the evolution over time of the median market volatility for different levels of correlation between the two types of heterogeneity. The dotted (resp. dash-dot, thin, dashed, solid) line corresponds to a median correlation at time $t = 1$ of -0.75 (resp. -0.25 , 0 , 0.25 , 0.75). We use the same baseline parameter values as in Figure 1.1.

Figure 1.3: Time evolution of the median trading volume depending on the time preferences



The figure presents the evolution over time of the median trading volume of three rational agents having different time preference rates. The figure shows the uncorrelated case (Panel A), the strongly negatively correlated case (Panel B), and the strongly positively correlated case (Panel C). The dotted (resp. solid, dashed) line corresponds to the patient (resp. neutral, impatient) agent whose time preference rate is given by the 10th (resp. 50th, 90th) percentile of the Gamma distribution used to describe the time preference rate heterogeneity. We use the same baseline parameter values as in Figure 1.1.

Figure 1.4: Time evolution of the aggregate consumption share of the optimistic agents depending on the trajectory



The figure presents the evolution over time of the aggregate consumption share of the optimistic agents with three different trajectories. The dotted (resp. solid, dashed) line corresponds to the good (resp. neutral, bad) trajectory defined by $W_t = \sqrt{t}$ (resp. $W_t = 0$, $W_t = -\sqrt{t}$) at each date. We use the same baseline parameter values as in Figure 1.1.

Table 1.1: Average half-lives of some specific agents

	$\delta_{10\%}$	$\delta_{25\%}$	$\delta_{50\%}$	$\delta_{75\%}$	$\delta_{90\%}$
$\rho_{10\%}$	399.1	562.3	731.1	562.3	399.1
$\rho_{50\%}$	84.8	110.2	133.1	110.2	84.8
$\rho_{90\%}$	19.5	22.4	24.6	22.4	19.5

The table contains the average half-life of 15 specific agents. The characteristics of the agents are such that they partition the initial wealth share distribution of the continuum of agents: $\delta_{i\%}$ and $\rho_{j\%}$ are such that Agent $(\delta_{i\%}, \rho_{j\%})$ is more optimistic than $i\%$ of the agents and more impatient than $j\%$ of the agents. We use the same baseline parameter values as in Figure 1.1.

Chapter 2

Disagreeing forever: a testable model
with non-vanishing belief
heterogeneity

Abstract

I develop an overlapping generations model where each generation is constituted of a continuum of agents with heterogeneous beliefs. This belief heterogeneity persists across successive generations, and I thus obtain stationary results. The model can equivalently be described as an economy with a sliding horizon where infinitely-lived investors continuously revise their consumption plans, which eases the construction of the intertemporal representative agent. I further derive that the equilibrium stock mean return and volatility increase with belief dispersion. Using analyst forecasts from the IBES database, I show suggestive evidence that these positive relations are empirically verified and hold with data sampled at various frequencies when considering the market.

Keywords: Heterogeneous beliefs, Overlapping generations, Continuum of agents, Asset pricing

JEL classification: D53 - D90 - G12 - G40

2.1 Introduction

It is well-established that there is a great belief heterogeneity among stock market participants. Using a recent survey administered to a panel of retail investors, Giglio et al. (2020) for instance show that beliefs are characterized by a large and persistent individual heterogeneity, and that investors are likely to exhibit a willingness to “agree to disagree.” While early studies argued that traders with biased beliefs could be neglected, the recent literature has shown these views to be wrong, and has studied, both from empirical and theoretical perspectives, the implications of such belief dispersion. In particular, Jouini and Napp (2011) show theoretically that heterogeneity is important per se, as an economy with biased investors who are rational on average differs markedly from an economy with rational investors only. Moreover, surveys suggest that investors do not agree more on future stock market returns nowadays than they did in the past. Heterogeneous belief models should thus incorporate this disagreement persistence, which is unfortunately not often the case (see, e.g., Atmaz and Basak, 2018). Not only is this feature important for the sake of empirical relevance, but also because vanishing belief dispersion prevents the belief heterogeneity implications to be tested relevantly over long periods.

In this paper, I develop and empirically test a tractable model in which belief heterogeneity does not progressively vanish. More precisely, in a dynamic general equilibrium framework, I develop an overlapping heterogeneous generations model evolving in continuous time, where each generation is constituted of a continuum of constant relative risk aversion (CRRA) investors who agree to disagree. As time goes by, each member of a given generation gives birth to an investor of the next generation, who is endowed with the same beliefs. At the end of their life, the agents then consume part of their wealth and bequeath the remaining part to the next generation. The combination of the intra-family belief transmission assumption and of a dynamic family budget constraint ensures that the belief dispersion remains persistent across generations, as empirically observed. My main theoretical contribution is thus to derive stationary results with non-vanishing belief heterogeneity

that confirm part of the existing results found in models where the agents unrealistically tend towards rationality. It also allows me to study if the model implications are verified empirically, and if the belief heterogeneity impacts persist over long horizons. Moreover, the model incorporates continuous effective consumption, which differs from Atmaz and Basak (2018) where there is a single consumption date when uncertainty is resolved.

An important feature of the model is that it considers a large number of agents, which has two main advantages. First, it allows me to consider an unbounded investor type space, or, stated differently, to take into account every possible existing belief. In a model with bounded belief biases, the most biased agents (towards optimism or pessimism) would eventually control almost all of the wealth in the economy in the most extreme states (i.e., the very good or very bad ones), and belief heterogeneity would disappear. The choice of a continuum thus guarantees that belief heterogeneity persists not only for all dates but also in all states of the world. The second advantage of considering an infinite number of investors is that I can use a statistical distribution to describe their wealth shares. Similarly to, e.g., Cvitanic and Malamud (2011), I assume that the initial distribution is a Gaussian one with exogenously determined parameters, and I show that the normality of the wealth share distribution remains as time goes by. Thus, I only need to estimate two parameters, namely the average belief bias and the belief dispersion, in order to describe the agents' beliefs, which eases the model empirical assessment. It also allows me to easily disentangle between the impacts of the first and second moments. This is interesting because, as underlined by Brandon and Wang (2020), most models focus on the effects of the latter.

Before discussing its theoretical implications, let me briefly talk about a dual approach of the model. It is based on an alternative model populated by a continuum of infinitely-lived CRRA investors, who continuously revise their consumption plans and, as they are assumed to keep a constant prevision horizon of length T , continuously shift their consumption date, leading to a model with a sliding horizon (and no effective consumption). Roughly speaking, in this framework, each of the agents corresponds to an entire family in the overlapping

heterogeneous generations model, and the successive plan revisions coincide with the different plans made by the successive generations. In fact, this alternative approach echoes the seminal work of Lindahl (1939), who observes that “the plans of the economic subjects at any given point of time are neither fully consistent with one another nor with the external conditions, and therefore they must be successively revised.”¹ This is also in line with the more general temporary general equilibrium theory of Grandmont (1977, 2008). Importantly, this continuous plan revision feature ensures that the belief heterogeneity is persistent as these revisions prevent every investor to go extinct. This dual approach, which is fully equivalent to the main model, further helps to construct an intertemporal representative agent, defined as the fictitious agent who, if endowed with the total wealth of the economy, would have a marginal utility equal to the equilibrium price.

I now turn to the implications of the model, and see how they relate to the existing literature. Note that most of the theoretical results are similar to those in Atmaz and Basak (2018) evaluated at $t = 0$, i.e., when the heterogeneity has not started to vanish. Again, one of the theoretical contributions of my study is to show that simple mechanisms can prevent such belief dispersion vanishing to happen.

Looking at the stock price, I infer that it depends positively on the average belief bias, which is in line with the studies of Jouini and Napp (2007) and Kurz and Motolese (2011). I further derive that the belief dispersion impact is positive for sufficiently good states of the world and negative for sufficiently bad ones. Note that the sign of this impact only depends on t through the current state of the world W_t , and not through the remaining time before

¹Erik Lindahl, *Studies in the theory of money and capital*, 1939: p.38. Lindahl also indicates a general procedure to construct a solution consistent with such successive plan revisions:

Starting from the plans and the external conditions valid at the initial point of time, we have first to deduce the development that will be the result of these data for a certain period forward during which no relevant changes in the plans are assumed to occur. Next we have to investigate how far the development during this first period—involving as it must various surprising for the economic subjects—will force them to revise their plans of action for the future, the principles for such a revision being assumed to be included in the data of the problem. And since on this basis the development during the second period is determined in the same manner as before, fresh deductions must be made concerning the plans for the third period, and so on.

the economy ending date as this is the case in a finite horizon setting. Finally, similarly to Atmaz and Basak (2018), I find a convex relation between the stock price and the cashflow news. This price convexity implies that the stock price reacts more to good news than to bad news, and that the stock price reaction to any type of news is stronger in relatively good states. Basu (1997) and Nagel (2005) provide empirical evidence for the first prediction, and, consistent with the second one, Conrad et al. (2002) show that the market responds more strongly to bad news in good times than in bad times. Other theoretical studies derive this convex relation in a model with incomplete common information (Veronesi, 1999), or assuming short-sale constraints (Xu, 2007).

I also study the relation between belief heterogeneity and the stock mean return, and observe that the higher the heterogeneity is, the higher the expected returns are. Thus, I contribute to the literature on this topic which derives conflicting results. More precisely, the positive relation that I document is in line with the conjecture of Williams (1977) that more dispersion of opinion represents more risk, and therefore that agents should be more compensated for holding a riskier asset. Banerjee and Kremer (2010) confirm this predicted positive relation in a dynamic model in which investors disagree on the interpretation of public information, and Buraschi and Jiltsov (2006) derive a similar result linking heterogeneity in beliefs to option open interest. Conversely, another strand of the literature, based on the seminal work of Miller (1977), documents a negative link.² This negative relation critically depends on the presence of market frictions. For example, Chen et al. (2002) obtain this result by developing a model with differences of opinion and short-sales constraints. Alternatively, Atmaz and Basak (2018) theoretically derive that higher dispersion leads to higher returns when the view on the stock is sufficiently pessimistic, and to lower returns when the view is sufficiently optimistic. Coming back to my model's implications, I further obtain that the stock mean return unconditionally decreases with risk aversion. This is because,

²In a model with short-sale constraints and differences of opinion, Miller (1977) argues that the stock is overpriced as it reflects the view of the optimistic agents. In fact, because of the short-sale constraints, pessimistic agents stay out of the market. The higher the differences of opinion, the higher this effect, and therefore the higher the stock overpricing, resulting in lower subsequent returns.

in a heterogeneous economy, more risk averse investors speculate less aggressively, and thus earn lower returns.

Looking at the heterogeneity impacts on the stock volatility, I additionally derive that it monotonically increases with belief dispersion, and that it is higher than the production process volatility. As stated in Atmaz and Basak (2018), this is because higher fluctuations in the average belief bias translate to additional stock price fluctuations, and therefore increase the stock volatility. This monotonic positive relation between belief dispersion and stock volatility is well-documented in the theoretical literature (see, e.g., Shalen, 1993 in a two-period rational expectations model, Scheinkman and Xiong, 2003 in a model with short sale constraints, Buraschi and Jiltsov, 2006 in a model with rational agents with incomplete and heterogeneous information, Andrei et al., 2019 in a model with disagreement on the length of business cycles). I complement these findings by deriving a stationary formula where the heterogeneity effects on volatility remain persistent over time.

Because the formulas are stationary, I then translate the main theoretical implications of the model into testable hypotheses, and, turning to the empirical part of the paper, see if they are verified using real data and running ordinary least squares (OLS) regressions. Note that I focus on market-wide implications because there is only one risky stock available in the model. More precisely, I study if a higher market belief dispersion empirically predicts higher market returns and a higher market volatility. While most studies in this literature focus on monthly data, I also ask whether these relations hold for data computed over longer horizons. These additional hypotheses are thereby directly supported by the non-vanishing belief heterogeneity feature of my model. Lastly, rolling window regressions with quarterly data complete the analysis and allow a more careful study of the time evolution of the belief dispersion predictive ability.

I use analyst monthly forecasts of the earnings-per-share (EPS) long-term growth rate (LTG) of individual stocks from the Institutional Brokers Estimate System (IBES) Unadjusted Summary database from January 1982 to December 2019 as a proxy for investors'

beliefs. Building on Yu (2011), I then obtain a monthly market belief dispersion variable, which is defined as the cross-sectional average of individual stock disagreements. I consider both a value- and an equally-weighted variable, and, as a robustness check, I also construct an alternative monthly belief dispersion variable defined as the standard deviation of individual stock disagreements. The variables for longer horizons (from one quarter to two years) are then obtained by averaging the monthly values over the period of interest. I further use data on individual stock prices from the Center for Research in Security Prices (CRSP) database to construct two market indices (I create a value- and an equally-weighted index), and I compute their simple returns—referred to as the market returns—for various holding periods. Importantly, my empirical analysis differs from other existing ones because I specifically study the returns of indices constituted of the individual stocks used in the construction of the market belief dispersion variables. This allows me to more directly capture the link between the market characteristics and the investor beliefs. Similarly, I construct market volatility data from the daily returns of these market indices.

The empirical tests confirm the predicted positive relation between market disagreement and market returns for most specifications and horizons considered. Thus, considering a model-grounded non-vanishing belief dispersion framework and using more tightly linked data, I show new results on the long-run impacts. They contribute to a large empirical debate, already discussed above from the theoretical point of view. Diether et al. (2002) for instance report that high dispersion stocks earn lower returns. Interestingly, Doukas et al. (2006) replicate their results, and find that the relation becomes positive when controlling for uncertainty in analysts' earnings forecasts. An empirical positive link is also found in, e.g., Anderson et al. (2005) or Banerjee (2011). Other studies derive mixed results or no relation. In particular, Buraschi et al. (2014) find that the relation is ambiguous and leverage-dependent: it is positive and significant for high leverage firms, but can turn negative and non-significant for moderately leveraged firms. Finally, Avramov et al. (2009) find that financial distress drives the negative dispersion effect, and show that it is a facet of non-

investment grade firms which account for less than 5% of the total market capitalization, and that the effect is virtually non-existent otherwise.

Lastly, the results regarding the impacts on the market volatility are more mixed. In fact, while I mostly derive positive coefficients (controlling for lagged volatility), they are not statistically significant. The rolling window regressions further show that the sign and intensity of the belief dispersion impact varies through time, which might explain the weakness of the results. Empirical evidence nevertheless points more towards the approval of the positive model-implied relation than towards its rejection. Note that other empirical works study this link either in the cross-section or using shorter time periods and confirm that it is positive (see, e.g., Ajinkya and Gift, 1985 using data over a 10-month period, Anderson et al., 2005 using monthly data over a 7-year period, Banerjee, 2011 in a cross-sectional analysis).

The paper is organized as follows. Section 2.2 presents the theoretical analysis and translates the main theoretical implications into testable hypotheses. Using empirical data, I then provide a test of these hypotheses in Section 2.3. Section 2.4 concludes. All proofs are reported in Appendix 2.A, and Appendix 2.B contains additional empirical results.

2.2 Theoretical part

This section presents the overlapping heterogeneous generations model and derives its equilibrium. It also describes a dual approach based on a model with a sliding horizon, which allows the construction of an intertemporal representative agent. Lastly, it further contains the theoretical results relative to the stock price, its mean return, and its volatility, and formulates testable hypotheses.

2.2.1 An overlapping heterogeneous generations model

Consider a pure-exchange security market economy evolving in continuous time with an infinite horizon. The economy is populated by overlapping generations of heterogeneous investors who maximize their expected utility from future endowment. They consume part of their wealth and bequeath the remaining part to the subsequent generation.

Uncertainty is modeled by a filtered probability space $(\Omega, F, (F_t), \mathbb{P})$, where Ω is the set of states of nature, F is the σ -algebra of observable events, (F_t) describes how information is revealed through time, and \mathbb{P} is the true probability measure giving the likelihood of occurrence of the different events in F . I assume that there is a single source of risk, modeled by a $((F_t), \mathbb{P})$ -Brownian motion W .

Let y denote the production process in the economy, and assume that, under the probability measure \mathbb{P} , it follows a geometric Brownian motion with drift μ and volatility σ .³

Each generation of investors is constituted of a large number of heterogeneous agents. More precisely, I assume that each generation is composed of a continuum of agents, endowed with a fraction of the production process and having heterogeneous beliefs, who are born at the same date, and who all have a life of length T . These agents have standard CRRA preferences, characterized by $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$ the coefficient of relative risk aversion, and both consume and bequeath their wealth according to some given proportions (common to all investors). Moreover, they disagree on the dynamic of the production process and are characterized by their own subjective beliefs, which give the subjective likelihood of occurrence of the different events in the economy. Formally, these subjective beliefs are indexed by $\delta \in \mathbb{R}$, and, for a given δ , the subjective beliefs are defined by the subjective probability measure \mathbb{Q}^δ , which is assumed to be equivalent to the true probability measure \mathbb{P} . I call Agent- δ of a given generation the agent characterized by these beliefs. Concretely, in all generations, all agents agree on the volatility of the production process σ but disagree

³Note that this drift μ takes into account that part of the production process is consumed at each date as explained further. This is shown formally in the proof of Lemma 2.1.

on its drift.⁴ Instead of considering that it equals μ , Agent- δ believes that the drift of the production process is given by $\mu + \delta$. Thus, δ represents her belief bias, and she is relatively optimistic (resp. pessimistic) compared to an agent with true beliefs if her bias is positive (resp. negative). In fact, Agent- δ believes that the production process is given by

$$dy_t = (\mu + \delta) y_t dt + \sigma y_t dW_{\delta,t},$$

where W_{δ} is a standard unidimensional $((\mathcal{F}_t), \mathbb{Q}^{\delta})$ -Brownian motion, such that $W_{\delta,t} = W_t - \frac{\delta t}{\sigma}$.

As each generation is populated by a continuum of agents, I use probability density functions to describe their wealth share distribution, similarly to, e.g., Beddock and Jouini (2021). At time $t = 0$, I assume that the wealth share distribution of the investors of Generation-0 is given by a Gaussian distribution with parameters $\bar{\delta}$ and $\bar{\omega}$, which are given exogenously. I further show in Proposition 2.1 that, at any time t and in any state of the world W_t , the wealth share distribution of Generation- t —denoted by ν_{δ,t,W_t} —is still given by a normal probability density function with parameters $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$, whose expressions are endogenously determined.⁵

Let now more precisely describe the timeline of a given generation in the model.

At time t , the investors of Generation- t are born and forecast that the production process will deliver a payoff y_{t+T} at time $t + T$. Because their common lifespan is T , they thus make their consumption and bequest plans to consume and bequeath at the end of their life. More formally, at time t , Agent- δ of Generation- t plans to consume an amount $c_{\delta,t+T}$ of her endowment $y_{\delta,t+T}$ at time $t + T$ and to bequeath the remaining part $b_{\delta,t+T} = y_{\delta,t+T} - c_{\delta,t+T}$

⁴This assumption is in line with the literature which shows that the expected return is harder to estimate than the variance (see, e.g., Williams, 1977, Merton, 1980).

⁵Because she inherits later in her life, the time- t wealth share of a given agent of Generation- t more precisely describes this agent's expected share of total endowment at time $t + T$ seen from date t .

to the next generation. Her expected utility is given by

$$\mathbb{E}_t (M_{\delta,t+T} (au(c_{\delta,t+T}) + u(b_{\delta,t+T}))), \quad (2.1)$$

where $M_{\delta,t+T}$ is the Radon-Nikodym derivative of her subjective probability measure \mathbb{Q}^δ with respect to \mathbb{P} and a is an exogenously given non-negative coefficient, common to all agents and all generations, that allows determining the proportions of consumed and bequeathed wealth, as shown in Proposition 2.1. Setting $a = 1$ signifies that the agents equally weight the utility they derive from consumption and bequest, and thus allocate half of their wealth to each. For higher (resp. lower) values, it means that they prefer to consume (resp. bequeath), which reflects some selfishness (resp. altruism).

Next, at time $t + dt$, Agent- δ of Generation- t gives birth to Agent- δ of Generation- $t + dt$. These two agents are therefore part of Family- δ . Note that this intra-family belief transmission assumption, which implies that all members of the same family keep the same beliefs, is key because it allows belief heterogeneity to persist as time goes by.

Finally, at time $t + T$, Agent- δ of Generation- t consumes $c_{\delta,t+T}$, bequeaths $b_{\delta,t+T}$ to Agent- δ of Generation- $t + dt$, and dies.

The Generation- t temporary equilibrium—defined by a continuum of consumption and bequest plans, denoted by $(c_{\delta,t+T})_{\delta \in \mathbb{R}}$ and $(b_{\delta,t+T})_{\delta \in \mathbb{R}}$ respectively, and a positive density price p_{t+T} —is obtained when each agent of Generation- t maximizes her expected utility according to her beliefs such that both her static budget constraint and the market clearing condition are satisfied. Additionally, the sum of the consumption and the bequest of a given agent must not exceed her endowment.

It is worth mentioning how this temporary equilibrium evolves from a generation to another. In fact, at time $t + dt$, agents of Generation- $t + dt$ plan to consume at time $t + dt + T$ a part of the bequeathed wealth of Generation- t that they will inherit later in their lives and to bequeath the remaining part to the subsequent generation. Thus, it leads to the

Generation- $t + dt$ temporary equilibrium, characterized by $(c_{\delta,t+dt+T})_{\delta \in \mathbb{R}}$, $(b_{\delta,t+dt+T})_{\delta \in \mathbb{R}}$, and p_{t+dt+T} . Applying the same reasoning to all subsequent generations results in a continuum of consumption and bequest plans made by all successive generations. Because each generation is born at a different time and in a different state of the world, the successive plans of the members of a given family can differ markedly even if they share the same beliefs.

The intra-family belief transmission assumption further implies an additional dynamic family budget constraint. More precisely, seen by Agent- δ of Generation- t , Family- δ 's expected endowment should remain unchanged between $t + T$ and $t + dt + T$, which means that the expected value of the Family- δ time- $t + dt + T$ endowment should equal the expected value of its adjusted time- $t + T$ endowment evaluated at date t and at the time- $t + dt + T$ price.⁶ Formally, for each time t and each Family- δ , this implies that

$$\mathbb{E}_t(p_{t+dt+T}y_{\delta,t+dt+T}) = \mathbb{E}_t\left(p_{t+dt+T}y_{\delta,t+T}\frac{y_{t+dt+T}}{y_{t+T}}\right). \quad (2.2)$$

Equation (2.2) simply states that the budget of each family evolves according to the evolution of the total production. Solving this equation together with the Generation- t temporary equilibrium equations delivers explicit solutions for the Generation- t wealth share distribution parameters $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$. They are given in Proposition 2.1, where I also report the equilibrium features.

Proposition 2.1 (Equilibrium and wealth share distribution).

In equilibrium, at time t and in state of the world W_t :

1. *The Generation- t investors' consumption and bequest plans for time $t + T$ are given by*

$$c_{\delta,t+T} = \frac{1}{1 + a^{-\frac{1}{\gamma}}}y_{\delta,t+T}, \quad \text{and} \quad b_{\delta,t+T} = \frac{a^{-\frac{1}{\gamma}}}{1 + a^{-\frac{1}{\gamma}}}y_{\delta,t+T},$$

⁶Family- δ 's adjusted time- $t + T$ endowment is given by $y_{\delta,t+T}\frac{y_{t+dt+T}}{y_{t+T}}$ and takes into account the growth of the production process between $t + T$ and $t + dt + T$.

with

$$y_{\delta,t+T} = y_{t+T} \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^{-1} (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}},$$

where $M_{\delta,t+T} = \exp\left(\frac{\delta}{\sigma}(W_{t+T} - W_t) - \frac{1}{2}\frac{\delta^2}{\sigma^2}T\right)$ is the Generation- t Radon-Nikodym derivative of the subjective probability measure \mathbb{Q}^δ with respect to \mathbb{P} , and

$$\lambda_{\delta,t,W_t} = \frac{1}{\sqrt{2\pi}^\gamma} \frac{\sqrt{\sigma^2 + T\varphi(\bar{\omega}_{t,W_t}^2)}^{\gamma-1}}{\sqrt{\varphi(\bar{\omega}_{t,W_t}^2)}^\gamma \sqrt{\sigma^2}^{\gamma-1}} \exp\left(-\frac{(\delta - \bar{\delta}_{t,W_t} + (1-\gamma)T\varphi(\bar{\omega}_{t,W_t}^2))^2}{2\varphi(\bar{\omega}_{t,W_t}^2)}\right)$$

is the inverse of the Generation- t Lagrange multiplier with

$$\varphi(x) = \frac{x}{2} - \frac{\gamma\sigma^2}{2T} + \sqrt{\left(\frac{x}{2} - \frac{\gamma\sigma^2}{2T}\right)^2 + \frac{\sigma^2 x}{T}}.$$

2. The state price density is given by

$$p_{t+T} = y_{t+T}^{-\gamma} \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^\gamma.$$

3. The Generation- t wealth share distribution is given by a normal probability density function with a standard deviation and a mean given respectively by

$$\bar{\omega}_{t,W_t} = \bar{\omega}, \quad \text{and} \quad \bar{\delta}_{t,W_t} = \bar{\delta} + \frac{\varphi(\bar{\omega}^2) W_t}{\sigma},$$

where $\bar{\delta}$ and $\bar{\omega}$ are exogenously given constants describing the Generation-0 wealth share distribution.

A first point to notice is that the higher a is, i.e., the more selfish each generation of investors is, the larger is the share of wealth they consume before they die, and the smaller is the one they bequeath to the subsequent generation. On the other extreme, setting

$a = 0$ means that the investors do not consume and bequeath all their wealth to the next generation.

Additionally, from the third item, I observe that the Generation- t wealth share distribution's standard deviation is constant and does not vanish as time goes by, which ensures that the agents differ persistently in their beliefs. This feature is in sharp contrast with Atmaz and Basak (2018) and other finite-horizon models with one terminal production (and consumption) date, where belief dispersion consistently decreases with time. Such a result is of great empirical importance as surveys show that beliefs are mostly characterized by large and persistent individual heterogeneity (see, e.g., Meeuwis et al., 2019, Giglio et al., 2020, Das et al., 2020). Moreover, as the mean of the wealth share distribution depends positively on W_t , the average belief bias fluctuates around the initial value $\bar{\delta}$. On average, a generation of investors is more optimistic when it is born in a good state of the world, and vice versa.⁷

2.2.2 A dual approach of the model

Alternatively, one can study an equivalent model, which leads to a similar equilibrium and analogous implications, and whose main feature is to consider a continuum of infinitely-lived agents who continuously revise their plans in a sliding horizon framework. While it results in a model without effective consumption—as explained below the investors maximize utility derived from planned consumption but continuously postpone the consumption date, and they thus never consume—the main advantage of this approach is to permit a natural construction of an intertemporal representative agent. In addition, the interpretation and the novelty of further derived results are easier to explain under this alternative view. In this section, I therefore briefly describe the settings of this model which builds on Lindahl (1939)'s theory, before constructing the intertemporal representative agent.

This dual model shares most of its characteristics with the overlapping heterogeneous generations model presented above. Indeed, I still consider a pure-exchange economy, based

⁷More precisely, I say that a generation of investors is more optimistic if a larger share of the expected future total endowment is held by more optimistic agents.

on an expected production process similar to the process y , and evolving in continuous time. Moreover, there is also a continuum of heterogeneous CRRA agents who maximize their expected utility from future (expected) consumption, and I use a normal probability density function with an exogenously given mean $\bar{\delta}$ and standard deviation $\bar{\omega}$ to characterize their initial wealth share distribution. A major difference, however, is that there is now a single generation of infinitely-lived agents who continuously shift their (expected) consumption date and revise their plans as time goes by. Roughly speaking, in this framework, a given Agent- δ , whose beliefs are defined as in Section 2.2.1, corresponds to the entire Family- δ , and her successive plan revisions coincide with the different plans made by the successive generations of this family.

The mechanism behind the successive plan revisions is the following. At time t and in state of the world W_t , the investors forecast that the expected production process will deliver a payoff at time $t + T$, and they thus plan to consume at this date. The parameter T can therefore be associated to the agents' prevision horizon or, more generally, to the agents' investment horizon. At time $t + dt$, the market reopens and new information W_{t+dt} comes in. Assuming that their prevision horizon is fixed, the agents now forecast that the expected production process will deliver a payoff at time $t + dt + T$. They therefore shift the date of their expected consumption, and update their plans because of the new information W_{t+dt} . The same mechanism applies to all successive dates and, as time goes by, this leads to a continuum of temporary equilibria—defined by a continuum of expected consumption plans and a state price density—with a sliding horizon, where consumption plans are continuously revised, and where effective consumption never occurs as it is continuously postponed. It should be noted that assuming that the agents continuously forecast that the (expected) production process will deliver a payoff after a duration T is a key ingredient of this model. In fact, it ensures to maintain a persistent belief heterogeneity. More precisely, the investors keep their beliefs and consider a similar maximization program at each date, but in a different state of the world, where they plan to consume after a period of length T (under an additional

dynamic budget constraint similar to the one defined by (2.2)).

Let now properly state the equivalence of the two approaches in the following lemma.

Lemma 2.1. *The sliding horizon model is equivalent to the overlapping heterogeneous generations model under the assumption that the successive generations only bequeath and do not consume (i.e., $a = 0$). In case effective consumption is allowed, the two models still yield analogous implications up to a small adjustment in the drift of the underlying production processes.*

As stated before, the main interest of considering this alternative equivalent approach is that it allows a natural definition and construction of the intertemporal representative agent. Such an agent is constructed as the fictitious (infinitely-lived) agent who, if endowed with the total wealth of the economy, would have a marginal utility equal to the equilibrium price. She has the same utility function as the other investors and is characterized by the Radon-Nykodym derivative of her subjective probability measure with respect to the true probability measure.

Proposition 2.2 (Representative agent). *The intertemporal representative agent of the economy is the fictitious investor whose time- $t + T$ Radon-Nikodym derivative of the subjective probability measure \mathbb{Q}^δ with respect to \mathbb{P} seen from date t is given by*

$$M_{RA,t+T} = \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^\gamma,$$

where $M_{\delta,t+T}$ and λ_{δ,t,W_t} are defined as in Proposition 2.1.

On the technical side, this result, combined with the equivalence of the two models, is important because it allows the complete characterization of the inverse of the Lagrange multipliers and thus the full determination of the models' equilibrium.

2.2.3 The stock price and its dynamics

I now derive the stock price and its dynamics in the presence of belief heterogeneity. As stated in the previous section, the results are derived under the view of the alternative model with a sliding horizon, which eases the interpretation and the comparison with the existing literature. For the sake of clarity, and because both models are equivalent, I nevertheless use the notations defined in Section 2.2.1

I assume that a risky stock S is available for trading. The stock is in positive net supply of one unit and, at time t , is a claim to the payoff y_{t+T} expected to be paid at time $t + T$. Studying its properties leads to the following proposition.

Proposition 2.3 (Equilibrium stock price, mean return, and volatility).

In equilibrium, at time t and in state of the world W_t :

1. *The stock price is given by*

$$S_t = \bar{S}_t \exp \left(\bar{\delta}_{t,W_t} T - \frac{\varphi(\bar{\omega}^2) T^2}{2} \right),$$

2. *The mean stock return is given by*

$$\mu_{S_t} = \bar{\mu} + \sigma \left(\frac{\partial \bar{\delta}_{t,W_t} T}{\partial W_t} \right) + \frac{1}{2} \left(\frac{\partial \bar{\delta}_{t,W_t} T}{\partial W_t} \right)^2,$$

3. *The stock volatility is given by*

$$\sigma_{S_t} = \bar{\sigma} + \frac{\partial \bar{\delta}_{t,W_t} T}{\partial W_t},$$

where $\bar{S}_t = y_t \exp((\mu - \gamma\sigma^2) T)$, $\bar{\mu} = \mu$, and $\bar{\sigma} = \sigma$ are the equivalent quantities obtained in a similar standard economy without belief heterogeneity.

The formulas share similarities with those in Atmaz and Basak (2018). However, the two models differ along one important dimension: unlike them, I design an economy where

the effective consumption date is continuously postponed so that the time- t remaining time before consumption always equals T instead of $T - t$. As the time- t heterogeneity impacts depend on the remaining time before consumption, I thus obtain a stationary model where the heterogeneity effects are not smoothed as time goes by. In particular, as implied by their framework, their time- T stock price formula reduces to $S_T^{AB} = y_T$. The heterogeneity effects that they observe thus completely vanish when approaching date T , and their model is not stationary. Conversely, my time- T stock price is still fully impacted by the time- T investors heterogeneity, characterized by $\bar{\delta}_{T,W_T}$ and $\bar{\omega}$.

The results stationarity and persistence are of first interest as they more accurately reflect the heterogeneous market participants reality. These features also allow me to test the model empirically over a long period and to use data sampled at various frequencies. Before doing so in Section 2.3, I now discuss the properties of the market characteristics derived in Proposition 2.3. I first focus on the equilibrium mean stock return μ_S .

As stated before, one of the consequences of the sliding horizon methodology is that the time- t remaining time before expected consumption does not depend on t . Thus, the derivative of the stock price with respect to t differs markedly from the one obtained in Atmaz and Basak (2018), leading to different mean returns.

I find that a higher belief dispersion leads to a higher equilibrium mean stock return. More precisely, the belief dispersion has an impact on the sensitivity of the average belief bias to news: the higher the belief dispersion, the higher this sensitivity, and thus the higher the mean return. This is consistent with the recent work of Brandon and Wang (2020), who find that the average return on stocks with high sensitivity to earning belief shocks is 7.14% per year higher than that in stocks with low sensitivity. Conversely, Atmaz and Basak (2018) derive that the sign of the relation is state-dependent, and that higher dispersion leads to lower returns when the view is sufficiently optimistic. They further show that the relation between the mean return and the relative risk aversion coefficient depends on the level of optimism, while I derive an unambiguous negative relation. The intuition behind

this result is simple: in an economy populated by heterogeneous agents, more risk averse agents speculate less aggressively, and thus earn lower returns. Lastly, I observe that the mean equilibrium stock return increases as the investment horizon T increases.

For the sake of completeness, I now briefly report the properties of the first and third items of Proposition 2.3, which are mostly similar to those in Atmaz and Basak (2018) except that the heterogeneity effects remain persistent. I refer the reader to their paper for more detailed explanations of the underlying mechanisms behind these results.

Specifically, for a given time t , I find analogous impacts of the belief distribution parameters on the stock price. First, the stock price depends positively on the time- t average belief bias $\bar{\delta}_{t,W_t}$. Second, the sign of the belief dispersion impact is state-dependent: the impact is positive for sufficiently good states of the world and negative for sufficiently bad ones.⁸ I also similarly derive that the stock price is convex in the time- t expected production level y_t . Finally, unlike the standard case, the impact of the coefficient of relative risk aversion γ is not always negative, but can be positive for sufficiently bad states of the world.⁹

The third item considers the stock volatility σ_S . Several observations are in order. First, in a heterogeneous economy, it is higher than the production process volatility σ , in line with empirical observations (see, e.g., Ajinkya and Gift, 1985, Anderson et al., 2005, and Banerjee, 2011). Recall that an important difference with Atmaz and Basak (2018) is that, although both formulas have the same shape, this excess volatility effect does not decrease as time goes by. This is because the investors' heterogeneity remains persistent. Second, all else equal, the higher the belief dispersion, the higher the state-sensitivity of the average belief

⁸Two effects are at play. On the one hand, as the function φ is increasing in belief dispersion, there is a direct negative effect. On the other hand, there is an indirect effect of the belief dispersion through the average belief bias. For bad states, this indirect effect is negative and reinforces the first effect. For good ones, this is the opposite, and the overall effect can even be positive in case of sufficiently good states. Formally, as the average belief bias is state-dependent, I derive that the stock price increases in belief dispersion when

$$\bar{\delta}_{t,W_t} > \bar{\delta} + \frac{\varphi(\bar{\omega}^2)T}{2}.$$

⁹Formally, the stock price increases in risk aversion when the following condition holds

$$\bar{\delta}_{t,W_t} < \bar{\delta} + \frac{\varphi(\bar{\omega}^2)T}{2} - \sqrt{(\gamma\sigma^2 - T\bar{\omega}^2)^2 + 4T\sigma^2\bar{\omega}^2}.$$

bias, and thus the higher the excess volatility. In fact, a higher fluctuation in the average belief bias translates to additional stock price fluctuations, and therefore increases the stock volatility. Additionally, the coefficient of relative risk aversion has a negative impact on σ_S . Finally, a higher investment horizon T leads to a higher stock volatility.

Overall, a rich set of predictions can be derived from Proposition 2.3. In particular, it allows considering the impacts of belief heterogeneity on the mean stock return and on the stock volatility. Moreover, as observed above, the impact of the belief bias on the model characteristics is indirect and depends on the belief dispersion, which therefore appears to be the relevant predictor to consider. The persistence of belief dispersion that I obtain—which is not a feature of other existing models—also allows me to use data sampled at various frequencies, leading to new testable implications. In sum, these theoretical results can be converted into testable hypotheses stated below that are empirically tested in the next section.

Hypothesis. *The main testable implications of the model are the following:*

H1. A higher belief dispersion predicts higher mean stock returns.

H2. The ability of belief dispersion to predict mean stock returns remains over long horizons.

H3. A higher belief dispersion predicts a higher stock volatility.

H4. The ability of belief dispersion to predict the stock volatility remains over long horizons.

2.3 Empirical test of the model

This section provides an empirical test of the hypotheses $H1-H4$. As the model considers an economy where there is only one risky stock available for trading, it is more relevant to use an index—referred to as the market index and whose definition is given below—rather than individual stocks. I thus focus on studying the empirical ability of market belief dispersion

to predict future market index returns and volatility over time, using variables sampled at various frequencies (from one month to two years) from January 1982 to December 2019.

2.3.1 Market belief dispersion data

Let me first explain how the data on market belief dispersion, expressed in percentages, is constructed. I use analyst monthly forecasts of individual stocks as a proxy for investors' beliefs, and I aggregate them using different methodologies (described below) to obtain market belief dispersion variables. As underlined by Yu (2011), this bottom-up approach has the advantage of taking into account hundreds of forecasts at any given time, and thus likely has a good signal-to-noise ratio. The large number of forecasts also echoes the choice of considering a continuum of agents in the theoretical model. Additionally, using forecasts of individual stocks allows defining and studying the returns and volatility of a market index based on these stocks, which ensures a direct link between beliefs and market characteristics.

The data comes from two databases: I use the analyst monthly forecasts of the EPS LTG of individual stocks from the IBES database,¹⁰ and the CRSP database to obtain monthly market capitalizations. The IBES data is winsorized at the 1% and 99% levels to account for potential outliers or data errors. I also winsorize the prices at the 99% level. Furthermore, I exclude stocks whose price is below five dollars at portfolio formation to avoid that extreme returns on penny stocks drive the results and stocks for which less than two analysts provide EPS LTG forecasts during the month to focus on stocks that exhibit some forecast dispersion. The data is available from January 1982 to December 2019. Throughout the sample, the average number of stocks used to compute the variable is 965, and each stock at any given time is followed on average by five to six analysts.¹¹ This large number of stocks alleviates

¹⁰I thus use earnings data to measure cashflows rather than dividends data. This choice is motivated by Da (2009), who argues that potential problems of working with dividends could arise because of the dividend payout policy of some firms. Campbell (2000) further highlights other empirical difficulties. On the theoretical side, using the accounting clean surplus identity, Vuolteenaho (1999) shows that if one looks at the infinite horizon, cash flow and earnings contain the same information. Thus, earnings are both theoretically equivalent and empirically better-behaved than dividends.

¹¹The average number of analysts who provide an EPS LTG forecast for any given stock at any given

the possibility that idiosyncratic firm disagreement drives the variations of the market belief dispersion.

For each common stock i listed on the NYSE/Amex/Nasdaq in each month t that meets the above-mentioned requirements, I obtain the standard deviation of the analyst forecasts—that I refer to as the stock disagreement and that I denote $\tilde{\omega}_{i,t}$ —from the IBES Unadjusted Summary database.¹² Additionally, I obtain the market capitalization of each of these stocks at the end of each month—that I denote $MKTCAP_{i,t}$ —using the closing price and the number of shares outstanding of the stock considered from the CRSP database.

With this data in hand, I construct my first value-weighted measure of monthly market belief dispersion, $\bar{\omega}_{mean,1M}^{VW}$, which is similar to the one defined in Yu (2011). For a given month t , it is defined as the cross-sectional (value-weighted) average of individual stock disagreements

$$\bar{\omega}_{mean,1M,t}^{VW} = \frac{\sum_i MKTCAP_{i,t} \times \tilde{\omega}_{i,t}}{\sum_i MKTCAP_{i,t}}.$$

To further rule out the possibility that the market belief dispersion is driven by idiosyncratic firm disagreement, I also consider an alternative value-weighted monthly measure: the cross-sectional value-weighted standard deviation of individual stock disagreements, $\bar{\omega}_{std,1M}^{VW}$. A larger dispersion of individual stock disagreements indeed likely reflects a higher market disagreement among investors. Letting N_t denote the number of stocks that meet the requirements in month t , it is obtained with the following formula

$$\bar{\omega}_{std,1M,t}^{VW} = \frac{\sum_i MKTCAP_{i,t} \times (\tilde{\omega}_{i,t} - \bar{\omega}_{mean,1M,t}^{VW})^2}{\frac{N_t-1}{N_t} \sum_i MKTCAP_{i,t}}.$$

Lastly, for the sake of robustness, I construct similar market belief dispersion variables, denoted by $\bar{\omega}_{mean,1M}^{EW}$ and $\bar{\omega}_{std,1M}^{EW}$ respectively, whose only difference with $\bar{\omega}_{mean,1M}^{VW}$ and $\bar{\omega}_{std,1M}^{VW}$ is to use equal-weighting rather than value-weighting.

time is 5.33.

¹²Similarly to Buraschi et al. (2014), I use unadjusted data to circumvent the problem of using stock-split adjusted data.

The variables for longer horizons (one quarter ($3M$), six months ($6M$), one year ($12M$), and two years ($24M$)) are then obtained by averaging the monthly values over the period of interest. In the remainder of the analysis, I therefore use those market belief dispersion variables defined over various horizons as the predictors, and examine different specifications depending on the horizon considered and the type of weighting used in the dependent variable.

Insert Figure 2.1 here.

To give an example of how the market belief dispersion variables vary as time goes by, Figure 2.1 shows the time series of the quarterly variables. I also report their summary statistics over the full sample in Table 2.1.

Insert Table 2.1 here.

2.3.2 Predicting market returns

To begin with, I test the hypotheses $H1-H2$. Formally, I therefore use empirical data to see if a higher market belief dispersion leads to higher market returns, and if this positive model-implied relation holds for data sampled at long frequencies.

Let first properly define the value-weighted market index whose returns are used in the analysis. Each month, it is constituted of all individual stocks, weighted by their market capitalization, whose price is above five dollars and for which at least two monthly EPS LTG forecasts are provided in the IBES database. In other words, the assets that constitute this market index are those used to construct the market belief dispersion variables, which allows me to more precisely capture the link between the investor beliefs and the market characteristics. In addition, I construct a similar market index using equal weights. I then compute the (raw) simple returns for various holding periods of both indices, and report their summary statistics, expressed in percentages, in Table 2.2, where the subscript indicates the holding period considered and the superscript the type of weighting.

Insert Table 2.2 here.

I can now test the model-implied hypotheses *H1-H2*. More specifically, *H1* implies that a higher market belief dispersion in a given period should result in higher market returns in the subsequent period. If *H2* is verified, this positive relation should hold no matter the different horizon considered (from one month to two years). In order to check these hypotheses, I thus run the following standard OLS regression

$$RET_{i,t}^k = \gamma_{j,i}^k + \theta_{j,i}^k \bar{\omega}_{j,i,t-1}^k + \xi_{j,i,t}^k, \quad (2.3)$$

where t refers to the period t , $i = \{1M, 3M, 6M, 12M, 24M\}$, $j = \{mean, std\}$, and $k = \{VW, EW\}$. Inference is based on autocorrelation- and heteroskedasticity-robust standard errors (Newey and West, 1987), and all variables are standardized prior to estimation. Moreover, I consider non-overlapping returns for horizons longer than a month to avoid econometric issues. This leads to a total of 20 specifications, whose results are reported in Table 2.3.

Insert Table 2.3 here.

Several observations are in order. First, all coefficients are positive. More interestingly, most of them are statistically different from zero.¹³ Thus, *H1* seems to be verified in the data. This complements the mixed results regarding the impact of belief dispersion on returns found in the literature. While some studies also derive a positive relation (see, e.g., Doukas et al., 2006), others find that a higher belief dispersion predicts lower returns (see, e.g., Diether et al., 2002). The novelty in my empirical study is that I focus on the returns of two indices (value- and equally-weighted) that only contain stocks for which there is some belief dispersion, and that are thus used in the construction of the market belief dispersion data.

¹³Note that, because the model implies a positive relation, I consider one-sided tests where the alternative hypothesis is that the coefficient is positive.

Second, I observe that the model-implied positive relation holds for all horizons, meaning that the effects of belief dispersion do not vanish over long periods. In other words, the hypothesis $H2$ is verified, and the belief dispersion persistence that I document in my model is an important feature to be taken into account.

Lastly, note that since the model abstracts from interest rate issues (both under the overlapping generations view or the sliding horizon one), and thus does not allow an accurate definition of the risk premium, it is mostly suitable to study the ability of market belief dispersion to predict the market index raw returns. Although the link is less direct, I nevertheless run a similar analysis using market index excess returns in Appendix 2.B.1, which confirms the previous results.

2.3.3 Predicting market volatility

I now focus on the hypotheses $H3$ - $H4$ and study if they are verified empirically. Hence, I test if market belief dispersion positively predicts future market volatility, and if this relation holds for various horizons.

As commonly done in the literature (see, e.g., French et al., 1987, Schwert, 1989), I exploit daily stock returns—taken from the CRSP database—to obtain my market volatility data. More precisely, to obtain value- and equally-weighted variables for a given horizon, I consider the two market indices defined in Section 2.3.2, compute the sum of their squared daily returns over the period of interest (discounted by the average daily market index returns of the period), and thereby obtain the value- and equally-weighted market variances for this horizon. I then easily convert them into the annualized market volatilities, that I denote VOL^{VW} and VOL^{EW} respectively. As in the previous part, I further add subscripts to these variables to indicate the horizon considered.

Computing their descriptive statistics shows that they are highly positively skewed and leptokurtic.¹⁴ Similarly to Paye (2012), I thus define annualized market log volatility vari-

¹⁴For instance, skewness (VOL_{3M}^{VW}) = 2.89, and kurtosis (VOL_{3M}^{VW}) = 15.43.

ables, $LVOL^{VW}$ and $LVOL^{EW}$, as the natural logarithms of VOL^{VW} and VOL^{EW} respectively, whose distribution are approximately Gaussian (Andersen et al., 2001). Because the empirical analysis relies on linear models estimated by OLS regressions, this latter property is of first importance, and I therefore use $LVOL^{VW}$ and $LVOL^{EW}$ as the independent variables in the subsequent empirical analysis. I report their summary statistics in Table 2.4.

Insert Table 2.4 here.

Volatility processes are known to exhibit a high degree of persistence when the data is sampled over short periods of time. To take this persistence into account, and to ensure that a potential predictor contains valuable information, it is thus important to control for past (log) volatility when trying to predict future (log) volatility. Paye (2012) for instance includes six lagged values of (log) volatility in his monthly specification and two lagged values in his quarterly one. I thus adopt a similar approach when testing if a higher market belief dispersion in a given period predicts a higher market (log) volatility in the subsequent one. More precisely, I use six lags when dealing with monthly data, two lags when dealing with quarterly data, and one lag when dealing with data sampled every six months. For the longest horizons considered (one and two years), I do not control for past (log) volatility because the data is sampled over relatively long periods of time and is not strongly persistent. Formally, for, e.g., the quarterly specifications, I therefore consider the following regression for market (log) volatility

$$LVOL_{3M,t}^k = \alpha_{0,j,3M}^k + \alpha_{1,j,3M}^k LVOL_{3M,t-1}^k + \alpha_{2,j,3M}^k LVOL_{3M,t-2}^k + \beta_{j,3M}^k \bar{\omega}_{j,3M,t-1}^k + \epsilon_{j,3M,t}^k \quad (2.4)$$

where t refers to quarter t , $j = \{mean, std\}$, and $k = \{VW, EW\}$.

The main interest of (2.4) and of other similar OLS regressions studying other horizons is to test the hypothesis $H_0 : \beta = 0$ against the alternative $H_1 : \beta > 0$: rejecting the null indeed implies that, when controlling for past (log) volatility, belief dispersion positively predicts future market (log) volatility. Similarly to the previous part, I consider one-sided tests

because the predicted sign of the coefficient is supported by the theoretical implications of the model. Again, inference is based on autocorrelation- and heteroskedasticity-robust standard errors, and all variables are standardized prior to estimation. The results are reported in Table 2.5.

Insert Table 2.5 here.

While most coefficients are positive, only those related to the equally-weighted specification for variables sampled at a yearly frequency are statistically significant. Thus, over the full sample, neither $H3$ nor $H4$ seem to be strongly validated by the data.

One of the potential explanations for this lack of statistical significance could be that the sign of the relation varies through the sample period considered. To further investigate this issue, I run rolling window regressions of Equation (2.4) using subsamples of 15 years.¹⁵ For the sake of concision, I focus on data sampled at the quarterly frequency, which leads to a total of 91 regressions for each specification.¹⁶ I therefore obtain time series of the estimated values of $\beta_{mean,3M}^{VW}$, $\beta_{std,3M}^{VW}$, $\beta_{mean,3M}^{EW}$, and $\beta_{std,3M}^{EW}$, and report them in Figure 2.2. In each panel, the horizontal axis shows the end date of the subsamples, and thicker rounds (resp. crosses) indicate statistically significant positive (resp. negative) values at the 10% level.

Insert Figure 2.2 here.

As reported in Panel A of Table 2.6, between 52 and 83 of the 91 estimated β 's are positive depending on the market belief dispersion variable considered. More interestingly, around a third of these estimated values are statistically significant at the 10% level. Moreover, none of the negative coefficients obtained from the value-weighted variables is significantly different from zero. The evidence therefore points towards the approval of the model-implied

¹⁵Conrad and Glas (2018) provide a similar analysis to test if macroeconomic variables predict volatility in the cross-section of industry portfolios. Note that, for the sake of completeness, I also provide in Appendix 2.B.2 an analogous rolling window analysis concerning the market index quarterly raw returns to see how the ability of market belief dispersion to predict them evolves over time.

¹⁶Because I need to have the market (log) volatility data of the two previous quarters, my first subsample starts in the third quarter of 1982 and ends in the second quarter of 1997.

positive effect of belief dispersion on the market volatility, and, while most of the significantly positive β 's are obtained for subsamples ending between 2000 and 2010, the weak results found in Table 2.5 likely result from the fact that negative coefficients are obtained from the most recent subsamples.

Insert Table 2.6 here.

As robustness checks, Panel B and Panel C of Table 2.6 further show results of alternative specifications that use rolling windows of 10 and 20 years respectively. Overall, they confirm the results found in the main specification.

2.4 Conclusion

In this paper, I define an infinite horizon economy populated by overlapping generations of investors who differ in their beliefs. For any generation, the wealth share distribution of its heterogeneous members can be described by a Gaussian distribution with a state-dependent mean and a constant standard deviation, implying a persistent belief heterogeneity over time. I compute the model equilibrium, and, using a fully equivalent approach based on continuous plan revisions of infinitely-lived agents and a sliding horizon, construct an intertemporal representative agent. I then study the implications of the belief heterogeneity on various quantities of interest, namely the stock price, its mean return, and its volatility. In particular, I derive that both the stock mean return and volatility monotonically increase with belief dispersion. Importantly, the theoretical framework that I use leads to stationary results with non-vanishing heterogeneity. The contribution of such a modeling is twofold. First, the heterogeneity persistence is empirically observed, and, consequently, the model results are more in line with reality (see, e.g., Giglio et al., 2020). Second, it allows me to test the model empirically over long periods. Using analyst forecasts from the IBES database, I show that the documented positive relation between the returns and the belief dispersion is verified in the data when considering a market index. This is true for data sampled at various

frequencies (from one month to two years), and the persistence of belief dispersion that I document in my model is thus an important feature to be taken into account. Additionally, empirical evidence points more towards the approval of the positive model-implied relation between the market volatility and the market belief dispersion than towards its rejection.

Lastly, note that the model only considers a single stock in the economy. It would thus be interesting to extend it to the case of a multi-stocks economy to derive testable cross-sectional relations. I leave this for future research.

2.A Proofs

It is more convenient to formally compute the model equilibrium and the representative agent under the dual approach based on the model presented in Section 2.2.2. Thus, in this appendix, I first prove Lemma 2.1, which states the equivalence of the two models. I then state and prove Theorem 2.1 (which is mostly an analogous version of Proposition 2.1 in the context of the sliding horizon model). Proposition 2.1 and Proposition 2.2 are then corollaries of this theorem. Lastly, I provide a proof of Proposition 2.3.

Proof of Lemma 2.1 The equivalence of the two models comes from the fact that they share similar settings, an analogous maximization program, and the same constraints. To see this, let formally present Agent- δ 's maximization program in the context of both models at time t and in state of the world W_t .

Let start with the overlapping heterogeneous generations model. Using (2.1), the maximization program of Agent- δ at time t and in state of the world W_t is given by

$$\max_{c_{\delta,t+T}, b_{\delta,t+T}} \mathbb{E}_t (M_{\delta,t+T} (au(c_{\delta,t+T}) + u(b_{\delta,t+T}))), \quad (2.5)$$

where a is a given non-negative coefficient, common to all agents and all generations, that represents the agents' degree of selfishness, and $M_{\delta,t+T}$ is the Radon-Nikodym derivative of her subjective probability measure \mathbb{Q}^δ with respect to \mathbb{P} , such that

$$M_{\delta,t+T} = \exp \left(\frac{\delta}{\sigma} (W_{t+T} - W_t) - \frac{1}{2} \frac{\delta^2}{\sigma^2} T \right). \quad (2.6)$$

Denoting by $y_{\delta,t+T}$ the endowment (at time $t+T$) of Agent- δ of Generation- t , solving (2.5) under the condition that the sum of her consumption and her bequest equals her endowment ($c_{\delta,t+T} + b_{\delta,t+T} = y_{\delta,t+T}$) leads to

$$c_{\delta,t+T} = \frac{1}{1 + a^{-\frac{1}{\gamma}}} y_{\delta,t+T}, \quad \text{and} \quad b_{\delta,t+T} = \frac{a^{-\frac{1}{\gamma}}}{1 + a^{-\frac{1}{\gamma}}} y_{\delta,t+T}. \quad (2.7)$$

Because $au(c_{\delta,t+T}) + u(b_{\delta,t+T}) = a\left(1 + a^{-\frac{1}{\gamma}}\right)^\gamma u(y_{\delta,t+T})$, maximizing (2.5) is therefore equivalent to solving

$$\max_{y_{\delta,t+T}} \mathbb{E}_t (M_{\delta,t+T} u(y_{\delta,t+T})), \quad (2.8)$$

under the same market clearing condition and budget constraint, before dividing it into consumption and bequest according to the above-mentioned proportions.

The market clearing condition and the budget constraint that both (2.5) and (2.8) need to satisfy are respectively given by

$$y_{t+T} = \int y_{\delta,t+T} d\delta, \quad (2.9)$$

and

$$(\nu_{\delta,t,W_t} y_{t+T} - y_{\delta,t+T}) p_{t+T} \geq 0. \quad (2.10)$$

Additionally, (2.2) needs to be verified for all families.

Let now more formally present the model with a sliding horizon. I denote by \tilde{y} the expected production process of this model. It follows a geometric Brownian motion with the same volatility σ as the process y and a drift denoted by $\tilde{\mu}$. I also define $\tilde{\nu}_{\delta,t,W_t}$ which stands for the time- t wealth share distribution of the infinity lived agents. Recall that the initial investors' wealth share distributions are the same in the two models. Recall also that, as stated in Section 2.2.2, the way heterogeneous beliefs are defined are identical in both models: in the sliding horizon model, $M_{\delta,t+T}$ is defined as in (2.6).

At time t and in state of the world W_t , Agent- δ 's maximization program is given by

$$\max_{\tilde{y}_{\delta,t+T}} \mathbb{E}_t (M_{\delta,t+T} u(\tilde{y}_{\delta,t+T})), \quad (2.11)$$

under the following constraints:

- the market clearing condition: $\tilde{y}_{t+T} = \int \tilde{y}_{\delta,t+T} d\delta$.
- the time- t static budget constraint: $(\tilde{\nu}_{\delta,t,W_t} \tilde{y}_{t+T} - \tilde{y}_{\delta,t+T}) \tilde{p}_{t+T} \geq 0$, where \tilde{p}_{t+T} is the time- $t + T$ equilibrium state price density.
- the time- t dynamic budget constraint: seen from date t , and before trading and consumption reallocation, Agent- δ 's expected consumption share should remain unchanged between $t + T$ and $t + T + dt$ because the endowment of each agent is fixed. Formally, it implies that $\mathbb{E}_t(\tilde{p}_{t+T+dt} \tilde{y}_{\delta,t+T+dt}) = \mathbb{E}_t\left(\tilde{p}_{t+T+dt} \tilde{y}_{\delta,t+T} \frac{\tilde{y}_{t+T+dt}}{\tilde{y}_{t+T}}\right)$.

It is then clear that (2.8) and (2.11) are similar. Moreover, the three above-mentioned constraints are similar to (2.9), (2.10), and (2.2) respectively. This thus leads to equivalent implications for the two models. Note, however, that the alternative model with a sliding horizon does not allow effective consumption. A small adjustment thus needs to be made in the drift of the production process \tilde{y} to obtain a full equivalence.

Let come back to the overlapping heterogeneous generations model. At time $t + dt + T$, the total endowment of the economy is given by $\frac{a^{-\frac{1}{\gamma}}}{1 + a^{-\frac{1}{\gamma}}} y_{t+dt+T}$ (because all the agents of Generation- t have consumed a fraction $\frac{1}{1 + a^{-\frac{1}{\gamma}}}$ of their individual time- $t + T$ endowments, given that a is common to all investors and all generations). Denoting by $\hat{\mu}$ the drift of the production process in this model if there was no consumption, and defining $\frac{1}{1 + a^{-\frac{1}{\gamma}}} = \alpha dt$, the overall endowment process of the model is therefore given by $dy_t = (\hat{\mu} - \alpha) y_t dt + \sigma y_t dW_t$. As stated in footnote 3, I thus define $\mu = \hat{\mu} - \alpha$ to take this consumption into account.

From there, it is clear that setting $\tilde{\mu} = \hat{\mu} - \alpha$ allows a full equivalence of the two models. Finally, note that the case $a = 0$ (meaning that agents only bequeath) implies $\alpha = 0$, and thus $\tilde{\mu} = \mu = \hat{\mu}$. ■

Theorem 2.1 (Equilibrium and representative agent in the model with a sliding horizon).

In equilibrium, at time t and in state of the world W_t :

1. The investors' consumption plans and the state price density are given by

$$\begin{aligned}\tilde{y}_{\delta,t+T} &= \tilde{y}_{t+T} \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^{-1} (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}}, \\ \tilde{p}_{t+T} &= \tilde{y}_{t+T}^{-\gamma} \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^{\gamma},\end{aligned}$$

where $M_{\delta,t+T} = \exp\left(\frac{\delta}{\sigma}(W_{t+T} - W_t) - \frac{1}{2}\frac{\delta^2}{\sigma^2}T\right)$ is the time- $t + T$ Radon-Nikodym derivative of the subjective probability measure \mathbb{Q}^δ with respect to \mathbb{P} (seen from date t), and

$$\lambda_{\delta,t,W_t} = \frac{1}{\sqrt{2\pi}^\gamma} \frac{\sqrt{\sigma^2 + T\varphi(\bar{\omega}_{t,W_t}^2)}^{\gamma-1}}{\sqrt{\varphi(\bar{\omega}_{t,W_t}^2)}^\gamma \sqrt{\sigma^2}^{\gamma-1}} \exp\left(-\frac{(\delta - \bar{\delta}_{t,W_t} + (1-\gamma)T\varphi(\bar{\omega}_{t,W_t}^2))^2}{2\varphi(\bar{\omega}_{t,W_t}^2)}\right)$$

is the inverse of the time- t Lagrange multiplier with

$$\varphi(x) = \frac{x}{2} - \frac{\gamma\sigma^2}{2T} + \sqrt{\left(\frac{x}{2} - \frac{\gamma\sigma^2}{2T}\right)^2 + \frac{\sigma^2 x}{T}}.$$

2. The intertemporal representative agent of the economy is the fictitious investor whose time- $t + T$ Radon-Nikodym derivative of the subjective probability measure \mathbb{Q}^δ with respect to \mathbb{P} (seen from date t) is given by

$$M_{RA,t+T} = \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^\gamma.$$

3. The wealth share distribution of the continuum of investors is given by a normal probability density function with a standard deviation and a mean given respectively by

$$\bar{\omega}_{t,W_t} = \bar{\omega}, \quad \text{and} \quad \bar{\delta}_{t,W_t} = \bar{\delta} + \frac{\varphi(\bar{\omega}^2) W_t}{\sigma},$$

where $\bar{\delta}$ and $\bar{\omega}$ are given constants describing the time-0 wealth share distribution.

Proof of Theorem 2.1 The three items of the proposition are solved altogether in the following proof.

The maximization program of Agent- δ at time t and in state of the world t is given by (2.11) in the proof of Lemma 2.1, which also explicitly states the constraints that must be satisfied.

The first order conditions directly give

$$\begin{cases} \tilde{y}_{\delta,t+T} &= \tilde{y}_{t+T} \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^{-1} (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}}, \\ \tilde{p}_{t+T} &= \tilde{y}_{t+T}^{-\gamma} \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^{\gamma}, \end{cases}$$

where λ_{δ,t,W_t} is the inverse of the time- t Lagrange multipliers of the form

$$\lambda_{\delta,t,W_t} = K_{t,W_t} \exp \left(-\frac{1}{2a_{t,W_t}^2} \delta^2 - b_{t,W_t} \delta \right). \quad (2.12)$$

To further identify a_{t,W_t} and b_{t,W_t} , notice that the budget constraint of the maximization program can be equivalently formulated as

$$\tilde{\nu}_{\delta,t,W_t} = \frac{\mathbb{E}_t(\tilde{p}_{t+T} \tilde{y}_{\delta,t+T})}{\mathbb{E}_t(\tilde{p}_{t+T} \tilde{y}_{t+T})}. \quad (2.13)$$

Let then look for a solution of the form $\tilde{\nu}_{\delta,t,W_t} = \frac{1}{\sqrt{2\pi\bar{\omega}_{t,W_t}}} \exp \left(-\frac{(\delta - \bar{\delta}_{t,W_t})^2}{2\bar{\omega}_{t,W_t}^2} \right)$, and see if it verifies (2.13).

Explicit computations show that it is indeed the case, which thus proves that the time- t wealth share distribution is given by a normal probability density function. Formally, by identification, simple algebra leads to

$$\begin{cases} a_{t,W_t}^2 &= \frac{\bar{\omega}_{t,W_t}^2}{2} - \frac{\gamma\sigma^2}{2T} + \sqrt{\left(\frac{\bar{\omega}_{t,W_t}^2}{2} - \frac{\gamma\sigma^2}{2T} \right)^2 + \frac{\sigma^2\bar{\omega}_{t,W_t}^2}{T}}, \\ b_{t,W_t} &= (1 - \gamma) T - \frac{\bar{\delta}_{t,W_t}}{a_{t,W_t}^2}. \end{cases}$$

Note also that using a reasoning similar to the one in the proof of Theorem 4 in Bianchi et al. (2021), based on the homogeneity property of the CRRA utility function, one can show the uniqueness of this solution (see also Dana, 1995).

To complete the identification of λ_{δ,t,W_t} , let determine K_{t,W_t} , which is obtained via the construction of the intertemporal representative agent.

By construction, this agent is the fictitious investor who, if endowed with the total wealth of the economy, would have a marginal utility equal to the equilibrium price. Formally, because her utility function is given by $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$, it translates into $\tilde{p}_{t+T} = \tilde{y}_{t+T}^{-\gamma} M_{AR,t+T}$, which directly leads to

$$M_{AR,t+T} = \left(\int (\lambda_{\delta,t,W_t} M_{\delta,t+T})^{\frac{1}{\gamma}} d\delta \right)^{\gamma}.$$

Moreover, for the intertemporal representative agent beliefs to be well-defined, one needs to ensure that $\mathbb{E}_t(M_{AR,t+T}) = 1$. Easy computations then result in

$$K_{t,W_t} = \frac{1}{\sqrt{2\pi}^{\gamma}} \frac{\sqrt{\sigma^2 + T a_{t,W_t}^2}^{\gamma-1}}{\sqrt{a_{t,W_t}^2}^{\gamma} \sqrt{\sigma^2}^{\gamma-1}} \exp\left(-\frac{a_{t,W_t}^2 b_{t,W_t}^2}{2}\right).$$

Defining the function φ as given in Theorem 2.1 and plugging the expressions of a_{t,W_t} , b_{t,W_t} , and K_{t,W_t} into (2.12), before rearranging the terms, yield the results.

Lastly, I need to ensure that the time- t dynamic budget constraint is verified. To do so, let define $\mu_{\tilde{p}_{t+T}}$ and $\sigma_{\tilde{p}_{t+T}}$ (that depend on $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$) such that $d\tilde{p}_{t+T} = \mu_{\tilde{p}_{t+T}} \tilde{p}_{t+T} dt + \sigma_{\tilde{p}_{t+T}} \tilde{p}_{t+T} dW_t$. I further define $\mu_{\tilde{y}_{\delta,t+T}}$ and $\sigma_{\tilde{y}_{\delta,t+T}}$ (that also depend on the parameters of the Generation- t wealth share distribution) similarly.

Using Ito's lemma, this budget constraint leads to

$$\mathbb{E}_t \left(\mu_{\tilde{y}_{\delta,t+T}} + \sigma_{\tilde{p}_{t+T}} \sigma_{\tilde{y}_{\delta,t+T}} - \tilde{\mu} - \sigma \sigma_{\tilde{p}_{t+T}} \right) = 0. \quad (2.14)$$

Direct computations allow me to rewrite the left hand-side of (2.14) as a polynomial function of δ of degree two. As (2.14) must be verified for all agents, each coefficient of the polynomial form must equal zero. By identification, this leads to the expressions of $\bar{\delta}_{t,W_t}$ and $\bar{\omega}_{t,W_t}$ derived in the theorem. ■

Proof of Proposition 2.1 The proof of Proposition 2.1 directly follows from the equivalence of the two models given in Lemma 2.1 and from the proof of the first and third items in Theorem 2.1. One simply needs to replace \tilde{y}_{t+T} by y_{t+T} , $\tilde{y}_{\delta,t+T}$ by $y_{\delta,t+T}$, \tilde{p}_{t+T} by p_{t+T} , and $\tilde{\nu}_{\delta,t,W_t}$ by ν_{δ,t,W_t} . Moreover, at time t , the agents of the sliding horizon model correspond to the Generation- t of the overlapping generations model.

Finally, note that the proportions of consumed and bequeathed wealth are explicitly given by (2.7) in the proof of Lemma 2.1. ■

Proof of Proposition 2.2 The proof of Proposition 2.2 directly follows from the proof of the second item in Theorem 2.1. ■

Proof of Proposition 2.3

1. By no arbitrage, the time- t stock price is given by

$$S_t = \frac{\mathbb{E}_t(p_{t+T}y_{t+T})}{\mathbb{E}_t(p_{t+T})}.$$

Computing the numerator and the denominator, and rearranging the terms lead to the formula in Proposition 2.3.

To determine the benchmark economy stock price \bar{S}_t , I set $\bar{\delta}$ and $\bar{\omega}$ to zero, and substitute them into the stock price formula.

2. 3. Applying Ito's lemma to the time- t stock price formula yields the results.

Similarly to the first item, $\bar{\mu}$ and $\bar{\sigma}$ are defined by setting $\bar{\delta}$ and $\bar{\omega}$ to zero into the formulas.

■

2.B Additional empirical results

2.B.1 Predicting market excess returns

In this appendix, I run a similar analysis as the one presented in Section 2.3.2. The only difference is that the dependent variable (either value- or equally-weighted) is now given by the excess returns of the market index instead of its raw returns. The results are reported in Table 2.7. They further confirm the model-implied positive relation: the coefficients are all positive and statistically significant. These results thus provide additional support for the validation of $H1-H2$.

Insert Table 2.7 here.

2.B.2 Rolling window analysis of the raw simple market index returns at the quarterly horizon

In order to gain deeper insights of the predicted positive relation between market belief dispersion and market index raw returns, I run rolling window regressions of Equation (2.3) using quarterly data and subsamples of 15 years, i.e., with 60 quarterly observations. This is the same methodology as the one in Section 2.3.3, and the results are reported in Figure 2.3.

Insert Figure 2.3 here.

The framework leads to a total of 92 regressions for each specification, and I therefore obtain time series of the estimated values of θ .¹⁷ The horizontal axis shows the end date of the subsamples, while the vertical axis gives the estimated value. Thicker rounds (resp. crosses) indicate statistically significant positive (resp. negative) values at the 10% level. The graphs show that most of the subsamples ending between 2011 and 2019 yield positive estimated values, while subsamples ending during the previous decade lead to negative ones.

¹⁷The first subsample starts in the second quarter of 1982 and ends in the first quarter of 1997.


Additionally, Table 2.8 reports the number of positive estimated coefficients for all specifications, along with the number of significantly positive (second column) and negative (third column) ones. Panel A presents the results for rolling windows of 60 quarterly observations, and Panels B and C use different window lengths of 40 and 80 quarterly observations respectively.

Insert Table 2.8 here.


Overall, apart from those obtained with the variable $\theta_{std,3M}^{VW}$, the results mostly confirm that the model implication seems to be verified empirically.

Figure 2.1: Time series of the quarterly market belief dispersion variables

Panel A: cross-sectional average of individual stock disagreements




Panel B: cross-sectional standard deviation of individual stock disagreements



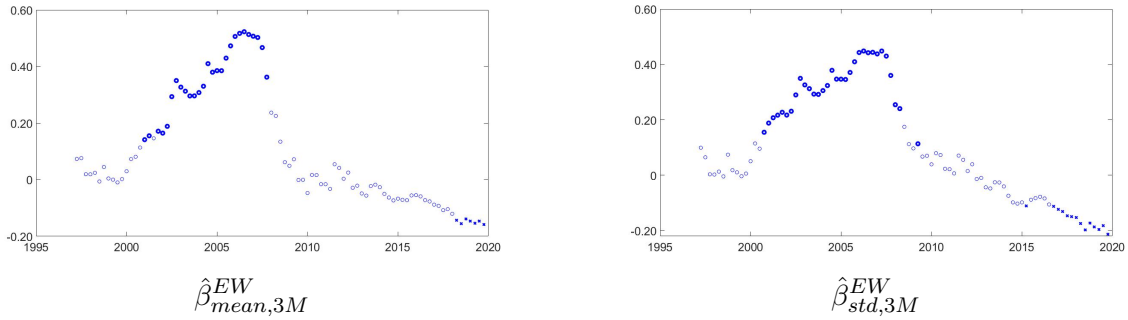
The figure presents the time evolution of the quarterly market belief dispersion variables, expressed in percentages. Panel A (resp. Panel B) shows the variables computed as the cross-sectional average (resp. standard deviation) of individual stock disagreements. The left column shows the value-weighted variables and the right column the equally-weighted ones. The sample goes from the first quarter of 1982 to the fourth quarter of 2019.

Figure 2.2: Time series of rolling window estimates of β

Panel A: Value-weighted specifications




Panel B: Equally-weighted specifications



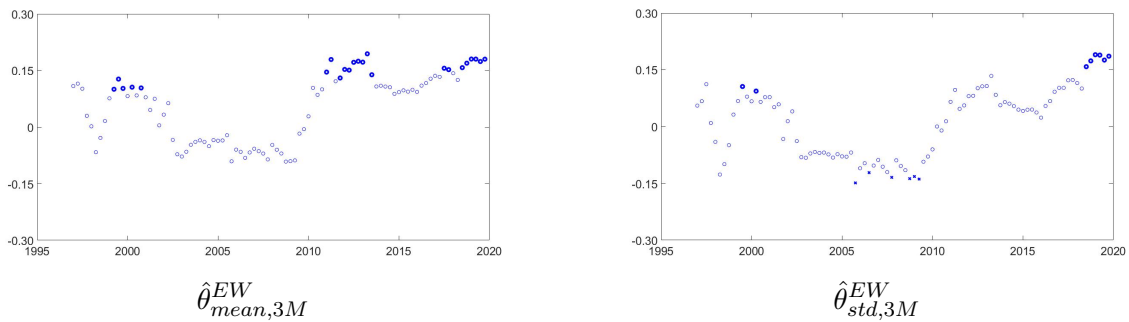
The figure presents the estimated values of $\beta_{mean,3M}^{VW}$, $\beta_{std,3M}^{VW}$, $\beta_{mean,3M}^{EW}$, and $\beta_{std,3M}^{EW}$ obtained from rolling window regressions with 60 quarterly observations. The horizontal axis shows the end date of the subsamples (the first one ends in the second quarter of 1997). Thicker rounds (resp. crosses) indicate statistically significant positive (resp. negative) values at the 10% level (using one-sided tests).

Figure 2.3: Time series of rolling window estimates of θ

Panel A: Value-weighted specifications



Panel B: Equally-weighted specifications



The figure presents the estimated values of $\theta_{mean,3M}^{VW}$, $\theta_{std,3M}^{VW}$, $\theta_{mean,3M}^{EW}$, and $\theta_{std,3M}^{EW}$ obtained from rolling window regressions with 60 quarterly observations. The horizontal axis shows the end date of the subsamples (the first one ends in the first quarter of 1997). Thicker rounds (resp. crosses) indicate statistically significant positive (resp. negative) values at the 10% level (using one-sided tests).

Table 2.1: Summary statistics of the quarterly market belief dispersion variables

	Mean	St. Dev.	Skewness	Kurtosis	ρ_1	ρ_2
Panel A: value-weighted specification						
$\bar{\omega}_{mean,3M}^{VW}$	3.26	0.43	0.84	3.28	0.89	0.77
$\bar{\omega}_{std,3M}^{VW}$	2.80	0.74	0.39	2.01	0.94	0.85
Panel B: equally-weighted specification						
$\bar{\omega}_{mean,3M}^{EW}$	3.78	0.37	0.53	2.63	0.88	0.74
$\bar{\omega}_{std,3M}^{EW}$	3.48	0.61	0.34	1.87	0.94	0.84

The table contains descriptive statistics of the quarterly belief dispersion variables, expressed in percentages. The mean, standard deviation, skewness, and kurtosis are reported for each variable, as well as the first- and second-order sample autocorrelations (ρ_1 and ρ_2). The sample goes from the first quarter of 1982 to the fourth quarter of 2019.

Table 2.2: Summary statistics of the market index returns for various holding periods

	Mean	St. Dev.	Skewness	Kurtosis
Panel A: value-weighted specification				
RET_{1M}^{VW}	0.29	3.79	-0.55	5.26
RET_{3M}^{VW}	0.86	6.23	-0.25	4.47
RET_{6M}^{VW}	1.74	9.36	0.23	5.29
RET_{12M}^{VW}	3.25	12.71	-0.20	2.39
RET_{24M}^{VW}	6.30	14.94	-0.48	2.80
Panel B: equally-weighted specification				
RET_{1M}^{EW}	0.35	4.07	-0.87	7.00
RET_{3M}^{EW}	1.04	7.01	-0.53	4.49
RET_{6M}^{EW}	2.03	9.85	-0.03	4.47
RET_{12M}^{EW}	3.92	13.66	0.03	2.95
RET_{24M}^{EW}	7.39	16.58	0.80	3.07

The table contains descriptive statistics of the raw market index simple returns for various holding periods, expressed in percentages. The mean, standard deviation, skewness, and kurtosis are reported for each variable. The subscript indicates the holding period considered and the superscript the type of weighting. The sample goes from January 1982 to December 2019.

Table 2.3: Market belief dispersion and market index raw returns

	1M	3M	6M	12M	24M
Panel A: value-weighted specification					
$\theta_{mean,..}^{VW}$	0.03 (0.81)	0.09* (1.39)	0.14* (1.60)	0.10 (0.75)	0.28** (2.03)
$\theta_{std,..}^{VW}$	0.05 (1.18)	0.10* (1.41)	0.14* (1.47)	0.13 (1.15)	0.31** (1.92)
Panel B: equally-weighted specification					
$\theta_{mean,..}^{EW}$	0.08** (1.78)	0.13** (1.98)	0.20*** (2.42)	0.25* (1.67)	0.42*** (2.58)
$\theta_{std,..}^{EW}$	0.07* (1.62)	0.10* (1.58)	0.17** (2.01)	0.25** (1.84)	0.48*** (3.01)

The table contains the results of regression (2.3). Inference is based on autocorrelation- and heteroskedasticity-robust standard errors (Newey and West, 1987), and all variables are standardized prior to estimation. For frequencies longer than a month, I use non-overlapping returns. I consider one-sided tests where the null is $H_0 : \theta = 0$ against the alternative $H_1 : \theta > 0$. The sample goes from January 1982 to December 2019.

Table 2.4: Summary statistics of the annualized market (log) volatility for various horizons

	Mean	St. Dev.	Skewness	Kurtosis
Panel A: value-weighted specification				
$LVOL_{1M}^{VW}$	-2.06	0.46	0.62	3.91
$LVOL_{3M}^{VW}$	-2.00	0.41	0.81	4.12
$LVOL_{6M}^{VW}$	-1.97	0.39	0.75	3.85
$LVOL_{12M}^{VW}$	-1.93	0.38	0.39	2.97
$LVOL_{24M}^{VW}$	-1.90	0.34	0.46	2.65
Panel B: equally-weighted specification				
$LVOL_{1M}^{EW}$	-2.09	0.49	0.70	3.95
$LVOL_{3M}^{EW}$	-2.01	0.43	0.92	4.29
$LVOL_{6M}^{EW}$	-1.98	0.41	0.94	4.28
$LVOL_{12M}^{EW}$	-1.94	0.39	0.73	3.31
$LVOL_{24M}^{EW}$	-1.91	0.37	0.80	3.56

The table contains descriptive statistics of the annualized market log volatility variables for various horizons. The mean, standard deviation, skewness, and kurtosis are reported for each variable. The subscript indicates the horizon considered and the superscript the type of weighting. The sample goes from January 1982 to December 2019.

Table 2.5: Market belief dispersion and annualized market (log) volatility

	<i>1M</i>	<i>3M</i>	<i>6M</i>	<i>12M</i>	<i>24M</i>
Panel A: value-weighted specification					
$\beta_{mean,..}^{VW}$	0.02	0.04	0.04	0.12	0.01
	(0.60)	(0.78)	(0.52)	(0.70)	(0.04)
$\beta_{std,..}^{VW}$	-0.00	0.01	0.03	0.07	-0.07
	(-0.09)	(0.26)	(0.43)	(0.37)	(-0.29)
Panel B: equally-weighted specification					
$\beta_{mean,..}^{EW}$	0.01	0.01	0.09	0.26*	0.18
	(0.18)	(0.15)	(0.71)	(1.56)	(0.86)
$\beta_{std,..}^{EW}$	0.01	0.04	0.10	0.26*	0.11
	(0.34)	(0.55)	(0.95)	(1.41)	(0.69)

The table contains the results of regression (2.4) and other similar regressions that use (log) volatility as the dependent variable and market belief dispersion as the predictor (controlling for past (log) volatility). Inference is based on autocorrelation- and heteroskedasticity-robust standard errors (Newey and West, 1987), and all variables are standardized prior to estimation. For frequencies longer than a month, I use non-overlapping data. I consider one-sided tests where the null is $H_0 : \beta = 0$ against the alternative $H_1 : \beta > 0$. The sample goes from January 1982 to December 2019.

Table 2.6: Statistics of rolling window estimates of β

	$\#\hat{\beta} > 0$	$\#\hat{\beta}^* > 0$	$\#\hat{\beta}^* < 0$
Panel A: 15 year rolling window (91 samples)			
$\hat{\beta}_{mean,3M}^{VW}$	83	24	0
$\hat{\beta}_{std,3M}^{VW}$	58	25	0
$\hat{\beta}_{mean,3M}^{EW}$	52	27	7
$\hat{\beta}_{std,3M}^{EW}$	59	32	14
Panel B: 10 year rolling window (111 samples)			
$\hat{\beta}_{mean,3M}^{VW}$	92	36	0
$\hat{\beta}_{std,3M}^{VW}$	73	22	10
$\hat{\beta}_{mean,3M}^{EW}$	60	29	10
$\hat{\beta}_{std,3M}^{EW}$	59	31	22
Panel C: 20 year rolling window (71 samples)			
$\hat{\beta}_{mean,3M}^{VW}$	53	24	0
$\hat{\beta}_{std,3M}^{VW}$	49	27	4
$\hat{\beta}_{mean,3M}^{EW}$	52	18	4
$\hat{\beta}_{std,3M}^{EW}$	50	40	10

The table contains some statistics on the estimated values of $\beta_{j,3M}^k$ ($j = \{mean, std\}$ and $k = \{VW, EW\}$) obtained from rolling window regressions for several settings. Panel A refers to the specification defined by Equation (2.4) and uses 15 year rolling windows (60 quarterly data). Panel B and C show results with 10 year rolling windows (40 quarterly data) and 20 year rolling windows (80 quarterly data) respectively. The first column reports the number of positive coefficients obtained from these regressions. The second (resp. third) column report the number of significantly positive (resp. negative) coefficients at the 10% level obtained from these regressions (using one-sided tests). The overall sample goes from the first quarter of 1982 to the fourth quarter of 2019.

Table 2.7: Market belief dispersion and market index excess returns

	<i>1M</i>	<i>3M</i>	<i>6M</i>	<i>12M</i>	<i>24M</i>
Panel A: value-weighted specification					
$\theta_{mean,..}^{VW}$	0.06*	0.14**	0.20***	0.19*	0.43***
	(1.45)	(2.16)	(2.45)	(1.41)	(2.67)
$\theta_{std,..}^{VW}$	0.07**	0.15**	0.20***	0.23**	0.46**
	(1.92)	(2.33)	(2.42)	(1.92)	(2.29)
Panel B: equally-weighted specification					
$\theta_{mean,..}^{EW}$	0.10**	0.17***	0.25***	0.32**	0.50***
	(2.27)	(2.53)	(2.99)	(2.02)	(2.62)
$\theta_{std,..}^{EW}$	0.09**	0.14**	0.23***	0.33**	0.56***
	(2.20)	(2.25)	(2.61)	(2.23)	(2.95)

The table contains the results of regression (2.3) when the dependent variable is given by the market index excess returns. Inference is based on autocorrelation- and heteroskedasticity-robust standard errors (Newey and West, 1987), and all variables are standardized prior to estimation. For frequencies longer than a month, I use non-overlapping returns. I consider one-sided tests where the null is $H_0 : \theta = 0$ against the alternative $H_1 : \theta > 0$. The sample goes from January 1982 to December 2019.

Table 2.8: Statistics of rolling window estimates of θ

	$\#\hat{\theta} > 0$	$\#\hat{\theta}^* > 0$	$\#\hat{\theta}^* < 0$
Panel A: 15 year rolling window (92 samples)			
$\hat{\theta}_{mean,3M}^{VW}$	44	25	27
$\hat{\theta}_{std,3M}^{VW}$	43	3	37
$\hat{\theta}_{mean,3M}^{EW}$	60	23	0
$\hat{\theta}_{std,3M}^{EW}$	55	8	6
Panel B: 10 year rolling window (112 samples)			
$\hat{\theta}_{mean,3M}^{VW}$	55	8	24
$\hat{\theta}_{std,3M}^{VW}$	61	16	36
$\hat{\theta}_{mean,3M}^{EW}$	75	34	2
$\hat{\theta}_{std,3M}^{EW}$	64	24	18
Panel C: 20 year rolling window (72 samples)			
$\hat{\theta}_{mean,3M}^{VW}$	40	13	12
$\hat{\theta}_{std,3M}^{VW}$	22	8	19
$\hat{\theta}_{mean,3M}^{EW}$	59	22	0
$\hat{\theta}_{std,3M}^{EW}$	40	2	2

The table contains some statistics on the estimated values of $\theta_{j,3M}^k$ ($j = \{mean, std\}$ and $k = \{VW, EW\}$) obtained from rolling window regressions defined similarly as Equation (2.3). Panel A uses 15 year rolling windows (60 quarterly data). Panel B and C show results with 10 year rolling windows (40 quarterly data) and 20 year rolling windows (80 quarterly data) respectively. The first column reports the number of positive coefficients obtained from these regressions. The second (resp. third) column report the number of significantly positive (resp. negative) coefficients at the 10% level obtained from these regressions (using one-sided tests). The overall sample goes from the first quarter of 1982 to the fourth quarter of 2019.

Chapter 3

Two skewed risks

Joint work with Paul Karehnke.

Abstract

We analyze the joint effects of skewness and correlation in a simple two-asset framework. Returns follow the split bivariate normal distribution, which combines bivariate normal distributions with different standard deviations and provides a good empirical fit. We show that equilibrium risk premia deviate from the CAPM if assets differ in skewness. Moreover, if the more positively skewed asset is more volatile, it underperforms and its beta, maximum return, idiosyncratic and systematic skewnesses are all higher—consistent with empirical evidence. We also derive formulas and analyze the role of skewness for portfolio choice and recently proposed conditional risk metrics.

Keywords: Skewness, Co-skewness, Conditional expected shortfall, Conditional value-at-risk, Portfolio choice, Asset pricing

JEL classification: G11 - G12 - G32

3.1 Introduction

Many settings relevant for investors and in risk management can be described in terms of two risks. Examples include the decision to add a risky asset to a risky portfolio, a default probability conditional on a market downturn, or a bank's contribution to the risk of the financial system. The role of the first two (co-)moments of asset returns in these settings is well understood. Less is known, however, about the impact of the next (co-)moment. At the same time, there is mounting evidence that the third moment is of key interest to investors and that most assets have skewed returns.¹

This paper aims at filling this gap and studies in a two-asset framework how skewness and its interaction with correlation affect portfolio choice, asset prices, and popular risk metrics. Our focus is on simplicity and intuition and our tool is the split bivariate normal (SBN) distribution. This distribution generalizes the bivariate normal (BN) distribution and introduces skewness by allowing the standard deviation parameters to differ above and below the mode. It is particularly appealing for finance applications because it takes into account that means and variances of stock returns are related empirically (Duffee, 1995). Moreover, it allows for many explicit computations due to its similarities with the BN distribution. We use the SBN distribution to demonstrate that low levels of asymmetry suffice to modify many standard predictions in an economically important way and reconcile them with empirical evidence.

In a pure exchange economy, we show that capital asset pricing model (CAPM) alphas are non-zero when the assets have different levels of skewness. Empirically, deviations from the CAPM have been related to systematic skewness (Harvey and Siddique, 2000), idiosyncratic

¹Recent evidence of skewness preference comes from, e.g., experiments (Ebert and Wiesen, 2011), capital budgeting decisions (Schneider and Spalt, 2016), and option-based decompositions of the market risk premium (Schneider, 2019). In addition, survey evidence from financial professionals finds that skewness systematically affects risk perceptions, while variance and kurtosis do not (Holzmeister et al., 2020). There is also evidence that, e.g., stock returns should theoretically be and empirically are skewed. For example, Albuquerque (2012) documents and reconciles positive skewness in individual stocks and negative skewness in aggregate stock returns and Del Viva et al. (2017) show that real options explain the skewness in individual stock returns.

skewness (Boyer et al., 2010), and maximum returns (Bali et al., 2011). Consistent with this empirical evidence, our contribution is to show theoretically that the asset that underperforms also has high systematic skewness, high idiosyncratic skewness, and high maximum returns. Acknowledging that more positively skewed stocks are also more volatile (see, e.g., Kumar, 2009), the framework can accommodate the high beta low return anomaly (Frazzini and Pedersen, 2014; Bali et al., 2017). The differences in maximum returns thereby directly arise from the underlying skewed return distribution, and the differences in idiosyncratic and systematic skewnesses result from the market equilibrium conditions. The representative investor’s skewness preference then leads to asset prices that deviate from the CAPM. The framework thereby shows that high max asset is both a lottery (Bali et al., 2011) and insurance (Barinov, 2018). More specifically, while the asset’s underperformance relates to high maximum returns and high idiosyncratic skewness, it is mainly driven by high systematic skewness. This is supported empirically by Barinov (2018) and Langlois (2020).

As far as portfolio choice is concerned, we show that the utility gains of considering skewness usually decrease with correlation when skewnesses are equal. When skewnesses differ, however, the relation becomes U-shaped and the gains are greater for highly correlated than for moderately correlated assets. In addition, underdiversification then becomes optimal, especially for high levels of correlation. Our portfolio choice analysis therefore highlights the role of correlation in the diversification-skewness trade-off and thus extends Mitton and Vorkink (2007) who consider only one skewed asset.

Recent research argues that expected utility (EU) preferences assign a low importance to skewness, while many popular behavioral theories like Tversky and Kahneman’s (1992) cumulative prospect theory (CPT) assign a higher—first-order importance—to skewness (see Ebert and Karehnke, 2020). Hence, we study portfolio choice not only for EU preferences but also for CPT. We show that these preferences magnify the importance of skewness for portfolio weights, certainty equivalents, and deviations from the CAPM.

Because the SBN distribution shares similarities with the BN distribution, it also allows

for a simple generalization of conditional risk metric formulas to skewed returns. In particular, Adrian and Brunnermeier (2016) propose a measure called ΔCoVaR , which is defined as the difference between the value-at-risk (VaR) of a given asset when a conditioning risk is under distress and this VaR when the conditioning risk is in its median state. In addition, Acharya et al. (2017) propose a risk measure—here referred to as the conditional expected shortfall (CoES)—that is defined as the expected value of a risky asset conditional on another risk being in the left tail of its distribution. We provide simple expressions for both ΔCoVaR and CoES when risks follow the SBN distribution. Using these expressions, we show that for positively correlated risks, both metrics are larger in absolute value when the main asset of interest is negatively skewed. Our contribution is to demonstrate that skewness thereby has a large quantitative impact. For, e.g., a correlation of 0.6, a 5% probability, a symmetric conditioning risk, and a skewness of -0.6 rather than zero, ΔCoVaR is about 30% higher in absolute value. Our analysis also shows that the skewness of the conditioning risk only has a limited impact on both risk measures. This corroborates Adrian and Brunnermeier’s (2016) empirical finding that ΔCoVaR is weakly correlated with the VaR of the conditioning risk.

Throughout the paper, our analysis of the joint effects of skewness and correlation in a two-asset framework relies on the SBN. The distribution is a bivariate extension of Fechner’s (1897) split normal distribution, which is a special case of the more widely used skewed t distribution of Hansen (1994). Arguing that symmetry is the exception rather than the rule, Fechner introduced the distribution to describe empirical regularities observed in many different fields. The univariate distribution has since then been known under several names in the literature (see Wallis, 2014, for a detailed survey) and has been used in finance to model inflation risk (Britton et al., 1998), downside risk (Feunou et al., 2013), and in one-risky-asset portfolio choice, asset pricing, and risk management applications (De Roon and Karehnke, 2017). The bivariate version that we use is very similar to the more general developments in Geweke (1989), Bauwens and Laurent (2005), and Villani and Larsson (2006), and it coincides with these distributions when the two risks are independent.

Figure 3.1 illustrates the construction of the SBN distribution. The left-hand side figure shows in each quadrant the probability density function (pdf) of BN distributions with different standard deviations. These pdfs are scaled in the right-hand side figure to obtain a continuous skewed distribution. The main advantage of the distribution is its simplicity that it inherits from the BN distribution and its ability to be skewed. Compared to more sophisticated distributions (see, e.g., Babić et al., 2019, for a recent overview), it allows us to focus on relatively small deviations from the standard, symmetric framework.

Insert Figure 3.1 here.

While the central contribution of the paper is to derive economic insights on the effects of skewness within a simple and intuitive bivariate framework, it is important to show that the distribution is empirically relevant and to check whether it provides a reasonable description of asset returns. We thereby analyze index excess returns of different asset classes (large stocks, small stocks, bonds, commodities, and emerging markets stocks) and also consider related distributions: the BN distribution, Azzalini and Dalla-Valle’s (1996) bivariate skew-normal (BSN) distribution, and a bivariate version of Hansen’s (1994) skewed t (BST) distribution constructed with a Gaussian copula.² The parameter values are chosen such that the correlation and the first three moments (four moments for the BST) coincide with the empirical distribution. Bivariate Kolmogorov-Smirnov tests and a comparison of implied and empirical co-skewnesses show that the SBN distribution fits the data considerably better than the BN distribution. This fit is also good when compared to the other two skewed distributions. For the co-skewness comparison, we derive the standard error of the coefficient of co-skewness for normally distributed returns, which to our knowledge has usually been computed with simulations (as in, e.g., Harvey and Siddique, 2000). We further compare the theoretical and empirical CoES levels and show that the SBN distribution has the lowest sum of squared differences which indicates that its fit is especially good in the left tail. This

²Note that Hansen’s (1994) skewed t distribution is a particular case of the skewed generalized t distribution defined in Theodossiou (1998) and can be reparametrized as the skewed Student- t defined in Zhu and Galbraith (2010).

is of primary importance when dealing with risk, and the SBN distribution is therefore a useful tool to conduct our analysis.

Our paper contributes to the growing literature on non-normally distributed financial returns (see Jondeau et al., 2007, for a review). More precisely, it adds to the stream of literature interested in understanding the role of skewness. For example, Kraus and Litzenberger (1976), Kane (1982), Mitton and Vorkink (2007), and Chabi-Yo et al. (2014) study the role of skewness in portfolio choice and asset pricing by assuming investors with skewness preference and skewed returns without further specifying the return generating process. Such an approach is easily motivated with Taylor-series approximations to expected utility. We show that these approximations correctly assess deviations from CAPM pricing but fare less well in terms of predicting the magnitude of risk premia. Other studies have used skewed distributions in which returns are linear in the skewness parameters and thus obtain mutual fund separation theorems and a skewness factor that drives expected returns (see, e.g., Simaan, 1993; Langlois, 2013; Dahlquist et al., 2016). Return asymmetry is thereby generated by a common skewed factor. Thus, for a given correlation structure, the range of possible asset skewnesses is more limited, which makes it hard to simultaneously reproduce the stylized facts about skewness discussed above.

The paper is organized as follows. Section 3.2 presents the SBN distribution and its empirical fit to index excess returns. Readers more interested in the applications and results may skip that section and directly read Sections 3.3 and 3.4. These sections study the effects of skewness in portfolio choice and asset pricing and in risk management, respectively. Section 3.5 concludes. Appendix 3.A gathers technical details on the distribution, Appendix 3.B contains the proofs, and Appendix 3.C contains additional results such as the derivation of the standard error of co-skewness.

3.2 A simple skewed bivariate distribution

To conduct our analysis and derive our main results, our tool is the SBN distribution. We define it in Section 3.2.1 and provide evidence for its empirical relevance in Section 3.2.2.

3.2.1 Definition of the SBN distribution

We first recall the case in which X and Y follow the BN distribution that has means μ_X and μ_Y , standard deviations σ_X and σ_Y , and a correlation ρ . The corresponding pdf is

$$f(x, y; \mu_X, \mu_Y, \sigma_X, \sigma_Y, \rho) = \frac{1}{2\pi\sigma_X\sigma_Y\sqrt{1-\rho^2}} \exp \left[-\frac{1}{2(1-\rho^2)} \left(\left(\frac{x-\mu_X}{\sigma_X} \right)^2 + \left(\frac{y-\mu_Y}{\sigma_Y} \right)^2 - 2\rho \left(\frac{x-\mu_X}{\sigma_X} \right) \left(\frac{y-\mu_Y}{\sigma_Y} \right) \right) \right].$$

The SBN distribution is constructed with four BN distributions that all have the same mean vector (m_X, m_Y) and correlation parameter c and differ in their standard deviation parameters. Let $s_{X,1}$ denote the standard deviation parameter for X when X is below or equal to m_X , and let $s_{X,2}$ denote the standard deviation parameter when X is above m_X . The corresponding parameters for Y , $s_{Y,1}$ and $s_{Y,2}$, are defined similarly. Thus, the standard deviation parameters in increasing order from the first to the last quadrant of $(X - m_X, Y - m_Y)$ are $(s_{X,2}, s_{Y,2})$, $(s_{X,1}, s_{Y,2})$, $(s_{X,1}, s_{Y,1})$, and $(s_{X,2}, s_{Y,1})$.

Recall that any pdf needs to integrate to one and be non-negative. In addition, we require the pdf to be continuous—as in the univariate case. Therefore, each BN pdf is scaled with a positive weight $\lambda_{i,j}$ for $i = 1, 2$ and $j = 1, 2$. Direct computations give

Definition. *The pdf of the SBN distribution for X and Y is*

$$g(x, y; \Omega) = \begin{cases} \lambda_{1,1}f(x, y; m_X, m_Y, s_{X,1}, s_{Y,1}, c), & \text{if } x \leq m_X \text{ and } y \leq m_Y, \\ \lambda_{1,2}f(x, y; m_X, m_Y, s_{X,1}, s_{Y,2}, c), & \text{if } x \leq m_X \text{ and } y > m_Y, \\ \lambda_{2,1}f(x, y; m_X, m_Y, s_{X,2}, s_{Y,1}, c), & \text{if } x > m_X \text{ and } y \leq m_Y, \\ \lambda_{2,2}f(x, y; m_X, m_Y, s_{X,2}, s_{Y,2}, c), & \text{if } x > m_X \text{ and } y > m_Y, \end{cases}$$

where $\Omega = \{m_X, m_Y, s_{X,1}, s_{X,2}, s_{Y,1}, s_{Y,2}, c\}$ is the set of parameters, $\lambda_{1,1}$, $\lambda_{1,2}$, $\lambda_{2,1}$, and $\lambda_{2,2}$ are reported in Appendix 3.A.1, and f is the pdf of the BN distribution.

We report the contour plot of the distribution in Panel A of Figure 3.2. The parameters are chosen such that X and Y have zero means, unit variances, and correlation of 0.4 and that the (standardized) skewness of X , γ_X , is zero, and the (standardized) skewness of Y , γ_Y , is -0.6 . For comparison, we report the contour plots of other bivariate distributions that also have these moments for X and Y . Panels B-D show the BN distribution, the BSN distribution, and the BST distribution constructed with a Gaussian copula. Comparing Panel A with Panel B illustrates that a moderate skewness of -0.6 modifies the appearance of the contour plot significantly: the lines are tighter when the negatively skewed variable Y takes positive values and wider otherwise. The distribution used in this paper also differs markedly from Azzalini and Dalla-Valle's (1996) BSN distribution. In particular, the contour lines change appearance more abruptly when moving from one quadrant to the other. This feature might be useful to fit return distributions for which means and standard deviations are related empirically (Duffee, 1995). Finally, the distribution is most similar to the BST. This is perhaps unsurprising because these two distributions are identical in the limit in the univariate case and the kurtosises of the BST are chosen to fit the kurtosises of the SBN. Nevertheless, there are some slight differences between the two that are most evident at the borders of the quadrants. Here, the contour lines change more abruptly for the SBN.

Insert Figure 3.2 here.

In Appendix 3.A.2, we derive the moment generating function of the SBN and report closed-form formulas for the moments and co-moments of the distribution as functions of the parameters. Several observations about these formulas are in order. First, if $s_{X,1} = s_{X,2}$ and $s_{Y,1} = s_{Y,2}$, the skewnesses are zero and it is easy to check that we obtain the moments of the BN distribution. When skewnesses are non-zero, the distribution has the desirable feature that its modes, which correspond to the location parameters m_X and m_Y , differ from the

means. Second, only the ratios of the standard deviations matter when studying the skewnesses and the co-skewnesses of the distribution, not the standard deviations themselves. These ratios, e.g., $s_{X,2}/s_{X,1}$, are also easy to interpret as they correspond to the probability that a split normally distributed variable takes values above its mode relative to the probability that it takes values below its mode. Thus, a split normal distribution arises from a return generating process with a volatility that takes values $s_{X,2}$ and $s_{X,1}$ above and below its mode and has on average $s_{X,2}/s_{X,1}$ -times as many observations above its mode than below. Third, if $c = 0$, we recover the expressions of the univariate split normal moments given in, e.g., De Roon and Karehnke (2017). Fourth, although this problem does not arise in our empirical section, some skewnesses of X and Y cannot be matched by the distribution. Indeed, the skewnesses of X and Y are bounded and the range of admissible skewnesses depends on the parameters. However, there is no simple expression for this range.³ Finally, because the distribution has seven parameters, we can choose values for μ_X , μ_Y , σ_X , σ_Y , $\sigma_{X,Y}$, γ_X , and γ_Y and solve for the matching parameters m_X , m_Y , $s_{X,1}$, $s_{X,2}$, $s_{Y,1}$, $s_{Y,2}$, and c . Because the number of parameters is limited, we cannot match co-skewness.⁴ Nevertheless, the co-skewnesses can help, e.g., to assess whether the distribution can fit additional moments of empirical data.

In the remainder of the paper, we use the expressions for the moments and co-moments in Appendix 3.A.2 to link symmetric and skewed distributions and study how skewness and its interaction with correlation affect financial decisions.⁵

³Unlike the univariate case in which the skewness of X only depends on the ratio of standard deviation parameters, $z_X \equiv s_{X,2}/s_{X,1}$, and setting $s_{X,1}$ ($s_{X,2}$) close to 0 gives the highest possible positive (negative) skewness, this skewness also depends on z_Y (defined analogously to z_X) and c in the bivariate case. Depending on the values of $c \in]-1, 1[$ and $z_Y \in]0, +\infty[$, the impact of z_X differs. For example, setting $s_{X,1}$ close to 0 does not lead to the largest possible skewness of X for certain values of z_Y and c .

⁴However, we could solve for the parameters that match the co-skewnesses instead of the skewnesses. We do not consider this alternative approach in this paper and leave it for future research.

⁵It might be noteworthy that any linear combination of two variables that follow the SBN distribution does not follow such a distribution, even if the two variables are independently distributed. Thus, the SBN is not a stable distribution. Multivariate stable distributions that allow for skewness, however, cannot have finite variances (Bauwens and Laurent, 2005). It is therefore impossible to construct a skewed stable distribution with a closed form pdf and finite moments.

3.2.2 Empirical fit

While the main objective of the paper is to study the impact of skewness on financial decisions with correlated risks, we first study the empirical fit of the SBN distribution to ensure that our theoretical analysis is based on a reasonable distribution. In this subsection, we study the overall fit with bivariate Kolmogorov-Smirnov tests and the fit in terms of co-skewness (which is not a free parameter in the distribution) by comparing empirical and implied ones. We further study the implied CoESs, i.e., the fit in the conditional left tail of the distribution, in Section 3.4.2. We conduct this analysis not only for the SBN, but also for the BN distribution, the BSN distribution, and the BST distribution constructed with a Gaussian copula. We emphasize that the advantage of the SBN lies in its simple and intuitive construction inherited from the BN distribution and its ability to be skewed.

We use monthly excess returns over the risk-free rate in the period from January 1989 to June 2018 on five indices from different asset classes: large stocks (S&P 500 TR Index), small stocks (Russell 2000 TR Index), bonds (Barclays US Aggregate Bond TR Index), commodities (GSCI TR), and emerging markets stocks (MSCI Emerging Market TR Index). The proxy for the risk-free rate is the return on the 30-day t-bill from CRSP and the index data is from Datastream.

Insert Table 3.1 here.

The descriptive statistics in Table 3.1 show that the index excess returns are all negatively skewed and exhibit excess kurtosis. Except for bonds and commodities, all skewnesses are statistically significant under the null hypothesis of normally distributed returns. While a split normal distribution can reproduce exactly the skewness in the data, it has also excess kurtosis but the level is usually too low to fit the kurtosis in the data (De Roon and Karehnke, 2017). Nevertheless, as we show below and in Section 3.4.2 respectively, the distribution can still fit the overall shape and the conditional left tail of the empirical distributions reasonably well. The correlations range from negative but close to zero (-0.05

for bonds and commodities) to large and positive (0.81 for large and small stocks). The index pairs that include bonds thereby consistently have the lowest correlations.

To illustrate the dependence structure of small stocks, bonds, commodities, and emerging markets stocks with large stocks, we further report in Figure 3.3 contour plots of empirical pdfs of the monthly excess returns on these index pairs calculated with a kernel density estimator (dashed lines) and compare them to the contour plots of the SBN distribution with parameters set to match the empirical moments (solid lines).

Insert Figure 3.3 here.

We formally assess the overall fit of the distributions with the bivariate Kolmogorov-Smirnov test proposed by Peacock (1983). Table 3.2 reports the p -values of this test under the null hypothesis that realized excess returns are drawn from the SBN, BN, BSN, or BST.⁶ At the 5% significance level, we reject the null for SBN in three out of ten cases, whereas we reject the null in seven (resp. four) cases for BN (resp. BSN). Thus, in addition to its analytical convenience, the SBN distribution seems to provide a reasonable fit compared to its close competitors. Compared to its close cousin the BST distribution that has two additional parameters to match the empirical kurtosises, the fit is less good. Here, the null is rejected at the 5% significance level for none of the ten pairs.

Insert Table 3.2 here.

The four distributions we study can only match the correlation of index pairs. Thus, by comparing the implied and empirical co-skewnesses, we obtain another assessment of the empirical fit. More formally, we test whether the empirical co-skewnesses are significantly different from the co-skewnesses obtained under the null that the index excess returns are drawn from a given skewed distribution. For the test with the BN distribution, we can derive

⁶We use the Matlab function `msn_moment_fit` (available at <http://azzalini.stat.unipd.it/SN/sn-download.html>) to find the parameters of the BSN distribution that match the empirical moments of the index pairs. This function indicates that the empirical moments of two index pairs yield parameters outside of the admissible range. Thus, we exclude these index pairs for the BSN distribution.

the sample variance of co-skewness under the null. It is

$$\text{var}(\gamma_{Y,X}) = \text{var}(\gamma_{X,Y}) = \frac{2 + 4\rho^2}{n}, \quad (3.1)$$

where ρ is the correlation between the two indices and n is the number of observations. The expression avoids having to approximate the sample variance of co-skewness with random draws from the BN distribution as in, e.g., Harvey and Siddique (2000). We are not aware of a reference for this formula and we therefore report its derivation in Appendix 3.C.2. From the formula in (3.1), the widely used sample variance of skewness of $6/n$ is obtained for $\rho = 1$.

Insert Table 3.3 here.

Table 3.3 reports the empirical co-skewnesses and the p -values computed under the different null hypotheses. All twenty co-skewnesses are negative. The two largest co-skewnesses in absolute value are between large and small stocks with a magnitude of about -0.6 . These two co-skewnesses are highly significant under the null of a normal distribution, which implies a zero co-skewness, while they are insignificant under the null of each of the skewed distributions. Thus, although the distributions do not have parameters to match the third co-moments, they can generate sizable co-skewnesses. This happens because, as shown in Table 3.1, small and large stocks are highly correlated and have large negative skewness. More generally, at the 5% significance level, 14 out of 20 co-skewnesses are significant under the normal null. For the split normal, this number is reduced to 11 and the p -values in the SBN column are generally higher than those in the BN column. For co-skewness, the other two skewed distributions seem to provide a better fit and the null is rejected in three cases for BSN (out of 16 co-skewnesses instead of 20) and seven cases for BST.

Overall, the SBN has a good empirical fit and is more parsimonious than the BST. Moreover, we show in Section 3.4.2 that the SBN has the best fit among the four distribution studied in the conditional left tail. The main advantage over the BST and BSN, however,

is the simplicity of the SBN that we exploit next to derive the joint impact of skewness and correlation on portfolio choice, asset pricing, and in risk management.

3.3 Portfolio choice and asset pricing

We study the effects of skewness on portfolio choice in Section 3.3.1 and asset pricing in Section 3.3.2. For each setting, we consider EU preferences and Tversky and Kahneman's (1992) CPT preferences.

3.3.1 Optimal choice

The investor can invest in two risky assets that have excess returns of R_X and R_Y and a risk-free asset that has a gross return of $1 + r_f$. The two risky assets can represent, e.g., returns on two stocks or the return on a portfolio of many stocks and the return on one additional stock. Let ω_X (resp. ω_Y) denote the investment in the risky asset X (resp. Y). Thus, the overall portfolio return is $1 + r_f + \omega_X R_X + \omega_Y R_Y$. We normalize initial wealth to one and assume that the investor derives utility from terminal wealth. We first consider an EU investor with exponential utility that maximizes $u(x) = -\exp(-\theta x)$, where $\theta > 0$ is the coefficient of absolute risk aversion. This utility function implies a preference for skewness and allows for simple expressions for portfolio weights and risk premia when excess returns follow the SBN distribution.⁷

Let us assume that R_X and R_Y follow the SBN distribution. Direct computations that exploit the properties of the SBN distribution and exponential utility functions yield the following proposition.

Proposition 3.1 (Portfolio weights with the SBN distribution). *Assume that R_X and R_Y follow the SBN distribution and that the investor maximizes EU with exponential utility. The*

⁷For our analysis it is not important that the exponential utility function lacks wealth effects because (a) we do not study the comparative statics of wealth and (b) De Roon and Karehnke (2017) show that the effects of skewness in asset pricing and portfolio choice are magnified with constant relative risk aversion utility functions, which allow for wealth effects by construction.

optimal portfolio weights ω_X and ω_Y are then implicitly given by

$$\omega_X = \frac{\tilde{\mu}_X \tilde{\sigma}_Y^2 - \tilde{\mu}_Y \tilde{\sigma}_{X,Y}}{\theta \tilde{\sigma}_X^2 \tilde{\sigma}_Y^2 \left(1 - \frac{\tilde{\sigma}_{X,Y}^2}{\tilde{\sigma}_X^2 \tilde{\sigma}_Y^2}\right)} \quad \text{and} \quad \omega_Y = \frac{\tilde{\mu}_Y \tilde{\sigma}_X^2 - \tilde{\mu}_X \tilde{\sigma}_{X,Y}}{\theta \tilde{\sigma}_X^2 \tilde{\sigma}_Y^2 \left(1 - \frac{\tilde{\sigma}_{X,Y}^2}{\tilde{\sigma}_X^2 \tilde{\sigma}_Y^2}\right)}, \quad (3.2)$$

where $\tilde{\mu}_X$, $\tilde{\mu}_Y$, $\tilde{\sigma}_X^2$, $\tilde{\sigma}_Y^2$, and $\tilde{\sigma}_{X,Y}$ are modified first and second co-moments that depend on the portfolio weights and are reported in the proposition's proof in Appendix 3.B.

The portfolio weights for the SBN distribution have a form similar to those obtained with a BN distribution. In particular, when the assets are not skewed, $\tilde{\mu}_X$, $\tilde{\mu}_Y$, $\tilde{\sigma}_X^2$, $\tilde{\sigma}_Y^2$, and $\tilde{\sigma}_{X,Y}$ simplify to the means, variances, and covariance of R_X and R_Y . Thus, we recover the normal case. Moreover, when correlation is zero (implying $\tilde{\sigma}_{X,Y} = 0$), the two portfolio weights are independent from each other and simplify to the one-asset portfolio weights reported in De Roon and Karehnke (2017). Finally, the ratio of the portfolio weights does not depend on θ .⁸ In other words, two-fund separation holds, which is expected with a risk-free asset and exponential utility regardless of the distributional assumptions (see, e.g., Ingersoll, 1987, Chapter 6.4).

With (3.2) in hand, we conduct a detailed analysis of the effects of correlation and skewness on different portfolio choice and asset pricing metrics. Table 3.4 analyzes portfolio choice with two assets whose excess returns both have a mean of 12%, a standard deviation of 15%, and a correlation that varies through columns from -0.2 to 0.8 .⁹ The skewnesses are equal in the top panels of the table (benchmark case and Panels A-B) and differ in the bottom panels (Panels C-E). The table reports optimal weights, certainty equivalent returns of the optimal portfolio $CER(\omega_{SBN})$ and of a portfolio with normal weights $CER(\omega_{BN})$, and skewnesses of the optimal portfolio $\gamma_{\omega_{SBN}}$ and of a portfolio with normal weights $\gamma_{\omega_{BN}}$.

Insert Table 3.4 here.

⁸In fact, although the modified moments are functions of weights, it suffices to solve the implicit functions for $\theta = 1$ and scale those weight by the inverse of θ .

⁹We focus on moderately negative to positive correlations because this case is more relevant empirically as shown in Section 3.2.2. The results for very negative correlations are available upon request.

Let us first consider the top panels of the table in which the skewnesses of X and Y are equal.¹⁰ As a consequence, the same amount is invested in each asset, and we obtain the familiar mean-variance result that weights decrease with correlation. Correlation also impacts portfolio skewness. For example, in Panel A, where both assets have a skewness of 0.6, the portfolio skewness increases from 0.4 (for $\rho = -0.2$) to 0.6 (for $\rho = 0.8$).

Comparing the certainty equivalent returns for different levels of skewness, we observe a trade-off between holding highly correlated assets with high levels of skewness and holding less correlated assets with lower levels of skewness. For instance, the certainty equivalent return is 9.6% for $\rho = 0.6$ in Panel A, while it is 9.4% for $\rho = 0.2$ in Panel B. The intuition behind this result lies in the definition of skewness itself: large positive events are more likely to occur for positive skewnesses. Hence, even if the correlation between the assets is high, the higher likelihood of very high returns makes the investor better off than holding less correlated assets that are more likely to experience crashes.

As discussed in De Roon and Karehnke (2017), the total effect of skewness on certainty equivalent returns can be decomposed in a change due to the distribution and a change due to the weights. The first effect is a utility gain (resp. loss) due to more positive (resp. negative) skewness in asset returns. For instance, for $\rho = 0.4$, the certainty equivalent return changes from 9.1% in the benchmark case to 10.4% in Panel A. The second effect is a utility gain due to the portfolio adjustment the investor makes to hold the optimal weights. This utility gain tends to decrease with correlation and ranges from 1.3% to 0.4% in Panels A and B.

In Panels C-E the skewnesses of X and Y differ, and it therefore becomes optimal to underdiversify. To measure the degree of underdiversification, these panels contain the ratio of the portfolio weights. As shown in Appendix 3.C.1, a ratio of one maximizes the Sharpe ratio in our setting, and deviations from one measure deviations from mean-variance efficiency. To facilitate the interpretation, the skewnesses are chosen such that $\gamma_X > \gamma_Y$ and

¹⁰In the remainder of Section 3.3, we refer to the skewness of asset X 's excess return simply as the skewness of asset X and similarly for asset Y .

thus $\omega_{X,SBN}/\omega_{Y,SBN} > 1$. The table shows that this ratio increases with correlation and skewness difference reaching 10.5 for $\rho = 0.8$ in Panel C in which $\gamma_X = 0.6$ and $\gamma_Y = -0.6$. This sacrifice in terms of mean-variance efficiency enables the investor to obtain a portfolio skewness of 0.5 versus 0 for the portfolio computed with the normal distribution. Moreover, weights invested in the more positively skewed asset can now increase with correlation. This happens when the skewness difference and the correlation are sufficiently large. For instance, in Panel C, the weight in X increases from 1.08 for $\rho = 0.6$ to 1.27 for $\rho = 0.8$. Finally, the relation between the utility gains of considering skewness and correlation is now U-shaped: the gains are greater for large positive correlations than for moderate ones.

Comparing Panels C-E shows that a portfolio weight is not only sensitive to the asset's own skewness, but also to the other asset's skewness. For instance, when $\gamma_X = 0.6$ and $\rho = 0.4$, the weight in X decreases from 1.11 for $\gamma_Y = -0.6$ (Panel C) to 1.00 for $\gamma_Y = 0.6$ (Panel A). More generally, an increase in the other asset's skewness decreases the amount invested when the two assets are positively correlated. In fact, when the correlation between X and Y is positive, the assets are (imperfect) substitutes. Hence, when the skewness of Y improves and the investor optimally demands more of Y , he also decreases his demand of X . In the knife-edge case of no correlation and because the utility function is of the exponential type, demands for X and Y are independent of each other.

The comparative statics extend the analysis of the utility gains of considering skewness of De Roon and Karehnke (2017) to two assets and confirm previous results of portfolio choice and skewness. In particular, Mitton and Vorkink (2007) show theoretically that skewness preference induces investors to hold underdiversified portfolios. They further show empirically that retail investors hold underdiversified portfolios and that the loss of mean-variance efficiency of their portfolios is largely explained by a gain in skewness. We obtain similar results within the standard EU framework. In addition, by looking at two skewed assets, our analysis highlights the role of changes in skewness as a motive for trading. Recent research documents that subjective beliefs of important economic variables are skewed and that this

subjective skewness is time-varying (Ghosh and Roussellet, 2020). Therefore, changes in skewness combined with heterogeneous skewness preferences are a plausible explanation for trading and can contribute to reduce the puzzle of why investors trade so much.¹¹

Next, we consider an investor with Tversky and Kahneman’s (1992) CPT preferences. The CPT investor uses decision weights that transform cumulative objective probabilities with a probability weighting function and evaluates outcomes with an S-shaped value function. As shown in, e.g., Ebert and Strack (2015), the probability weighting function generates a strong preference for skewness. However, it complicates the calculations substantially and we therefore solve the CPT portfolio choice maximization problem numerically. The caption of Table 3.5 provides the details.

Insert Table 3.5 here.

Table 3.5 reports the results of the portfolio choice analysis with CPT and has the same structure as Table 3.4. A few results stand out. Portfolio weights in CPT are much more sensitive to correlation. This is due to loss aversion, and the table shows that, e.g., in the symmetric case it is optimal not to invest in the risky assets when the correlation is high enough. Skewness also has a stronger impact on portfolio weights in CPT. First, it is optimal not to invest in the risky assets for, e.g., $\rho = 0.6$ when the return distribution is symmetric, but when skewnesses are 0.6 in Panel A, the weight in each asset is 3.08. Second, the gain of taking skewness into account is now larger. Third, when assets differ in their skewness and when correlation is sufficiently high, it can now be optimal to short the asset with the lower skewness to obtain a higher portfolio skewness. This is illustrated in Panel C for $\rho = 0.8$, where the overall portfolio skewness is 0.75.

In sum, the comparative statics show that CPT magnifies the effects of changes in correlation and skewness compared to EU. Ebert and Karehnke (2020) show that CPT induces a first-order preference for skewness, while this preference is only third-order in EU. Tables 3.4

¹¹This literature has focused on differences in information and overconfidence as motives for trading (Odean, 1999; Kelley and Tetlock, 2013).

and 3.5 illustrate that this skewness preference order governs the magnitude of the utility gains of considering skewness. For example, in Panel A when both assets have a skewness of 0.6 and $\rho = 0.4$, this certainty equivalent return gain exceeds 4% for CPT in Table 3.5, while it is only about 0.5% in the corresponding EU case.

3.3.2 Equilibrium returns

We compute equilibrium returns in an exchange economy with a representative investor using the expressions for portfolio weights in (3.2). In particular, we solve for the expected excess returns such that the investor optimally holds the exogenous supply of each asset denoted by the scalars $\bar{\omega}_X$ and $\bar{\omega}_Y$, respectively. The following proposition characterizes the resulting equilibrium with respect to a market portfolio that invests $\bar{\omega}_X$ in asset X and $\bar{\omega}_Y$ in asset Y .

Proposition 3.2 (Equilibrium returns with the SBN distribution). *Assume that R_X and R_Y follow the SBN distribution, that the representative investor maximizes EU with exponential utility, and that markets clear. The expected equilibrium excess returns over the risk-free rate are*

$$\mu_X = \bar{\beta}_X \bar{\mu}_M \xi + \zeta_X \quad \text{and} \quad \mu_Y = \bar{\beta}_Y \bar{\mu}_M \xi + \zeta_Y, \quad (3.3)$$

where $\bar{\mu}_M = \theta \bar{\sigma}_M^2$ is the modified market risk premium, $\bar{\beta}_X$ and $\bar{\beta}_Y$ are the modified betas, ζ_X , ζ_Y , and ξ are scalars, and $\bar{\sigma}_M^2$ is the modified market variance. The expressions for all these variables are given in the proposition's proof in Appendix 3.B.

The equilibrium excess returns have a similar form than in the normal case. However, additional scalars, ξ and ζ_k (for $k = X, Y$), appear to obtain expressions for the risk premium μ_k instead of the modified risk premium $\bar{\mu}_k$ given by Equation (3.26) in Appendix 3.B. These scalars vanish, i.e., $\xi = 1$ and $\zeta_k = 0$, when both assets have symmetric returns. In that case, we also have $\bar{\mu}_M = \mu_M$ and $\bar{\beta}_k = \beta_k$ and thus recover the normal case.

We use the proposition to analyze the asset pricing implications of skewness, with a focus on deviations from the standard mean-variance framework and CAPM. We study the case in which asset X has a skewness of 0.6 and asset Y has a skewness of -0.6 . Indeed, if the two assets have the same skewnesses, the market portfolio is still mean-variance efficient (although risk premia have different magnitudes than predicted by a mean-variance framework). We consider an economy with a representative EU investor as in Proposition 3.2 and an economy with a representative CPT investor. In the latter case, expected excess returns are solved numerically using a similar procedure as in Table 3.5.

Insert Table 3.6 here.

In Panel A, $\bar{\omega}_X = \bar{\omega}_Y = 0.5$ and a volatility of 15%. These assumptions imply that the assets have a beta of one.¹² The table shows that the more positively skewed asset X has a negative CAPM alpha and the alpha is lower with CPT than with EU ($\alpha_X^{CPT} = -2.3\%$ vs. $\alpha_X^{EU} = -0.5\%$ for $\rho = 0.4$). Asset Y instead has a positive alpha that equals the alpha of X in absolute value due to equilibrium accounting. The bottom part of the table reports the expected market excess return. Notice that the expected market excess returns μ_M are approximately the same in the EU and CPT economies for $\rho = 0.4$. This happens because market skewness is zero in Panel A and because we choose the coefficient of absolute risk aversion for the EU investor in each panel such that the equilibrium market excess return with EU is the same than with CPT when returns are normally distributed and $\rho = 0.4$.¹³

The last two rows of the table assess the accuracy of the widely used Taylor approximation approach to study skewness effects. It reports the alpha of asset X and the market risk premium that are obtained in an equilibrium in which the marginal utility of the representative investor is approximated with a second-order Taylor expansion around the expected

¹²An asset's beta is defined as usual as the covariance of an asset's return with the market return divided by market variance, where the market portfolio is a supply-weighted portfolio of asset X and Y . Similarly, an asset's CAPM alpha is the difference between an asset's risk premium and the product of its beta and the market risk premium.

¹³Panel A highlights an important difference between our approach and those in which asset return skewness is generated from a common skewed factor as in, e.g., Simaan (1993) and Dahlquist et al. (2016). It is clear that the latter approach is incapable of generating asset skewness when market skewness is zero and $\bar{\omega}_X = \bar{\omega}_Y = 0.5$.

market return. These quantities can be compared to those obtained using the formulas in Proposition 3.2 reported in the rows α_X^{EU} and μ_M^{EU} . The table shows that the Taylor approximation correctly assesses the sign of alpha. However, it underestimates the magnitude of risk premia and deviations from mean-variance efficiency. This underestimation is economically significant for the market risk premium (up to 1% in absolute value) and increases with correlation.

Empirically, more positively skewed stocks are also more volatile (Kumar, 2009). Thus, we set the volatility of assets X and Y to 20% and 10% in Panel B. A byproduct of different volatilities is that asset X now has a beta greater than one, while the beta of asset Y is below one. The table shows that asset X still has a negative alpha. This is interesting in the context of the well documented low performance of high beta stocks. Indeed, asset X is not attractive because of its high beta, but because of its high skewness. Schneider et al. (2020) obtain the same conclusion with simulated economies populated with Merton-type firms in which skewness arises endogenously due to leverage and stochastic volatility, and they provide consistent empirical evidence. In addition, Bali et al. (2017) demonstrate that investor's demand for lottery-like stocks is an important driver of the beta anomaly and that the beta anomaly disappears once regressions control for lottery demand. Bali et al. (2017) measure lottery demand with the maximum return—a proxy for the right tail of a distribution. To relate our analysis to their evidence, the table reports the return expectation in the 1% tails. As expected, the more positively skewed and more volatile asset X has a larger right tail (63% vs. 21%). Panel B also shows that asset X has a more extreme left tail (41% vs. 32%). This happens despite the higher skewness because asset X is more volatile. Thus, skewness preference provides a rationale for the well documented underperformance of stocks with high maximum returns and the mixed evidence for the underperformance of stocks with low minimum returns (Bali et al., 2011).

The results in the table can also be related to the recent study of Barinov (2018) on whether stocks with high maximum returns are lotteries or insurance. Barinov demon-

strates that stocks with high maximum returns covary positively with changes in market volatility—i.e., offer high returns when market volatility increases. In our static framework, co-skewness—the covariance between an asset’s return and the squared market excess return—measures this ability to provide insurance. The table shows that asset X indeed has positive co-skewness—i.e., provides insurance—while asset Y has negative co-skewness. Thus, the max asset provides both a lottery and insurance.¹⁴

While Panel B is able to reproduce many features of the data, the market excess return is positively skewed, but it is negatively skewed empirically (Albuquerque, 2012). In addition, the idiosyncratic volatility of both assets is the same, and the underperformance of the more positively skewed asset can thus not be related to the idiosyncratic volatility puzzle (Ang et al., 2006). Therefore, we consider in Panel C an economy in which—consistent with empirically observed market capitalizations (Kumar, 2009)—the supply of the high volatility, high skewness asset X is lower than the supply of asset Y ($\bar{\omega}_X = 0.1$ vs. $\bar{\omega}_Y = 0.9$). As a result, the market excess return is negatively skewed and asset X has a higher idiosyncratic volatility than asset Y . Panel C is still able to reproduce the same takeaways as Panel B for $\rho \geq 0.4$.¹⁵

Skewness preference offers rich predictions to explain deviations from the CAPM—consistent with empirically documented deviations. Overall, the evidence in the table confirms the literature for which the relevant pricing metric is the co-skewness of the residual of the CAPM regression with the market (Harvey and Siddique, 2000; Back et al., 2018; Karehnke and de Roon, 2020). Indeed, this residual co-skewness—which is also referred to as systematic skewness—is reported in the table and it always has the opposite sign of alpha. Related theoretical studies have shown that more skewed assets underperform when investors optimally form their expectations (Brunnermeier et al., 2007), have CPT prefer-

¹⁴This is true generally because lotteries and insurance contracts offer positively skewed payoffs—a moderate cost and the possibility to receive a high payoff with a low probability.

¹⁵The co-skewness of X with the market is here negative for high correlations. However, the co-skewness of the residual of a CAPM regression with the market is positive. Because residual co-skewness is approximately proportional to the slope on the squared market excess return controlling for the market excess return, the evidence in Panel C still suggests that asset X provides insurance (controlling for the market excess return).

ences (Barberis and Huang, 2008), or have heterogeneous skewness preferences (Mitton and Vorkink, 2007). Empirically, Boyer et al. (2010) show that stocks with higher idiosyncratic skewness underperform relative to the CAPM. The table reports idiosyncratic skewness (in the form of the skewness of the residual from a CAPM regression). Asset X indeed has positive idiosyncratic skewness, while asset Y has negative idiosyncratic skewness. The evidence thus also points to a negative relation between idiosyncratic skewness and alphas, although the link is less perfect. In Panel C both assets' alphas have the opposite sign of idiosyncratic skewness. However, they are increasing in idiosyncratic skewness—partly contradicting a preference for idiosyncratic skewness. At the same time, these alphas have the opposite sign and are decreasing in residual skewness—in line with a preference for residual skewness. This is supported empirically by Langlois (2020) who shows that systematic/residual co-skewness is more robustly priced than idiosyncratic skewness.

3.4 Conditional risk measures

There has been a growing interest in bivariate risk measures—for example in the systemic risk literature (see Benoit et al., 2017, for a review). A two-risk focus is thereby attractive as it provides a good trade-off between tractability and relevance. Unlike the univariate case that focuses on a risk in isolation, the two risks can be correlated and potentially result from an aggregation of many individual risks.

This section uses the SBN distribution to analyze two of the most popular measures: ΔCoVaR (see Adrian and Brunnermeier, 2016) and CoES (see Acharya et al., 2017). Sophisticated econometric approaches have already been proposed to estimate these measures in practice. Thus, we rather focus on the comparative statics and ask how they vary as a function of the skewness and correlation of the underlying risks—assuming that these risks follow an SBN distribution. We provide support for this approach at the end of the section by showing that theoretical CoESs calculated with the SBN distribution are close to realized

CoESs.

3.4.1 ΔCoVaR

Adrian and Brunnermeier (2016) propose a measure of systemic risk, ΔCoVaR , which is defined as the difference between the VaR of the financial system conditional on an institution being under distress and this VaR conditional on the same institution being in its median state. Hence, ΔCoVaR measures the tail-dependency between two random variables X and Y . Below, we derive its formula for the SBN distribution and evaluate its sensitivity to skewness and correlation. For the ease of exposition, we assume in the derivation and illustration that X and Y both have zero mean and unit variance. ΔCoVaR at the $q\%$ probability is defined as

$$\Delta\text{CoVaR}_q^{Y|X} = \text{CoVaR}_q^{Y|X=\text{VaR}_q^X} - \text{CoVaR}_q^{Y|X=\text{VaR}_{50}^X},$$

where VaR_q^X is the VaR of X at the $q\%$ probability and $\text{CoVaR}_q^{Y|C(X)}$ is the VaR of Y at the $q\%$ probability conditional on some event $C(X)$ of X .¹⁶

To compute ΔCoVaR , we first determine the parameters m_X , m_Y , $s_{X,1}$, $s_{X,2}$, $s_{Y,1}$, $s_{Y,2}$, and c such that X and Y have zero means, unit variances, the target skewnesses, and the target correlation. Second, we compute numerically the two required VaRs of X .¹⁷ Third, we compute the conditional VaRs of Y using the distribution of Y conditional on X reported in Appendix 3.A.4. In this third step, four cases arise depending on whether the VaR of X is below or above m_X and whether the conditional VaR of Y is below or above m_Y . The next proposition characterizes $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ in all four cases.

¹⁶Formally, VaR_q^X is defined by $\mathbb{P}(X \leq \text{VaR}_q^X) = q\%$ and $\text{CoVaR}_q^{Y|C(X)}$ by $\mathbb{P}(Y|C(X) \leq \text{CoVaR}_q^{Y|C(X)}) = q\%$.

¹⁷We obtain VaR_q^X and VaR_{50}^X numerically by solving $\int_{-\infty}^{\text{VaR}_z^X} g_X(x; \Omega) dx = z\%$ for $z = q, 50$, where g_X is the marginal distribution of X reported in Appendix 3.A.3.

Proposition 3.3 (CoVaR for the SBN distribution). *Assume that X and Y follow the SBN distribution. Let $M_{q,1} = \frac{\text{VaR}_q^X - m_X}{s_{X,1}}$ and $M_{q,2} = \frac{\text{VaR}_q^X - m_X}{s_{X,2}}$. In addition, let $d_{1,1}$, $d_{1,2}$, $d_{2,1}$, and $d_{2,2}$ denote the scaling functions whose expressions are in Appendix 3.A.4, and let Φ^{-1} denote the inverse of the standard normal cumulative distribution function (cdf).*

1. *If $\text{VaR}_q^X \leq m_X$ and $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} \leq m_Y$, then*

$$\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} = m_Y + cs_{Y,1}M_{q,1} + \sqrt{1-c^2}s_{Y,1}\Phi^{-1}\left(\frac{q\%}{d_{1,1}(\text{VaR}_q^X)}\right). \quad (3.4)$$

2. *If $\text{VaR}_q^X \leq m_X$ and $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} > m_Y$, then*

$$\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} = m_Y + cs_{Y,2}M_{q,1} + \sqrt{1-c^2}s_{Y,2}\Phi^{-1}\left(\frac{q\%}{d_{1,2}(\text{VaR}_q^X)} + \left(1 - \frac{d_{1,1}(\text{VaR}_q^X)}{d_{1,2}(\text{VaR}_q^X)}\right)\Phi\left(-\frac{cM_{q,1}}{\sqrt{1-c^2}}\right)\right). \quad (3.5)$$

3. *If $\text{VaR}_q^X > m_X$ and $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} \leq m_Y$, then*

$$\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} = m_Y + cs_{Y,1}M_{q,2} + \sqrt{1-c^2}s_{Y,1}\Phi^{-1}\left(\frac{q\%}{d_{2,1}(\text{VaR}_q^X)}\right). \quad (3.6)$$

4. *If $\text{VaR}_q^X > m_X$ and $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} > m_Y$, then*

$$\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} = m_Y + cs_{Y,2}M_{q,2} + \sqrt{1-c^2}s_{Y,2}\Phi^{-1}\left(\frac{q\%}{d_{2,2}(\text{VaR}_q^X)} + \left(1 - \frac{d_{2,1}(\text{VaR}_q^X)}{d_{2,2}(\text{VaR}_q^X)}\right)\Phi\left(-\frac{cM_{q,2}}{\sqrt{1-c^2}}\right)\right). \quad (3.7)$$

Equations (3.4) to (3.7) can be used to compute ΔCoVaR . In the most common and relevant case, ΔCoVaR is calculated for a low q and for left-skewed risks. It is thus likely that the two VaRs of X are not above m_X and the two conditional VaRs of Y are not above m_Y . The ΔCoVaR for this case is stated in the next corollary. The other cases can easily be constructed with Proposition 3.3 and are omitted to save space.

Corollary 3.1. *Assume that the skewness of X is negative or zero and that q is sufficiently small such that $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ and $\text{CoVaR}_q^{Y|X=\text{VaR}_{50}^X}$ are below or equal to m_Y . Then*

$$\Delta\text{CoVaR}_q^{Y|X} = c s_{Y,1} \frac{\text{VaR}_q^X - \text{VaR}_{50}^X}{s_{X,1}} + \sqrt{1 - c^2 s_{Y,1}} \left[\Phi^{-1} \left(\frac{q\%}{d_{1,1}(\text{VaR}_q^X)} \right) - \Phi^{-1} \left(\frac{q\%}{d_{1,1}(\text{VaR}_{50}^X)} \right) \right]. \quad (3.8)$$

When the two assets are uncorrelated, i.e., $c = 0$, $d_{1,1}$ is a constant function, leading to ΔCoVaR equals zero as both terms in (3.8) vanish. When the two assets are correlated and have zero skewnesses, $d_{1,1}$ is also a constant function. Hence, the second term in (3.8) vanishes, and (3.8) simplifies to the formula given by Equation (II.6) in Adrian and Brunnermeier (2016).

Table 3.7 reports ΔCoVaR at the 5% probability in the case of Corollary 3.1 and other possible cases. Each panel considers a different level of skewness for X and Y . The correlation varies across columns from weakly negative ($\rho = -0.2$) to strongly positive ($\rho = 0.8$).

Insert Table 3.7 here.

The table shows that the relation between ΔCoVaR and correlation obtained for the BN distribution generalizes to the case with skewness. First, when the two assets are uncorrelated, ΔCoVaR is zero. This is because the conditional VaR of Y then corresponds to the unconditional one.¹⁸ Second, ΔCoVaR is decreasing in and has the opposite sign of correlation. The intuition is that, ceteris paribus, e.g., for positive correlations, X and Y tend to move together lowering $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ and increasing $\text{CoVaR}_q^{Y|X=\text{VaR}_{50}^X}$ and thus leading to a negative ΔCoVaR .

Comparing the values across panels, we see on the one hand that the higher the skewness of Y , the lower ΔCoVaR is in absolute value. This happens for positive correlations because when the skewness of Y increases, high positive realizations of Y become more likely increasing both $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ and, to a lower extent, $\text{CoVaR}_q^{Y|X=\text{VaR}_{50}^X}$ and thus reducing their difference. On the other hand, the skewness of X has only a limited impact. Indeed,

¹⁸Formally, this case leads to $\text{VaR}_q^Y = \text{CoVaR}_q^{Y|X=\text{VaR}_q^X} = \text{CoVaR}_q^{Y|X=\text{VaR}_{50}^X}$.

$\Delta\text{CoVaR}_q^{Y|X}$ assesses the risk of Y conditional on X , but it is not a risk measure of X . The fact that the assets are only partially correlated also mitigates this impact. When, e.g., the correlation is 0.4 and the skewness of X worsens from 0 to -0.6 , ΔCoVaR even increases slightly from -0.66 in the symmetric case to -0.63 in Panel E. These comparative statics corroborate Adrian and Brunnermeier's (2016) empirical finding that there is only a weak correlation between an institution's VaR_q^X and the $\Delta\text{CoVaR}_q^{Y|X}$ of the financial system. In fact, our analysis shows that the financial system's $\Delta\text{CoVaR}_q^{Y|X}$ depends primarily on its own skewness Y . This skewness has a large quantitative impact. For, e.g., $\rho = 0.4$, ΔCoVaR increases by over 25% in absolute value when the skewness of Y worsens from 0 to -0.6 in Panel E' ($\Delta\text{CoVaR} = -0.83$).

3.4.2 CoES

Acharya et al. (2017) propose a measure of systemic expected shortfall and show theoretically and empirically that it is suitable to measure systemic risk. Formally, this metric—that we denote CoES—is defined as the expected value of Y given that X is in distress

$$\text{CoES}_q(Y) = \mathbb{E}(Y|X \leq \text{VaR}_q^X),$$

where we use the same notation as in Section 3.4.1. The next proposition provides the explicit expression for the CoES.

Proposition 3.4 (CoES for the SBN distribution). *Assume that X and Y follow the SBN distribution, and let ϕ and Φ denote the standard normal pdf and cdf, respectively. We have*

1. *If $\text{VaR}_q^X \leq m_X$, then*

$$\begin{aligned} \text{CoES}_q(Y) = \frac{1}{q\%} & \left[m_Y [(\lambda_{1,1} - \lambda_{1,2}) L(\bar{0}, \overline{M_{q,1}}) + \lambda_{1,2} \Phi(M_{q,1})] + \frac{\lambda_{1,2} s_{Y,2} - \lambda_{1,1} s_{Y,1}}{\sqrt{2\pi}} \Phi\left(\frac{M_{q,1}}{\sqrt{1-c^2}}\right) \right. \\ & \left. + c\phi(M_{q,1}) \left[(\lambda_{1,2} s_{Y,2} - \lambda_{1,1} s_{Y,1}) \Phi\left(-\frac{cM_{q,1}}{\sqrt{1-c^2}}\right) - \lambda_{1,2} s_{Y,2} \right] \right]. \end{aligned} \quad (3.9)$$

2. If $\text{VaR}_q^X > m_X$, then

$$\begin{aligned} \text{CoES}_q(Y) = & \frac{1}{q\%} \left[m_Y \left(\lambda_{1,1} \left(\frac{1}{4} + \frac{\sin^{-1}(c)}{2\pi} \right) + \lambda_{1,2} \frac{\cos^{-1}(c)}{2\pi} + (\lambda_{2,1} - \lambda_{2,2}) \left(L(\bar{0}, \overline{M_{q,2}}) - \left(\frac{1}{4} + \frac{\sin^{-1}(c)}{2\pi} \right) \right) \right) \right. \\ & + \lambda_{2,2} \left(\Phi(M_{q,2}) - \frac{1}{2} \right) \left. + \frac{(1-c)(\lambda_{1,2} - \lambda_{2,2})s_{Y,2} - (1+c)(\lambda_{1,1} - \lambda_{2,1})s_{Y,1}}{2\sqrt{2\pi}} \right. \\ & \left. + \frac{\lambda_{2,2}s_{Y,2} - \lambda_{2,1}s_{Y,1}}{\sqrt{2\pi}} \Phi\left(\frac{M_{q,2}}{\sqrt{1-c^2}}\right) + c\phi(M_{q,2}) \left[(\lambda_{2,2}s_{Y,2} - \lambda_{2,1}s_{Y,1}) \Phi\left(-\frac{cM_{q,2}}{\sqrt{1-c^2}}\right) - \lambda_{2,2}s_{Y,2} \right] \right]. \end{aligned}$$

Both expressions use $L(\bar{x}, \bar{y}) \equiv \mathcal{L}(h_1, h_2, k_1, k_2)$, where $h_1 = k_1 = -\infty$, $h_2 = x$, $k_2 = y$, and \mathcal{L} is as in Ang and Chen (2002) given by

$$\mathcal{L}(h_1, h_2, k_1, k_2) = \int_{h_1}^{h_2} \int_{k_1}^{k_2} \frac{1}{2\pi\sqrt{1-c^2}} \exp\left(-\frac{x^2 - 2cxy + y^2}{2(1-c^2)}\right) dx dy.$$

Note that in the symmetric case assuming that VaR_q^X is below its mode, $\lambda_{1,1} = \lambda_{1,2} = 1$ and $s_{Y,1} = s_{Y,2} \equiv s_Y$. Thus, (3.9) simplifies to $\text{CoES}_q(Y) = m_Y - \frac{1}{q\%} c s_Y \phi(M_q)$.

Table 3.8 reports CoES at the 5% probability and is constructed in a similar fashion than the previous table. It shows that CoES is zero when the correlation is zero. This is expected because the expectation of Y is zero. In addition, CoES decreases with correlation: the higher the correlation, the lower the mean of Y conditional on a negative tail event of X .

Insert Table 3.8 here.

Comparing the values across panels, CoES increases with the skewness of Y . As the skewness increases, high values of Y are more likely to occur; hence, the conditional mean of Y increases. For positive correlations, a higher skewness of X instead tends to lower CoES. In fact, a higher skewness of X means that very low negative realizations of X occur less often. Hence, ceteris paribus, more negative realizations of Y have to occur for given tail events of X , leading to a lower CoES.¹⁹ This effect can be large. For example, increasing

¹⁹This effect is mitigated when the skewness of Y is low. In such cases, many negative realizations of Y already happen when X has tail events, and there is no big change induced by the higher skewness of X .

the skewness of X from zero to 0.6 lowers the CoES by almost 15% for $\rho = 0.4$ (Benchmark case vs. Panel D).

Finally, we compute empirical CoESs and compare them to those implied by the distributions studied in Section 3.2.2.²⁰ Exploiting the same data, we use large stocks' excess returns as the conditioning variable X and the excess returns on small stocks, bonds, commodities, and emerging markets stocks as the variable of interest Y . We first standardize X and Y to focus on the effects of higher-order moments and facilitate comparisons. We then compute for each of these pairs the empirical CoES and the theoretical CoESs under one of the four distributions at the 5%, 10%, and 25% probabilities. Table 3.9 reports the results along with the sum of squared differences between theoretical and empirical CoESs for each probability and distribution.

Insert Table 3.9 here.

SBN has the lowest sum of squared differences among the four distributions for all probabilities and thus seems to provide a good fit on average despite its simplicity. At the 25% level, SBN and BST still have a lower sum of squared differences than BN but the difference is tiny (0.005 and 0.006 vs. 0.009). At lower probabilities (5% and 10%), however, the magnitudes of sum of squared differences are considerably lower for the skewed distributions than for BN. We conclude that skewness is important for empirical CoES levels, especially at low probabilities.

3.5 Conclusion

This paper uses a two-asset framework to study the joint effects of skewness and correlation. Our tool is the SBN distribution that is based on scaled BN distributions with different standard deviations and that allows for many explicit calculations. Using index excess returns,

²⁰The choice of computing CoES rather than ΔCoVaR is motivated by the fact that CoES is easier to compute. More importantly, the CoES provides a clearer and more direct test of the fit of the underlying distributions.

we show that this distribution provides a good empirical fit and that its fit is best in the conditional left tail compared to its competitors. This corroborates its usefulness as a tool for studying and modeling risk. More importantly, our theoretical framework is sufficiently rich to examine, e.g., the utility gains of considering skewness in portfolio choice, equilibrium deviations from the standard CAPM, and the sensitivity of recently proposed conditional risk metrics to skewness. For instance, we show that skewness preference provides a unified explanation for well known empirical features of risk premia such as the underperformance of stocks with high beta, high maximum returns, and high idiosyncratic skewness.

We focus on a simple framework and consider only one deviation from the bivariate normal framework—skewness. For future research, it would be interesting to examine other deviations such as asymmetric correlations in returns. Such an extension might provide additional asset pricing and portfolio implications and provide more intuition on the sources of co-skewness in returns.

3.A Details on the distribution

3.A.1 Calculating the scaling factors

The SBN distribution is a combination of scaled BN distributions. Formally, the pdf g is

$$g(x, y; \Omega) = \begin{cases} \lambda_{1,1}f(x, y; m_X, m_Y, s_{X,1}, s_{Y,1}, c) \equiv \lambda_{1,1}f_{1,1}(x, y), & \text{if } x \leq m_X \text{ and } y \leq m_Y, \\ \lambda_{1,2}f(x, y; m_X, m_Y, s_{X,1}, s_{Y,2}, c) \equiv \lambda_{1,2}f_{1,2}(x, y), & \text{if } x \leq m_X, \text{ and } y > m_Y, \\ \lambda_{2,1}f(x, y; m_X, m_Y, s_{X,2}, s_{Y,1}, c) \equiv \lambda_{2,1}f_{2,1}(x, y), & \text{if } x > m_X \text{ and } y \leq m_Y, \\ \lambda_{2,2}f(x, y; m_X, m_Y, s_{X,2}, s_{Y,2}, c) \equiv \lambda_{2,2}f_{2,2}(x, y), & \text{if } x > m_X \text{ and } y > m_Y, \end{cases}$$

where $\lambda_{1,1}$, $\lambda_{1,2}$, $\lambda_{2,1}$, and $\lambda_{2,2}$ are designed such that g is continuous (1), integrates to one (2), and is non-negative (3).

We first look at Property (1). Because the pdf of a BN distribution is continuous, we only need to ensure that g is continuous at the limit points. Formally, the left-hand side of the following equivalences must be satisfied²¹

$$\lambda_{1,1}f_{1,1}(x, m_Y) = \lambda_{1,2}f_{1,2}(x, m_Y), \quad \forall x \in] - \infty, m_X] \quad \Leftrightarrow \quad \frac{\lambda_{1,1}}{\lambda_{1,2}} = \frac{s_{Y,1}}{s_{Y,2}}, \quad (3.10)$$

$$\lambda_{2,1}f_{2,1}(x, m_Y) = \lambda_{2,2}f_{2,2}(x, m_Y), \quad \forall x \in]m_X, +\infty[\quad \Leftrightarrow \quad \frac{\lambda_{2,1}}{\lambda_{2,2}} = \frac{s_{Y,1}}{s_{Y,2}}, \quad (3.11)$$

$$\lambda_{1,1}f_{1,1}(m_X, y) = \lambda_{2,1}f_{2,1}(m_X, y), \quad \forall y \in] - \infty, m_Y] \quad \Leftrightarrow \quad \frac{\lambda_{1,1}}{\lambda_{2,1}} = \frac{s_{X,1}}{s_{X,2}}, \quad (3.12)$$

$$\lambda_{1,2}f_{1,2}(m_X, y) = \lambda_{2,2}f_{2,2}(m_X, y), \quad \forall y \in]m_Y, +\infty[\quad \Leftrightarrow \quad \frac{\lambda_{1,2}}{\lambda_{2,2}} = \frac{s_{X,1}}{s_{X,2}}, \quad (3.13)$$

²¹We also need that $\lambda_{1,1}f_{1,1}(m_X, m_Y) = \lambda_{1,2}f_{1,2}(m_X, m_Y) = \lambda_{2,1}f_{2,1}(m_X, m_Y) = \lambda_{2,2}f_{2,2}(m_X, m_Y)$. This fifth condition is automatically verified when the four others are.

where the equivalences are obtained by exploiting the properties of $f_{\cdot,\cdot}$.

We now turn to Property (2): g needs to integrate to one. Replacing g by its expression and rewriting it in terms of a probability, we have

$$\begin{aligned} & \lambda_{1,1}\mathbb{P}(X_1 \leq m_X, Y_1 \leq m_Y) + \lambda_{2,1}\mathbb{P}(X_3 > m_X, Y_3 \leq m_Y) \\ & + \lambda_{1,2}\mathbb{P}(X_2 \leq m_X, Y_2 > m_Y) + \lambda_{2,2}\mathbb{P}(X_4 > m_X, Y_4 > m_Y) = 1, \end{aligned}$$

where (X_1, Y_1) , (X_2, Y_2) , (X_3, Y_3) , and (X_4, Y_4) follow the four different BN distributions that we combine. The probabilities in this equation are quadrant probabilities of BN distributions. Thus, we obtain (see, e.g., Johnson and Kotz, 1972)

$$(\lambda_{1,1} + \lambda_{2,2}) \left(\frac{1}{4} + \frac{\sin^{-1}(c)}{2\pi} \right) + (\lambda_{1,2} + \lambda_{2,1}) \frac{\cos^{-1}(c)}{2\pi} = 1. \quad (3.14)$$

Combining Equations (3.10) to (3.14) and solving the system, we obtain

$\lambda_{1,1} = \Delta s_{X,1}s_{Y,1}$, $\lambda_{1,2} = \Delta s_{X,1}s_{Y,2}$, $\lambda_{2,1} = \Delta s_{X,2}s_{Y,1}$, and $\lambda_{2,2} = \Delta s_{X,2}s_{Y,2}$, with

$$\Delta = \frac{4\pi}{(s_{X,1}s_{Y,1} + s_{X,2}s_{Y,2}) (\pi + 2\sin^{-1}(c)) + 2(s_{X,1}s_{Y,2} + s_{X,2}s_{Y,1}) \cos^{-1}(c)}.$$

Finally, $\lambda_{1,1}$, $\lambda_{1,2}$, $\lambda_{2,1}$, and $\lambda_{2,2}$ are non-negative and thus Property (3) is satisfied.

3.A.2 Moments of the distribution

We express the moments and co-moments of the distribution, i.e., the means μ_X and μ_Y , the variances σ_X^2 and σ_Y^2 , the covariance $\sigma_{X,Y}$, the (standardized) skewnesses γ_X and γ_Y , and the (standardized) co-skewnesses $\gamma_{Y,X}$ and $\gamma_{X,Y}$ of X and Y as functions of m_X , m_Y , $s_{X,1}$, $s_{X,2}$, $s_{Y,1}$, $s_{Y,2}$, and c .²² To do so, we first compute the moment generating function

²²Recall that the (standardized) co-skewnesses are given by $\gamma_{Y,X} = \mathbb{E}[(Y - \mathbb{E}[Y])(X - \mathbb{E}[X])^2] / (\sigma_Y\sigma_X^2)$ and $\gamma_{X,Y} = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])^2] / (\sigma_X\sigma_Y^2)$.

MGF_{SBN} of the distribution. By definition, it is given by

$$\begin{aligned} MGF_{SBN}(t_1, t_2) &= \mathbb{E}(\exp(t_1X + t_2Y)) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \exp(t_1x + t_2y) g(x, y; \Omega) \, dx dy \\ &= \lambda_{1,1} \int_{-\infty}^{m_Y} \int_{-\infty}^{m_X} \exp(t_1x + t_2y) f_{1,1}(x, y) \, dx dy + \lambda_{1,2} \int_{m_Y}^{+\infty} \int_{-\infty}^{m_X} \exp(t_1x + t_2y) f_{1,2}(x, y) \, dx dy \\ &+ \lambda_{2,1} \int_{-\infty}^{m_Y} \int_{m_X}^{+\infty} \exp(t_1x + t_2y) f_{2,1}(x, y) \, dx dy + \lambda_{2,2} \int_{m_Y}^{+\infty} \int_{m_X}^{+\infty} \exp(t_1x + t_2y) f_{2,2}(x, y) \, dx dy. \end{aligned}$$

Making the changes of variable $u = \frac{x - m_X}{s_{X,1}} - (t_1s_{X,1} + t_2s_{Y,1}c)$ and $v = \frac{y - m_Y}{s_{Y,1}} - (t_2s_{Y,1} + t_1s_{X,1}c)$ and rearranging the terms, we obtain that the first quadrant integral equals

$$\begin{aligned} &\exp(t_1m_X + t_2m_Y) \times \\ &\mathcal{L}(-\infty, -(t_2s_{Y,1} + t_1s_{X,1}c), -\infty, -(t_1s_{X,1} + t_2s_{Y,1}c)) \exp\left(\frac{1}{2} [t_1^2s_{X,1}^2 + t_2^2s_{Y,1}^2 + 2t_1t_2s_{X,1}s_{Y,1}c]\right), \end{aligned}$$

where \mathcal{L} is given as in Ang and Chen (2002) by

$$\mathcal{L}(h_1, h_2, k_1, k_2) = \int_{h_1}^{h_2} \int_{k_1}^{k_2} \frac{1}{2\pi\sqrt{1-c^2}} \exp\left(-\frac{x^2 - 2cxy + y^2}{2(1-c^2)}\right) \, dx dy.$$

The other quadrant integrals are obtained analogously. Hence, the moment generating function is

$$\begin{aligned} MGF_{SBN}(t_1, t_2) &= \\ &\exp(t_1m_X + t_2m_Y) \times \left(\sum_{i=1}^2 \sum_{j=1}^2 \lambda_{i,j} L_{i,j}(-t_1, -t_2) \exp\left(\frac{1}{2} [t_1^2s_{X,i}^2 + t_2^2s_{Y,j}^2 + 2t_1t_2s_{X,i}s_{Y,j}c]\right) \right), \end{aligned}$$

where

$$L_{1,1}(t_1, t_2) = \mathcal{L}(-\infty, t_2s_{Y,1} + t_1s_{X,1}c, -\infty, t_1s_{X,1} + t_2s_{Y,1}c),$$

$$L_{1,2}(t_1, t_2) = \mathcal{L}(t_2s_{Y,2} + t_1s_{X,1}c, +\infty, -\infty, t_1s_{X,1} + t_2s_{Y,2}c),$$

$$L_{2,1}(t_1, t_2) = \mathcal{L}(-\infty, t_2s_{Y,1} + t_1s_{X,2}c, t_1s_{X,2} + t_2s_{Y,1}c, +\infty), \text{ and}$$

$$L_{2,2}(t_1, t_2) = \mathcal{L}(t_2s_{Y,2} + t_1s_{X,2}c, +\infty, t_1s_{X,2} + t_2s_{Y,2}c, +\infty).$$

We then obtain the moments using the derivatives of MGF_{SBN} computed at $(t_1, t_2) = (0, 0)$. The moments for X are (expressions for Y are similar)

$$\mu_X = m_X + \frac{\lambda_{1,1}s_{X,1}A_X}{2\sqrt{2\pi}}, \quad (3.15)$$

$$\sigma_X^2 = \lambda_{1,1}s_{X,1}^2 \left(B_X - \frac{\lambda_{1,1}A_X^2}{8\pi} \right), \quad (3.16)$$

$$\gamma_X = \frac{C_X + \frac{\lambda_{1,1}^2 A_X^3}{4\pi} - 3\lambda_{1,1}A_X B_X}{2\sqrt{2\pi}\sqrt{\lambda_{1,1}} \left(B_X - \frac{\lambda_{1,1}A_X^2}{8\pi} \right)^{3/2}}, \quad (3.17)$$

where

$$A_X = (1 + z_X^2)(-1 + z_Y)c + (-1 + z_X^2)(1 + z_Y),$$

$$B_X = (1 + z_X^3 z_Y) \left(\frac{1}{4} + \frac{\sin^{-1}(c)}{2\pi} + \frac{c\sqrt{1-c^2}}{2\pi} \right) + (z_X^3 + z_Y) \left(\frac{\cos^{-1}(c)}{2\pi} - \frac{c\sqrt{1-c^2}}{2\pi} \right),$$

$$C_X = (1 + z_X^4)(1 - z_Y)(c^3 - 3c) + 2(-1 + z_X^4)(1 + z_Y),$$

and $z_X = \frac{s_{X,2}}{s_{X,1}}$, and A_Y, B_Y, C_Y , and z_Y are defined similarly. The co-moments are

$$\sigma_{X,Y} = \lambda_{1,1}s_{X,1}s_{Y,1} \left(B_{X,Y} - \frac{\lambda_{1,1}A_X A_Y}{8\pi} \right), \quad (3.18)$$

$$\gamma_{Y,X} = \frac{C_{X,Y} + \frac{\lambda_{1,1}^2 A_X^2 A_Y}{4\pi} - 2\lambda_{1,1}A_X B_{X,Y} - \lambda_{1,1}A_Y B_X}{2\sqrt{2\pi}\sqrt{\lambda_{1,1}} \left(B_X - \frac{\lambda_{1,1}A_X^2}{8\pi} \right) \sqrt{B_Y - \frac{\lambda_{1,1}A_Y^2}{8\pi}}}, \quad (3.19)$$

$$\gamma_{X,Y} = \frac{C_{Y,X} + \frac{\lambda_{1,1}^2 A_Y^2 A_X}{4\pi} - 2\lambda_{1,1}A_Y B_{X,Y} - \lambda_{1,1}A_X B_Y}{2\sqrt{2\pi}\sqrt{\lambda_{1,1}} \left(B_Y - \frac{\lambda_{1,1}A_Y^2}{8\pi} \right) \sqrt{B_X - \frac{\lambda_{1,1}A_X^2}{8\pi}}}, \quad (3.20)$$

where

$$B_{X,Y} = (1 + z_X^2 z_Y^2) \left(c \left(\frac{1}{4} + \frac{\sin^{-1}(c)}{2\pi} \right) + \frac{\sqrt{1-c^2}}{2\pi} \right) + (z_X^2 + z_Y^2) \left(c \frac{\cos^{-1}(c)}{2\pi} - \frac{\sqrt{1-c^2}}{2\pi} \right),$$

$$C_{X,Y} = (1 + z_X^3)(-1 + z_Y^2)(1 + c^2) + 2(-1 + z_X^3)(1 + z_Y^2)c, \text{ and}$$

$$C_{Y,X} = (1 + z_Y^3)(-1 + z_X^2)(1 + c^2) + 2(-1 + z_Y^3)(1 + z_X^2)c.$$

3.A.3 Marginal distribution

We derive the pdf g_X of the marginal distribution of X . We have

$$g_X(x; \Omega) = \int_{\mathbb{R}} g(x, y; \Omega) dy = \begin{cases} \lambda_{1,1} \int_{-\infty}^{m_Y} f_{1,1}(x, y; \Omega) dy + \lambda_{1,2} \int_{m_Y}^{+\infty} f_{1,2}(x, y; \Omega) dy, & \text{if } x \leq m_X, \\ \lambda_{2,1} \int_{-\infty}^{m_Y} f_{2,1}(x, y; \Omega) dy + \lambda_{2,2} \int_{m_Y}^{+\infty} f_{2,2}(x, y; \Omega) dy, & \text{if } x > m_X. \end{cases}$$

We compute the four integrals. To compute the integral with $f_{1,1}$, we use the change of variable $z = \frac{\frac{y-m_Y}{s_{Y,1}} - c\frac{x-m_X}{s_{X,1}}}{\sqrt{1-c^2}}$ and obtain

$$\int_{-\infty}^{m_Y} f_{1,1}(x, y; \Omega) dy = \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,1}}\right) f_X(x; m_X, s_{X,1}),$$

where Φ is the standard normal cdf and $f_X(\cdot; \mu, \sigma)$ is the pdf of a normal distribution with mean μ and standard deviation σ . Using similar changes of variables, we obtain

$$g_X(x; \Omega) = \begin{cases} \left(\lambda_{1,2} + (\lambda_{1,1} - \lambda_{1,2}) \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,1}}\right) \right) f_X(x; m_X, s_{X,1}), & \text{if } x \leq m_X, \\ \left(\lambda_{2,2} + (\lambda_{2,1} - \lambda_{2,2}) \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,2}}\right) \right) f_X(x; m_X, s_{X,2}), & \text{if } x > m_X. \end{cases}$$

In the specific case of two uncorrelated variables (i.e., $c = 0$), we have

$$g_X(x; \Omega) = \begin{cases} \frac{2s_{X,1}}{s_{X,1} + s_{X,2}} f_X(x; m_X, s_{X,1}), & \text{if } x \leq m_X, \\ \frac{2s_{X,2}}{s_{X,1} + s_{X,2}} f_X(x; m_X, s_{X,2}), & \text{if } x > m_X. \end{cases}$$

This corresponds to the pdf of the (univariate) split normal distribution (Wallis, 2014) with the parameterization of De Roon and Karehnke (2017). Thus, when $c = 0$, the pdf of the SBN distribution with parameters $(m_X, m_Y, s_{X,1}, s_{X,2}, s_{Y,1}, s_{Y,2}, 0)$ is given by the product of two split normal pdfs with parameters $(m_X, s_{X,1}, s_{X,2})$ and $(m_Y, s_{Y,1}, s_{Y,2})$, respectively.

3.A.4 Conditional distribution

We derive the pdf $g_{Y|X=x}$ of the conditional distribution of Y (i.e., the distribution of Y conditional on X). By definition, we have $g_{Y|X=x}(y; \Omega) = \frac{g(x, y; \Omega)}{g_X(x; \Omega)}$.

Defining

$$\begin{aligned} d_{1,1}(x) &= \frac{\lambda_{1,1}}{\lambda_{1,2} + (\lambda_{1,1} - \lambda_{1,2}) \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,1}}\right)}, \\ d_{1,2}(x) &= \frac{\lambda_{1,2}}{\lambda_{1,2} + (\lambda_{1,1} - \lambda_{1,2}) \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,1}}\right)}, \\ d_{2,1}(x) &= \frac{\lambda_{2,1}}{\lambda_{2,2} + (\lambda_{2,1} - \lambda_{2,2}) \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,2}}\right)}, \text{ and} \\ d_{2,2}(x) &= \frac{\lambda_{2,2}}{\lambda_{2,2} + (\lambda_{2,1} - \lambda_{2,2}) \Phi\left(-\frac{c(x-m_X)}{\sqrt{1-c^2}s_{X,2}}\right)}, \end{aligned}$$

we easily obtain

$$g_{Y|X=x}(y; \Omega) = \begin{cases} d_{1,1}(x) \frac{f(x, y; m_X, m_Y, s_{X,1}, s_{Y,1}, c)}{f_X(x; m_X, s_{X,1})}, & \text{if } x \leq m_X \text{ and } y \leq m_Y, \\ d_{1,2}(x) \frac{f(x, y; m_X, m_Y, s_{X,1}, s_{Y,2}, c)}{f_X(x; m_X, s_{X,1})}, & \text{if } x \leq m_X \text{ and } y > m_Y, \\ d_{2,1}(x) \frac{f(x, y; m_X, m_Y, s_{X,2}, s_{Y,1}, c)}{f_X(x; m_X, s_{X,2})}, & \text{if } x > m_X \text{ and } y \leq m_Y, \\ d_{2,2}(x) \frac{f(x, y; m_X, m_Y, s_{X,2}, s_{Y,2}, c)}{f_X(x; m_X, s_{X,2})}, & \text{if } x > m_X \text{ and } y > m_Y. \end{cases} \quad (3.21)$$

3.B Proofs

Proof of Proposition 3.1. Using the moment generating function of the SBN derived in Appendix 3.A.2, the EU of the investor is

$$\begin{aligned} & \mathbb{E}(u(1 + r_f + \omega_X R_X + \omega_Y R_Y)) \\ &= -\exp(-\theta(1 + r_f + \omega_X m_X + \omega_Y m_Y)) \\ & \quad \times \left(\sum_{i=1}^2 \sum_{j=1}^2 \lambda_{i,j} L_{i,j}(\theta\omega_X, \theta\omega_Y) \exp\left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c]\right) \right). \end{aligned}$$

Because the problem is concave, we can determine the optimal portfolio weights from the first-order conditions. Formally, we first compute $\frac{\partial L_{i,j}(\theta\omega_X, \theta\omega_Y)}{\partial \omega_k}$ (for $i = 1, 2, j = 1, 2$, and $k = X, Y$) using the Leibniz integral rule. For, e.g., $i = j = 1$ and $k = X$, it equals²³

$$\theta s_{X,1} \left(c\phi(\theta(\omega_Y s_{Y,1} + \omega_X s_{X,1}c)) \Phi(\sqrt{1 - c^2\theta\omega_X s_{X,1}}) + \phi(\theta(\omega_X s_{X,1} + \omega_Y s_{Y,1}c)) \Phi(\sqrt{1 - c^2\theta\omega_Y s_{Y,1}}) \right),$$

where ϕ and Φ denote the standard normal pdf and cdf, respectively. We then define $l_{i,j}^k$ such that $l_{i,j}^k(\omega_X, \omega_Y) = -\frac{1}{\theta} \frac{\partial L_{i,j}(\theta\omega_X, \theta\omega_Y)}{\partial \omega_k}$.²⁴ Focusing on the optimal weight for X and using its first order condition, we then obtain

$$\begin{aligned} & \sum_{i=1}^2 \sum_{j=1}^2 (m_X - \theta(\omega_X s_{X,i}^2 + \omega_Y s_{X,i} s_{Y,j} c)) \exp\left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c]\right) \lambda_{i,j} L_{i,j}(\theta\omega_X, \theta\omega_Y) \\ & \quad + \sum_{i=1}^2 \sum_{j=1}^2 \exp\left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c]\right) \lambda_{i,j} l_{i,j}^k(\omega_X, \omega_Y) = 0. \quad (3.22) \end{aligned}$$

²³Other explicit computations are available upon request.

²⁴When X and Y are symmetric, $\sum_{i=1}^2 \sum_{j=1}^2 l_{i,j}^k(\omega_X, \omega_Y) = 0$ for $k = X, Y$.

Rearranging Equation (3.22) yields

$$\omega_X = \frac{m_X \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)}{\theta \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i}^2 \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)} + \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} l_{i,j}^X (\omega_X, \omega_Y)}{\theta \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i}^2 \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)} - \frac{c \omega_Y \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i} s_{Y,j} \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)}{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i}^2 \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y)}.$$

(3.23)

By symmetry, we obtain a similar expression for ω_Y

$$\omega_Y = \frac{m_Y \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)}{\theta \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{Y,j}^2 \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)} + \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} l_{i,j}^Y (\omega_X, \omega_Y)}{\theta \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{Y,j}^2 \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)} - \frac{c \omega_X \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i} s_{Y,j} \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y) \right)}{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{Y,j}^2 \lambda_{i,j} L_{i,j} (\theta \omega_X, \theta \omega_Y)}.$$

(3.24)

Plugging Equation (3.24) into Equation (3.23) gives an implicit equation for ω_X (and vice versa for ω_Y). Formally, we obtain that the optimal weights are solution of

$$\omega_X = \frac{\tilde{\mu}_X (\omega_X, \omega_Y) \tilde{\sigma}_Y^2 (\omega_X, \omega_Y) - \tilde{\mu}_Y (\omega_X, \omega_Y) \tilde{\sigma}_{X,Y} (\omega_X, \omega_Y)}{\theta \tilde{\sigma}_X^2 (\omega_X, \omega_Y) \tilde{\sigma}_Y^2 (\omega_X, \omega_Y) \left(1 - \frac{\tilde{\sigma}_{X,Y}^2 (\omega_X, \omega_Y)}{\tilde{\sigma}_X^2 (\omega_X, \omega_Y) \tilde{\sigma}_Y^2 (\omega_X, \omega_Y)} \right)}$$

and

$$\omega_Y = \frac{\tilde{\mu}_Y (\omega_X, \omega_Y) \tilde{\sigma}_X^2 (\omega_X, \omega_Y) - \tilde{\mu}_X (\omega_X, \omega_Y) \tilde{\sigma}_{X,Y} (\omega_X, \omega_Y)}{\theta \tilde{\sigma}_X^2 (\omega_X, \omega_Y) \tilde{\sigma}_Y^2 (\omega_X, \omega_Y) \left(1 - \frac{\tilde{\sigma}_{X,Y}^2 (\omega_X, \omega_Y)}{\tilde{\sigma}_X^2 (\omega_X, \omega_Y) \tilde{\sigma}_Y^2 (\omega_X, \omega_Y)} \right)},$$

where

$$\begin{aligned}
 \tilde{\mu}_X(\omega_X, \omega_Y) &= \frac{m_X \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} L_{i,j}(\theta \omega_X, \theta \omega_Y) \right)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)} \\
 &\quad + \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} l_{i,j}^X(\omega_X, \omega_Y)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)}, \\
 \tilde{\mu}_Y(\omega_X, \omega_Y) &= \frac{m_Y \left(\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} L_{i,j}(\theta \omega_X, \theta \omega_Y) \right)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)} \\
 &\quad + \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) \lambda_{i,j} l_{i,j}^Y(\omega_X, \omega_Y)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)}, \\
 \tilde{\sigma}_X^2(\omega_X, \omega_Y) &= \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i}^2 \lambda_{i,j} L_{i,j}(\theta \omega_X, \theta \omega_Y)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)}, \\
 \tilde{\sigma}_Y^2(\omega_X, \omega_Y) &= \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{Y,j}^2 \lambda_{i,j} L_{i,j}(\theta \omega_X, \theta \omega_Y)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)}, \text{ and} \\
 \tilde{\sigma}_{X,Y}(\omega_X, \omega_Y) &= \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right) s_{X,i} s_{Y,j} \lambda_{i,j} L_{i,j}(\theta \omega_X, \theta \omega_Y)}{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right)} c.
 \end{aligned}$$

Note that $\tilde{\mu}_X$, $\tilde{\mu}_Y$, $\tilde{\sigma}_X^2$, $\tilde{\sigma}_Y^2$, and $\tilde{\sigma}_{X,Y}$ are all scaled by

$$\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp \left(\frac{\theta^2}{2} [\omega_X^2 s_{X,i}^2 + \omega_Y^2 s_{Y,j}^2 + 2\omega_X \omega_Y s_{X,i} s_{Y,j} c] \right).$$

This ensures that we have $\tilde{\mu}_X(\omega_X, \omega_Y) = m_X$, $\tilde{\mu}_Y(\omega_X, \omega_Y) = m_Y$, $\tilde{\sigma}_X^2(\omega_X, \omega_Y) = s_X^2$, $\tilde{\sigma}_Y^2(\omega_X, \omega_Y) = s_Y^2$, and $\tilde{\sigma}_{X,Y}(\omega_X, \omega_Y) = c s_X s_Y$ in the normal case (i.e., when $s_X = s_{X,1} = s_{X,2}$ and $s_Y = s_{Y,1} = s_{Y,2}$). We thus call these functions the modified moment functions. ■

Proof of Proposition 3.2. Let denote respectively by $\bar{\mu}_X$, $\bar{\mu}_Y$, $\bar{\sigma}_X^2$, $\bar{\sigma}_Y^2$, and $\bar{\sigma}_{X,Y}$ the modified moment functions $\tilde{\mu}_X$, $\tilde{\mu}_Y$, $\tilde{\sigma}_X^2$, $\tilde{\sigma}_Y^2$, and $\tilde{\sigma}_{X,Y}$ defined in the proof of Proposition 3.1 and evaluated at the portfolio weights $\bar{\omega}_X$ and $\bar{\omega}_Y$ that clear the market. For example, $\bar{\mu}_X := \tilde{\mu}_X(\bar{\omega}_X, \bar{\omega}_Y)$ denotes the modified risk premium of X . Using Proposition 3.1, we then

have

$$\bar{\mu}_X = \bar{\omega}_X \theta \bar{\sigma}_X^2 \left(1 - \frac{\bar{\sigma}_{X,Y}^2}{\bar{\sigma}_X^2 \bar{\sigma}_Y^2} \right) + \bar{\mu}_Y \frac{\bar{\sigma}_{X,Y}}{\bar{\sigma}_Y^2} \quad \text{and} \quad \bar{\mu}_Y = \bar{\omega}_Y \theta \bar{\sigma}_Y^2 \left(1 - \frac{\bar{\sigma}_{X,Y}^2}{\bar{\sigma}_X^2 \bar{\sigma}_Y^2} \right) + \bar{\mu}_X \frac{\bar{\sigma}_{X,Y}}{\bar{\sigma}_X^2}.$$

Combining the two expressions, we obtain

$$\bar{\mu}_X = \theta (\bar{\omega}_X \bar{\sigma}_X^2 + \bar{\omega}_Y \bar{\sigma}_{X,Y}) \quad \text{and} \quad \bar{\mu}_Y = \theta (\bar{\omega}_Y \bar{\sigma}_Y^2 + \bar{\omega}_X \bar{\sigma}_{X,Y}). \quad (3.25)$$

Let us define the modified covariance between asset X (asset Y) and the market $\bar{\sigma}_{X,M} = \bar{\omega}_X \bar{\sigma}_X^2 + \bar{\omega}_Y \bar{\sigma}_{X,Y}$ ($\bar{\sigma}_{Y,M} = \bar{\omega}_Y \bar{\sigma}_Y^2 + \bar{\omega}_X \bar{\sigma}_{X,Y}$), the modified market risk premium $\bar{\mu}_M = \bar{\omega}_X \bar{\mu}_X + \bar{\omega}_Y \bar{\mu}_Y$, and the modified market variance $\bar{\sigma}_M^2 = \bar{\omega}_X \bar{\sigma}_{X,M} + \bar{\omega}_Y \bar{\sigma}_{Y,M}$. We have $\bar{\mu}_M = \theta \bar{\sigma}_M^2$ and we can rewrite (3.25) as

$$\bar{\mu}_X = \bar{\beta}_X \bar{\mu}_M \quad \text{and} \quad \bar{\mu}_Y = \bar{\beta}_Y \bar{\mu}_M, \quad (3.26)$$

where $\bar{\beta}_k = \frac{\bar{\sigma}_{k,M}}{\bar{\sigma}_M^2}$ is the modified beta for $k = X, Y$. From Equation (3.15) in Appendix 3.A.2 we also have for $k = X, Y$

$$m_k = \mu_k - \frac{\lambda_{1,1} s_{k,1} A_k}{2\sqrt{2\pi}}. \quad (3.27)$$

Thus, plugging (3.27) into the definition of $\bar{\mu}_X$, we obtain for $k = X, Y$

$$\bar{\mu}_k = \frac{\mu_k}{\xi} - \frac{\zeta_k}{\xi}, \quad (3.28)$$

where $\xi = \frac{\frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \exp\left(\frac{\theta^2}{2} [\bar{\omega}_X^2 s_{X,i}^2 + \bar{\omega}_Y^2 s_{Y,j}^2 + 2\bar{\omega}_X \bar{\omega}_Y s_{X,i} s_{Y,j} c]\right)}{\sum_{i=1}^2 \sum_{j=1}^2 \exp\left(\frac{\theta^2}{2} [\bar{\omega}_X^2 s_{X,i}^2 + \bar{\omega}_Y^2 s_{Y,j}^2 + 2\bar{\omega}_X \bar{\omega}_Y s_{X,i} s_{Y,j} c]\right)} \lambda_{i,j} L_{i,j}(\bar{\omega}_X, \bar{\omega}_Y)$ and

$$\zeta_k = \frac{\lambda_{1,1} s_{k,1} A_k}{2\sqrt{2\pi}} - \frac{\sum_{i=1}^2 \sum_{j=1}^2 \exp\left(\frac{\theta^2}{2} [\bar{\omega}_X^2 s_{X,i}^2 + \bar{\omega}_Y^2 s_{Y,j}^2 + 2\bar{\omega}_X \bar{\omega}_Y s_{X,i} s_{Y,j} c]\right) \lambda_{i,j} l_{i,j}^k(\bar{\omega}_X, \bar{\omega}_Y)}{\sum_{i=1}^2 \sum_{j=1}^2 \exp\left(\frac{\theta^2}{2} [\bar{\omega}_X^2 s_{X,i}^2 + \bar{\omega}_Y^2 s_{Y,j}^2 + 2\bar{\omega}_X \bar{\omega}_Y s_{X,i} s_{Y,j} c]\right) \lambda_{i,j} L_{i,j}(\bar{\omega}_X, \bar{\omega}_Y)}.$$

Combining (3.26) and (3.28) leads to

$$\mu_X = \bar{\beta}_X \bar{\mu}_M \xi + \zeta_X \quad \text{and} \quad \mu_Y = \bar{\beta}_Y \bar{\mu}_M \xi + \zeta_Y.$$

■

Proof of Proposition 3.3. $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ is defined such that

$$\int_{-\infty}^{\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}} g_{Y|X=\text{VaR}_q^X}(y; \Omega) dy = q\%.$$

Given the definition of $g_{Y|X=\text{VaR}_q^X}$ in Appendix 3.A.4, four cases need to be considered depending on the position of VaR_q^X in the marginal distribution of X and on the position of $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ in the conditional distribution of Y .

Proof of statement 1. If $\text{VaR}_q^X \leq m_X$ and $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} \leq m_Y$, we have

$$\int_{-\infty}^{\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}} d_{1,1}(\text{VaR}_q^X) \frac{f(\text{VaR}_q^X, y; m_X, m_Y, s_{X,1}, s_{Y,1}, c)}{f_X(\text{VaR}_q^X; m_X, s_{X,1})} dy = q\%.$$

Making the change of variable $z = \frac{y - m_Y - c \frac{\text{VaR}_q^X - m_X}{s_{X,1}}}{\sqrt{1 - c^2}}$ and doing some computations, we obtain

$$\text{CoVaR}_q^{Y|X=\text{VaR}_q^X} = m_Y + c s_{Y,1} \frac{\text{VaR}_q^X - m_X}{s_{X,1}} + \sqrt{1 - c^2} s_{Y,1} \Phi^{-1} \left(\frac{q\%}{d_{1,1}(\text{VaR}_q^X)} \right). \quad (3.29)$$

Proof of statement 2. If $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ is in the right part of the conditional distribution of Y and VaR_q^X is below its mode, $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ is

$$\begin{aligned} \int_{m_Y}^{\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}} d_{1,2}(\text{VaR}_q^X) \frac{f(\text{VaR}_q^X, y; m_X, m_Y, s_{X,1}, s_{Y,2}, c)}{f_X(\text{VaR}_q^X; m_X, s_{X,1})} dy = \\ q\% - \int_{-\infty}^{m_Y} d_{1,1}(\text{VaR}_q^X) \frac{f(\text{VaR}_q^X, y; m_X, m_Y, s_{X,1}, s_{Y,1}, c)}{f_X(\text{VaR}_q^X; m_X, s_{X,1})} dy. \end{aligned}$$

From this expression and using similar computations, we find that $\text{CoVaR}_q^{Y|X=\text{VaR}_q^X}$ equals

$$m_Y + cs_{Y,2} \frac{\text{VaR}_q^X - m_X}{s_{X,1}} + \sqrt{1-c^2} s_{Y,2} \Phi^{-1} \left(\frac{q\%}{d_{1,2}(\text{VaR}_q^X)} + \left(1 - \frac{d_{1,1}(\text{VaR}_q^X)}{d_{1,2}(\text{VaR}_q^X)} \right) \Phi \left(-\frac{c(\text{VaR}_q^X - m_X)}{\sqrt{1-c^2} s_{X,1}} \right) \right).$$

Proof of statements 3 and 4. The last two cases in which VaR_q^X is in the right part of the marginal distribution of X are computed in the same way. We just need to replace $s_{X,1}$ by $s_{X,2}$, $d_{1,1}(\text{VaR}_q^X)$ by $d_{2,1}(\text{VaR}_q^X)$, and $d_{1,2}(\text{VaR}_q^X)$ by $d_{2,2}(\text{VaR}_q^X)$.

Introducing $M_{q,1}$ and $M_{q,2}$, we find the expressions reported in Proposition 3.3. ■

Proof of Corollary 3.1. Taking the difference between (3.29) evaluated at $X = \text{VaR}_q^X$ and $X = \text{VaR}_{50}^X$ yields the result. ■

Proof of Proposition 3.4. $\text{CoES}_q(Y)$ is defined as

$$\text{CoES}_q(Y) = \mathbb{E}(Y|X \leq \text{VaR}_q^X) = \frac{1}{q\%} \int_{-\infty}^{+\infty} \int_{-\infty}^{\text{VaR}_q^X} yg(x, y) \, dx dy.$$

Proof of statement 1. If $\text{VaR}_q^X \leq m_X$, we have

$$\text{CoES}_q(Y) = \frac{1}{q\%} \left(\lambda_{1,1} \int_{-\infty}^{m_Y} \int_{-\infty}^{\text{VaR}_q^X} y f_{1,1}(x, y) \, dx dy + \lambda_{1,2} \int_{m_Y}^{+\infty} \int_{-\infty}^{\text{VaR}_q^X} y f_{1,2}(x, y) \, dx dy \right). \quad (3.30)$$

To compute the first integral of Equation (3.30), we make the following change of variables $u = \frac{x - m_X}{s_{X,1}}$ and $v = \frac{y - m_Y}{s_{Y,1}}$. The integral can therefore be rewritten as

$$\frac{s_{Y,1}}{2\pi\sqrt{1-c^2}} \int_{-\infty}^0 \int_{-\infty}^{\frac{\text{VaR}_q^X - m_X}{s_{X,1}}} v \exp\left(-\frac{u^2 + v^2 - 2cuv}{2(1-c^2)}\right) \, dudv + m_Y L\left(-\infty, 0, -\infty, \frac{\text{VaR}_q^X - m_X}{s_{X,1}}\right).$$

To compute the first term of this expression we make the change of variables $u = u$ and

$z = \frac{v - cu}{\sqrt{1 - c^2}}$. The first term is then given by

$$\begin{aligned} & \frac{s_{Y,1}\sqrt{1-c^2}}{2\pi} \int_{-\infty}^{-\frac{cu}{\sqrt{1-c^2}}} \int_{-\infty}^{\frac{\text{VaR}_q^X - m_X}{s_{X,1}}} \left(z + \frac{cu}{\sqrt{1-c^2}} \right) \exp\left(-\frac{u^2 + z^2}{2}\right) dudz \\ &= \frac{s_{Y,1}\sqrt{1-c^2}}{2\pi} \int_{-\infty}^{\frac{\text{VaR}_q^X - m_X}{s_{X,1}}} \exp\left(-\frac{u^2}{2}\right) \left(\int_{-\infty}^{-\frac{cu}{\sqrt{1-c^2}}} z \exp\left(-\frac{z^2}{2}\right) dz \right) du \\ & \quad + \frac{s_{Y,1}c}{2\pi} \int_{-\infty}^{\frac{\text{VaR}_q^X - m_X}{s_{X,1}}} u \exp\left(-\frac{u^2}{2}\right) \left(\int_{-\infty}^{-\frac{cu}{\sqrt{1-c^2}}} \exp\left(-\frac{z^2}{2}\right) dz \right) du. \end{aligned}$$

This sum can be rewritten to

$$-\frac{s_{Y,1}\sqrt{1-c^2}}{2\pi} \int_{-\infty}^{\frac{\text{VaR}_q^X - m_X}{s_{X,1}}} \exp\left(-\frac{u^2}{2(1-c^2)}\right) du + \frac{s_{Y,1}c}{\sqrt{2\pi}} \int_{-\infty}^{\frac{\text{VaR}_q^X - m_X}{s_{X,1}}} u \exp\left(-\frac{u^2}{2}\right) \Phi\left(-\frac{cu}{\sqrt{1-c^2}}\right) du.$$

We then compute the first part of this sum directly, and we integrate by parts to compute the second part. After rearranging the terms we finally get that the first integral of Equation (3.30) equals

$$\begin{aligned} & -\frac{s_{Y,1}}{\sqrt{2\pi}} \Phi\left(\frac{\text{VaR}_q^X - m_X}{s_{X,1}\sqrt{1-c^2}}\right) - \frac{cs_{Y,1}}{\sqrt{2\pi}} \Phi\left(-\frac{c(\text{VaR}_q^X - m_X)}{s_{X,1}\sqrt{1-c^2}}\right) \exp\left(-\frac{(\text{VaR}_q^X - m_X)^2}{2s_{X,1}^2}\right) \\ & \quad + m_Y L\left(-\infty, 0, -\infty, \frac{\text{VaR}_q^X - m_X}{s_{X,1}}\right). \end{aligned}$$

Using a similar reasoning, we obtain that the second integral of Equation (3.30) equals

$$\begin{aligned} & \frac{s_{Y,2}}{\sqrt{2\pi}} \Phi\left(\frac{\text{VaR}_q^X - m_X}{s_{X,1}\sqrt{1-c^2}}\right) - \frac{cs_{Y,2}}{\sqrt{2\pi}} \left(1 - \Phi\left(-\frac{c(\text{VaR}_q^X - m_X)}{s_{X,1}\sqrt{1-c^2}}\right)\right) \exp\left(-\frac{(\text{VaR}_q^X - m_X)^2}{2s_{X,1}^2}\right) \\ & \quad + m_Y L\left(0, +\infty, -\infty, \frac{\text{VaR}_q^X - m_X}{s_{X,1}}\right). \end{aligned}$$

Hence, we obtain

$$\begin{aligned} \text{CoES}_q(Y) = & \frac{1}{q\%} \left[m_Y \left(\lambda_{1,1} L \left(-\infty, 0, -\infty, \frac{\text{VaR}_q^X - m_X}{s_{X,1}} \right) + \lambda_{1,2} L \left(0, +\infty, -\infty, \frac{\text{VaR}_q^X - m_X}{s_{X,1}} \right) \right) \right. \\ & + \frac{\lambda_{1,2}s_{Y,2} - \lambda_{1,1}s_{Y,1}}{\sqrt{2\pi}} \Phi \left(\frac{\text{VaR}_q^X - m_X}{s_{X,1}\sqrt{1-c^2}} \right) - c\phi \left(\frac{\text{VaR}_q^X - m_X}{s_{X,1}} \right) \times \\ & \left. \left(\lambda_{1,1}s_{Y,1} \Phi \left(-\frac{c(\text{VaR}_q^X - m_X)}{s_{X,1}\sqrt{1-c^2}} \right) + \lambda_{1,2}s_{Y,2} \left(1 - \Phi \left(-\frac{c(\text{VaR}_q^X - m_X)}{s_{X,1}\sqrt{1-c^2}} \right) \right) \right) \right]. \end{aligned}$$

Proof of statement 2. If $\text{VaR}_q^X > m_X$, $\text{CoES}_q(Y)$ is

$$\begin{aligned} \text{CoES}_q(Y) = & \frac{1}{q\%} \left(\lambda_{1,1} \int_{-\infty}^{m_Y} \int_{-\infty}^{m_X} y f_{1,1}(x, y) \, dx dy + \lambda_{1,2} \int_{m_Y}^{+\infty} \int_{-\infty}^{m_X} y f_{1,2}(x, y) \, dx dy \right. \\ & \left. \lambda_{2,1} \int_{-\infty}^{m_Y} \int_{m_X}^{\text{VaR}_q^X} y f_{2,1}(x, y) \, dx dy + \lambda_{2,2} \int_{m_Y}^{+\infty} \int_{m_X}^{\text{VaR}_q^X} y f_{2,2}(x, y) \, dx dy \right). \end{aligned}$$

Using similar computations we obtain

$$\begin{aligned} \text{CoES}_q(Y) = & \frac{1}{q\%} \left[m_Y \left(\lambda_{1,1} \left(\frac{1}{4} + \frac{\sin^{-1}(c)}{2\pi} \right) + \lambda_{1,2} \frac{\cos^{-1}(c)}{2\pi} + \lambda_{2,1} L \left(-\infty, 0, 0, \frac{\text{VaR}_q^X - m_X}{s_{X,2}} \right) \right) \right. \\ & + \lambda_{2,2} L \left(0, +\infty, 0, \frac{\text{VaR}_q^X - m_X}{s_{X,2}} \right) \left. + \frac{(1-c)(\lambda_{1,2} - \lambda_{2,2})s_{Y,2} - (1+c)(\lambda_{1,1} - \lambda_{2,1})s_{Y,1}}{2\sqrt{2\pi}} \right. \\ & + \frac{\lambda_{2,2}s_{Y,2} - \lambda_{2,1}s_{Y,1}}{\sqrt{2\pi}} \Phi \left(\frac{\text{VaR}_q^X - m_X}{s_{X,2}\sqrt{1-c^2}} \right) - c\phi \left(\frac{\text{VaR}_q^X - m_X}{s_{X,2}} \right) \times \\ & \left. \left(\lambda_{2,1}s_{Y,1} \Phi \left(-\frac{c(\text{VaR}_q^X - m_X)}{s_{X,2}\sqrt{1-c^2}} \right) + \lambda_{2,2}s_{Y,2} \left(1 - \Phi \left(-\frac{c(\text{VaR}_q^X - m_X)}{s_{X,2}\sqrt{1-c^2}} \right) \right) \right) \right]. \end{aligned}$$

Introducing $M_{q,1}$, $M_{q,2}$, and L and slightly rearranging the terms, we find the expressions reported in Proposition 3.4. ■

3.C Additional results

3.C.1 Sharpe ratio in the portfolio choice setting

The Sharpe ratio of the portfolio of two risky assets is

$$SR(R_p) = \frac{\omega_X \mu_X + \omega_Y \mu_Y}{(\omega_X^2 \sigma_X^2 + \omega_Y^2 \sigma_Y^2 + 2\omega_X \omega_Y \sigma_X \sigma_Y \rho)^{0.5}}.$$

Using $\mu_X = \mu_Y \equiv \mu$ and $\sigma_X = \sigma_Y \equiv \sigma$, we obtain

$$\begin{aligned} SR(R_p) &= \frac{\mu}{\sigma} \frac{\omega_X + \omega_Y}{(\omega_X^2 + \omega_Y^2 + 2\omega_X \omega_Y \rho)^{0.5}} \\ &= \frac{\mu}{\sigma} \frac{\omega_X + \omega_Y}{((\omega_X + \omega_Y)^2 - 2(1 - \rho)\omega_X \omega_Y)^{0.5}} \\ &= \frac{\mu}{\sigma} \left(1 - 2(1 - \rho) \frac{\frac{\omega_X}{\omega_Y}}{\left(\frac{\omega_X}{\omega_Y} + 1\right)^2} \right)^{-0.5}. \end{aligned}$$

Using the last expression, it is easy to check that the Sharpe ratio is increasing in ω_X/ω_Y for $\omega_X/\omega_Y < 1$ and decreasing in ω_X/ω_Y for $\omega_X/\omega_Y > 1$. It is also continuous in ω_X/ω_Y . Thus, $\omega_X/\omega_Y = 1$ maximizes the Sharpe ratio and deviations from one measure the mean-variance inefficiency of the portfolio.

3.C.2 Sample variance of co-skewness

Below we calculate the sample variance of co-skewness for the case in which returns are drawn from a BN distribution. Let $m_{ij} = \mathbb{E} \left[(X - \mathbb{E}(X))^i (Y - \mathbb{E}(Y))^j \right]$ denote the sample centralized co-moment of order ij of a sample of size n , and let μ_{ij} denote the corresponding population centralized co-moment. Using the formula (10.17) in Stuart and Ord (1994), the

sample variance of co-skewness $\gamma_{Y,X}$ is

$$\begin{aligned} \text{var}(\gamma_{Y,X}) &= \frac{m_{21}}{m_{20}m_{02}^{1/2}} \\ &= \left(\frac{\mu_{21}^2}{\mu_{20}^2\mu_{02}} \right) \left(\frac{\text{var}(m_{21})}{\mu_{21}^2} + \frac{\text{var}(m_{20}m_{02}^{1/2})}{\mu_{20}^2\mu_{02}} - 2 \frac{\text{cov}(m_{21}, m_{20}m_{02}^{1/2})}{\mu_{21}\mu_{20}\mu_{02}^{1/2}} \right). \end{aligned}$$

Under the assumption that the returns are drawn from a BN distribution, we have $\mu_{21} = 0$. Hence, using (10.23) in Stuart and Ord (1994), we can rewrite the previous expression to


$$\begin{aligned} \text{var}(\gamma_{Y,X}) &= \frac{\text{var}(m_{21})}{\mu_{20}^2\mu_{02}} \\ &= \frac{1}{n} \frac{\mu_{42} - \mu_{21}^2 + 8\mu_{20}\mu_{11}^2 + \mu_{20}^2\mu_{02} - 4\mu_{31}\mu_{11} - 2\mu_{20}\mu_{22}}{\mu_{20}^2\mu_{02}}. \end{aligned}$$

For the normal case, we have $\mu_{20} = \sigma_X^2$, $\mu_{02} = \sigma_Y^2$, $\mu_{42} = 3\sigma_X^4\sigma_Y^2(1 + 4\rho^2)$, $\mu_{11} = \rho\sigma_X\sigma_Y$, $\mu_{31} = 3\rho\sigma_X^3\sigma_Y$, and $\mu_{22} = \sigma_X^2\sigma_Y^2(1 + 2\rho^2)$. Plugging these expressions in the expression for $\text{var}(\gamma_{Y,X})$ above, we get


$$\text{var}(\gamma_{Y,X}) = \frac{2 + 4\rho^2}{n}.$$

Figure 3.1: Four quadrants of BN distributions

Panel A: Unscaled

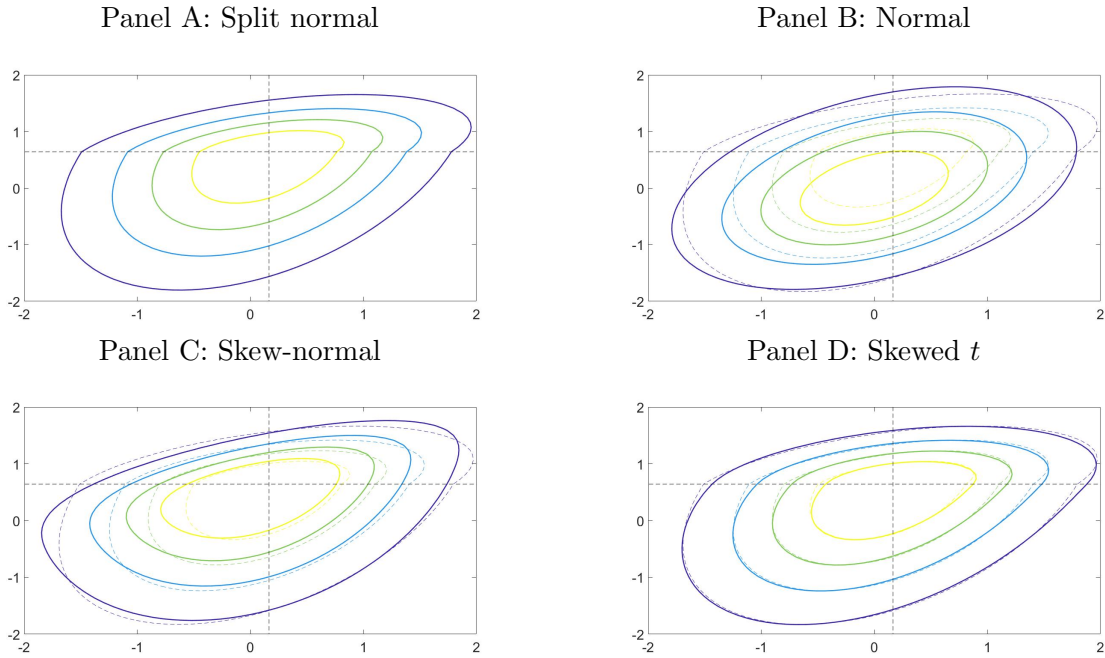


Panel B: Scaled



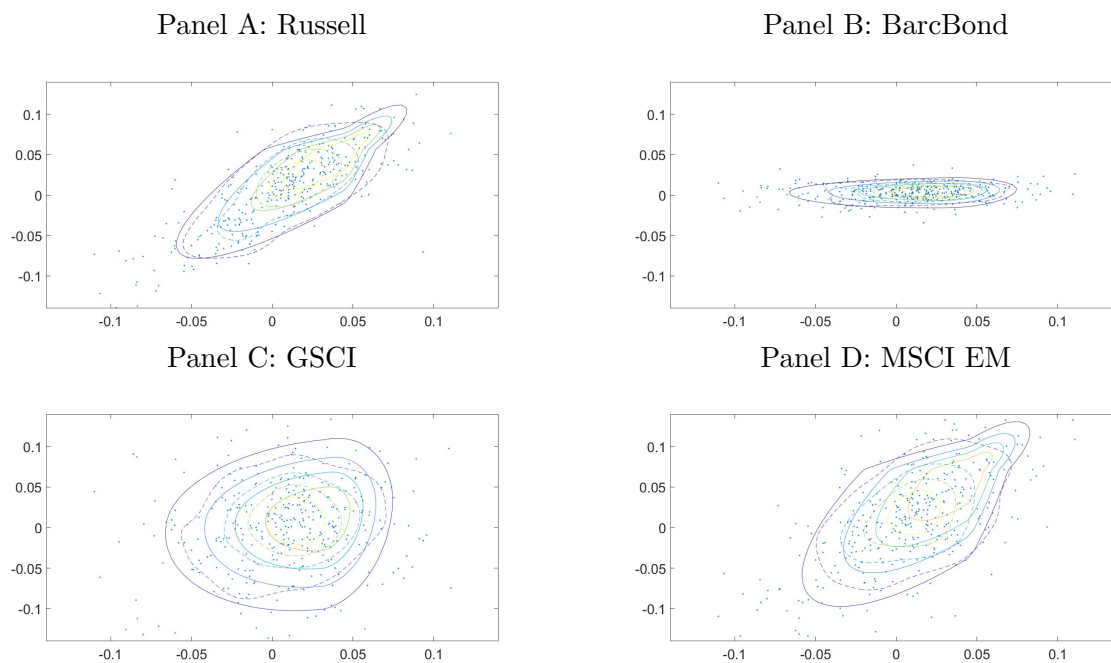
The blue (red, yellow, white) quadrant is the first (second, third, fourth) quadrant of a BN distribution in which X and Y have zero-means, no correlation, and variances of 1.2 and 0.6 (0.6 and 0.6, 0.6 and 1.2, 1.2 and 1.2). The quadrants are unscaled in Panel A and scaled to obtain a continuous skewed pdf in Panel B.

Figure 3.2: Comparison of bivariate distributions



The figure presents contour plots of the SBN pdf (Panel A), the BN pdf (Panel B), the BSN pdf (Panel C), and the BST pdf constructed with a Gaussian copula (Panel D). The distributions are parameterized such that the two random variables X and Y have an expected value of zero, unit variance, a correlation of 0.4, and skewnesses are zero for X and equal to -0.6 for Y (for the asymmetric distributions). The distribution in Panel D has two additional parameters that are chosen to match the kurtosises of X and Y in Panel A. All plots represent X on the horizontal axis and Y on the vertical axis. Straight dashed lines highlight the mode values of the SBN distribution, and the contour plot of the SBN distribution is included in Panels B-D to facilitate the comparison.

Figure 3.3: The bivariate distribution of returns on the SP500 and ...



Each contour plot presents the bivariate pdf of the monthly excess returns on the S&P 500 TR Index on the x -axis and the corresponding return of one of the four other indices on the y -axis (the Russell 2000 TR Index (Panel A), the Barclays US Aggregate Bond TR Index (Panel B), the GSCI TR (Panel C), and the MSCI Emerging Market TR Index (Panel D)). The sample contains returns from January 1989 to June 2018 and each observation is represented in the x - y plane with a dot. Each panel contains the empirical pdf estimated with a kernel density estimator (dashed lines) and the pdf of the SBN distribution with the sample means, standard deviations, correlation, and skewnesses (solid lines).

Table 3.1: Descriptive statistics

	SP500	Russell	BarcBond	GSCI	MSCI EM
Panel A: moments					
Mean	0.67%	0.71%	0.27%	0.15%	0.78%
St.dev.	4.07%	5.31%	1.04%	5.95%	6.58%
Skewness	-0.586	-0.519	-0.239	-0.166	-0.586
	0.00	0.00	0.07	0.20	0.00
Kurtosis	4.309	4.146	3.630	4.879	4.795
	0.00	0.00	0.02	0.00	0.00
Panel B: correlations					
Russell	0.806	-	-	-	-
BarcBond	0.106	-0.006	-	-	-
GSCI	0.175	0.205	-0.051	-	-
MSCI EM	0.670	0.655	0.018	0.277	-

The table contains the means, standard deviations, skewnesses, and kurtosises of the monthly index returns in excess of the 30-day t-bill return in Panel A and their correlations in Panel B. The indices are the S&P 500 TR Index (SP500), the Russell 2000 TR Index (Russell), the Barclays US Aggregate Bond TR Index (BarcBond), the GSCI TR (GSCI), and the MSCI Emerging Market TR Index (MSCI EM). The p -values for the null hypothesis that skewness (resp. kurtosis) is zero (resp. three) are reported below the skewness (resp. kurtosis) under the null of normally distributed returns. All statistics are monthly and for returns from January 1989 to June 2018.

Table 3.2: Kolmogorov-Smirnov tests

Index pair	SBN	BN	BSN	BST
SP500, Russell	0.06	0.01	0.15	0.36
SP500, BarcBond	0.24	0.05	0.03	0.81
SP500, GSCI	0.08	0.03	0.04	0.63
SP500, MSCI EM	0.02	0.01	0.03	0.28
Russell, BarcBond	0.26	0.06	NA	0.53
Russell, GSCI	0.07	0.00	0.17	0.32
Russell, MSCI EM	0.04	0.04	0.19	0.66
BarcBond, GSCI	0.49	0.24	0.26	0.99
BarcBond, MSCI EM	0.13	0.12	NA	0.99
GSCI, MSCI EM	0.01	0.03	0.04	0.28

The table contains the p -values of the bivariate Kolmogorov-Smirnov goodness-of-fit test proposed by Peacock (1983) for the null hypothesis that the excess returns are drawn from the distribution indicated in the column header for all the possible pairs of the five indices. NA indicates the cases in which the parameters of the BSN distribution are outside of their admissible range.

Table 3.3: Tests on co-skewness

Index pair	Emp. co-skewness		Test $\gamma_{Y,X}$				Test $\gamma_{X,Y}$			
	$\gamma_{Y,X}$	$\gamma_{X,Y}$	SBN	BN	BSN	BST	SBN	BN	BSN	BST
SP500, Russell	-0.634	-0.624	0.12	0.00	0.52	0.11	0.11	0.00	0.46	0.15
SP500, BarcBond	-0.015	-0.200	0.82	0.84	0.00	0.73	0.02	0.01	0.20	0.04
SP500, GSCI	-0.322	-0.446	0.00	0.00	0.52	0.00	0.00	0.00	0.03	0.00
SP500, MSCI EM	-0.589	-0.570	0.05	0.00	0.91	0.05	0.07	0.00	0.96	0.14
Russell, BarcBond	-0.082	-0.110	0.29	0.28	NA	0.35	0.14	0.14	NA	0.18
Russell, GSCI	-0.261	-0.490	0.02	0.00	0.33	0.06	0.00	0.00	0.01	0.00
Russell, MSCI EM	-0.536	-0.549	0.04	0.00	0.97	0.12	0.05	0.00	0.96	0.16
BarcBond, GSCI	-0.173	-0.084	0.02	0.02	0.67	0.04	0.24	0.26	0.23	0.36
BarcBond, MSCI EM	-0.134	-0.095	0.09	0.07	NA	0.13	0.26	0.21	NA	0.36
GSCI, MSCI EM	-0.381	-0.274	0.00	0.00	0.15	0.01	0.03	0.00	0.26	0.13

The table contains empirical co-skewnesses of the excess returns of ten pairs of indices from different asset classes (large stocks, small stocks, bonds, commodities, and emerging markets stocks) and the p -values of the tests for both $\gamma_{Y,X}$ and $\gamma_{X,Y}$ that the excess returns are drawn from the distribution indicated in the column header. For the BN distribution, we compute the p -values using the variance of co-skewness under the null derived in Appendix 3.C.2. For the other distributions, we compute the p -values by generating 100,000 times the statistics (using samples of 354 observations) under the null. NA indicates the cases in which the parameters of the BSN distribution are outside of their admissible range.

Table 3.4: Portfolio choice

Correlation	-0.2	0	0.2	0.4	0.6	0.8	-0.2	0	0.2	0.4	0.6	0.8	-0.2	0	0.2	0.4	0.6	0.8
$\gamma_X = \gamma_Y \equiv \gamma$	Benchmark case: $\gamma = 0$						Panel A: $\gamma = 0.6$						Panel B: $\gamma = -0.6$					
$\omega_{X,SBN}, \omega_{Y,SBN}$	1.33	1.07	0.89	0.76	0.67	0.59	1.86	1.44	1.17	1.00	0.88	0.78	1.10	0.88	0.74	0.63	0.54	0.48
$CER(\omega_{SBN})$ in %	16.00	12.80	10.67	9.14	8.00	7.11	19.76	15.54	12.76	10.88	9.58	8.53	13.93	11.20	9.37	8.01	6.94	6.14
$CER(\omega_{BN})$ in %	16.00	12.80	10.67	9.14	8.00	7.11	18.43	14.63	12.11	10.36	9.10	8.10	13.19	10.64	8.92	7.60	6.51	5.73
$\gamma_{\omega_{SBN}}, \gamma_{\omega_{BN}}$	0.00	0.00	0.00	0.00	0.00	0.00	0.40	0.42	0.44	0.48	0.54	0.60	-0.40	-0.42	-0.44	-0.48	-0.54	-0.60
$\gamma_X \neq \gamma_Y$	Panel C: $\gamma_X = 0.6$ and $\gamma_Y = -0.6$						Panel D: $\gamma_X = 0.6$ and $\gamma_Y = 0$						Panel E: $\gamma_X = 0$ and $\gamma_Y = -0.6$					
$\omega_{X,SBN}$	1.81	1.44	1.23	1.11	1.08	1.27	1.82	1.44	1.21	1.07	1.02	1.14	1.32	1.07	0.91	0.80	0.73	0.72
$\omega_{Y,SBN}$	1.14	0.88	0.70	0.54	0.39	0.12	1.37	1.07	0.86	0.69	0.53	0.29	1.11	0.88	0.72	0.60	0.50	0.36
$\omega_{X,SBN}/\omega_{Y,SBN}$	1.58	1.64	1.77	2.04	2.74	10.49	1.33	1.35	1.42	1.55	1.91	3.93	1.19	1.21	1.25	1.32	1.47	2.01
$CER(\omega_{SBN})$ in %	16.84	13.37	11.10	9.52	8.40	7.71	17.86	14.17	11.72	10.00	8.76	7.94	14.98	12.00	10.03	8.62	7.54	6.67
$CER(\omega_{BN})$ in %	15.68	12.64	10.55	9.06	7.94	7.08	17.17	13.72	11.39	9.72	8.49	7.57	14.57	11.72	9.82	8.45	7.38	6.48
$\gamma_{\omega_{SBN}}$	0.32	0.29	0.28	0.30	0.35	0.52	0.32	0.31	0.31	0.33	0.37	0.48	-0.14	-0.16	-0.16	-0.16	-0.16	-0.16
$\gamma_{\omega_{BN}}$	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.21	0.22	0.22	0.23	0.27	-0.20	-0.21	-0.22	-0.22	-0.23	-0.27

The table analyzes the joint effects of skewness and correlation on portfolio choice. The table reports the optimal weights, $\omega_{X,SBN}$ and $\omega_{Y,SBN}$, the certainty equivalent return of the optimal portfolio, $CER(\omega_{SBN})$, the certainty equivalent return of the portfolio calculated with normally distributed returns evaluated with skewed returns, $CER(\omega_{BN})$, the optimal portfolio skewness, $\gamma_{\omega_{SBN}}$, and the portfolio skewness of the portfolio calculated with normally distributed returns, $\gamma_{\omega_{BN}}$. Excess returns on the risky assets X and Y follow the SBN distribution that matches means of 12%, standard deviations of 15%, a given correlation, and given skewnesses. This correlation varies through the columns from -0.2 to 0.8 , and the skewnesses vary through panels. The top part of the table, i.e., the BN benchmark case and Panel A-B, considers two assets with the same levels of skewness. The bottom part of the table, i.e., Panel C-E, analyzes portfolio choice with assets that have different levels of skewness. In this case, optimal weights in X and Y differ and we also report the ratio of the two weights, $\omega_{X,SBN}/\omega_{Y,SBN}$, as a measure of deviation from mean-variance efficiency. The investor has EU preferences with coefficient of absolute risk aversion 5.

Table 3.5: Portfolio choice with CPT

Correlation	-0.2	0	0.2	0.4	0.6	0.8	-0.2	0	0.2	0.4	0.6	0.8	-0.2	0	0.2	0.4	0.6	0.8	
$\gamma_X = \gamma_Y \equiv \gamma$	Benchmark case: $\gamma = 0$																		
$\omega_{X,SBN}, \omega_{Y,SBN}$	14.04	7.78	3.68	0.76	0.00	0.00	13.30	8.77	5.95	4.18	3.08	2.27	17.62	5.14	0.00	0.00	0.00	0.00	0.00
$CER(\omega_{SBN})$ in %	21.01	8.66	2.47	0.13	0.00	0.00	31.52	18.81	11.24	6.88	4.60	2.96	13.41	0.73	0.00	0.00	0.00	0.00	0.00
$CER(\omega_{BN})$ in %	21.01	8.66	2.47	0.13	0.00	0.00	31.47	18.66	9.94	2.53	0.00	0.00	13.07	0.55	-1.82	-0.95	0.00	0.00	
$\gamma_{\omega_{SBN}}, \gamma_{\omega_{BN}}$	0.00	0.00	0.00	0.00	NaN	NaN	0.40	0.42	0.44	0.48	0.54	0.60	-0.40	-0.42	NaN	NaN	NaN	NaN	
$\gamma_X \neq \gamma_Y$	Panel C: $\gamma_X = 0.6$ and $\gamma_Y = -0.6$																		
$\omega_{X,SBN}$	15.00	9.34	6.08	4.19	3.20	3.98	14.60	9.49	6.39	4.48	3.49	3.90	15.77	7.85	2.06	0.00	0.00	0.00	
$\omega_{Y,SBN}$	10.35	5.15	2.44	1.03	0.22	-0.71	11.51	6.54	3.54	1.71	0.57	-0.51	13.77	6.09	1.32	0.00	0.00	0.00	
$\omega_{X,SBN}/\omega_{Y,SBN}$	1.45	1.81	2.49	4.06	14.70	-5.62	1.27	1.45	1.80	2.62	6.18	-7.61	1.15	1.29	1.56	NaN	NaN	NaN	
$CER(\omega_{SBN})$ in %	24.20	12.38	6.14	3.13	1.77	2.15	26.96	14.69	7.70	3.99	2.31	2.21	17.60	5.25	0.41	0.00	0.00	0.00	
$CER(\omega_{BN})$ in %	21.20	8.63	2.28	0.08	0.00	0.00	25.79	13.29	5.83	1.18	0.00	0.00	17.04	4.47	-0.35	-0.41	0.00	0.00	
$\gamma_{\omega_{SBN}}$	0.27	0.34	0.40	0.47	0.56	0.75	0.30	0.33	0.38	0.44	0.52	0.67	-0.15	-0.14	-0.11	NaN	NaN	NaN	
$\gamma_{\omega_{BN}}$	0.00	0.00	0.00	0.00	NaN	NaN	0.20	0.21	0.22	0.22	NaN	NaN	-0.20	-0.21	-0.22	NaN	NaN	NaN	

The table is constructed in the same fashion as Table 3.4. Portfolio weights maximize CPT preferences implemented with the piecewise negative exponential specification of Giorgi and Hens (2006) and a parameterization that corresponds to the experiments of Tversky and Kahneman (1992).

The value function is

$$v(x) = \begin{cases} -6.52e^{-0.2x} + 6.52, & \text{if } x \geq 0, \\ 14.7e^{0.2x} - 14.7, & \text{if } x < 0, \end{cases}$$

and the Tversky and Kahneman (1992) weighting function is $T(p) = p^{0.65} / (p^{0.65} + (1-p)^{0.65})^{1/0.65}$. The reference point is the wealth generated by the risk-free return. Therefore, the CPT objective is $\int v(\omega_{X,SBN}R_X + \omega_{Y,SBN}R_Y) d(T \circ CDF)$, where CDF is the cumulative distribution function of the portfolio return. To solve the portfolio choice problem, we draw 10^7 return pairs from the SBN distribution and then maximize the sample equivalent of the CPT objective. The returns are drawn from a distribution with the same parameters as in Table 3.4, except that the expected excess returns are equal to 6.13% to obtain the same portfolio weights as in Table 3.4 when the return distribution is symmetric and $\rho = 0.4$.

Table 3.6: Asset pricing

Correlation	Panel A: same volatility, same supply					Panel B: different volatility, same supply					Panel C: different volatility, different supply							
	-0.2	0	0.2	0.4	0.6	0.8	-0.2	0	0.2	0.4	0.6	0.8	-0.2	0	0.2	0.4	0.6	0.8
Asset X																		
α^{CPT} in %	-2.90	-2.65	-2.44	-2.27	-2.09	-1.91	-1.88	-1.90	-1.88	-1.83	-1.75	-1.65	-1.13	-1.87	-2.45	-2.91	-3.29	-3.57
α^{EU} in %	-0.33	-0.38	-0.42	-0.45	-0.48	-0.49	-0.20	-0.25	-0.29	-0.33	-0.36	-0.38	-0.16	-0.31	-0.45	-0.59	-0.74	-0.87
β_X	1.00	1.00	1.00	1.00	1.00	1.00	1.71	1.60	1.52	1.45	1.41	1.37	0.05	0.47	0.82	1.13	1.39	1.62
$\gamma_{X,M}$	0.33	0.30	0.28	0.26	0.24	0.22	0.52	0.48	0.45	0.42	0.39	0.36	0.07	0.03	-0.01	-0.03	-0.05	-0.05
max_X in %	47.2	47.3	47.4	47.6	47.8	47.5	62.8	63.1	63.2	63.4	63.7	63.3	62.9	63.0	63.2	63.4	63.7	63.3
min_X in %	-30.2	-30.5	-30.8	-31.2	-31.6	-32.2	-40.4	-40.7	-41.0	-41.5	-42.1	-42.9	-40.3	-40.8	-41.1	-41.5	-42.1	-42.9
σ_{ϵ_X} in %	11.6	10.6	9.5	8.2	6.7	4.7	9.6	8.9	8.1	7.1	5.9	4.2	20.0	19.5	18.4	16.5	13.9	10.1
γ_{ϵ_X}	0.44	0.42	0.40	0.36	0.29	0.26	0.52	0.48	0.45	0.40	0.32	0.28	0.60	0.56	0.53	0.47	0.39	0.33
$\gamma_{\epsilon_X,M}$	0.42	0.42	0.44	0.47	0.53	0.69	0.28	0.32	0.36	0.41	0.49	0.67	0.09	0.15	0.22	0.29	0.39	0.60
Asset Y																		
α^{CPT} in %	2.90	2.65	2.44	2.27	2.09	1.91	1.88	1.90	1.88	1.83	1.75	1.65	0.13	0.21	0.27	0.32	0.37	0.40
α^{EU} in %	0.33	0.38	0.42	0.45	0.48	0.49	0.20	0.25	0.29	0.33	0.36	0.38	0.02	0.03	0.05	0.07	0.08	0.10
β_Y	1.00	1.00	1.00	1.00	1.00	1.00	0.29	0.40	0.48	0.55	0.59	0.63	1.11	1.06	1.02	0.99	0.96	0.93
$\gamma_{Y,M}$	-0.33	-0.30	-0.28	-0.26	-0.24	-0.22	-0.14	-0.12	-0.10	-0.09	-0.08	-0.07	-0.59	-0.57	-0.55	-0.53	-0.51	-0.48
max_Y in %	30.2	30.5	30.8	31.2	31.6	32.2	20.2	20.3	20.5	20.8	21.1	21.5	20.2	20.4	20.5	20.8	21.1	21.5
min_Y in %	-47.2	-47.3	-47.4	-47.6	-47.8	-47.5	-31.4	-31.5	-31.6	-31.7	-31.8	-31.7	-31.4	-31.5	-31.6	-31.7	-31.8	-31.7
σ_{ϵ_Y} in %	11.6	10.6	9.5	8.2	6.7	4.7	9.6	8.9	8.1	7.1	5.9	4.2	2.2	2.2	2.0	1.8	1.5	1.1
Market M																		
μ^{CPT} in %	4.61	5.15	5.65	6.12	6.53	6.93	4.08	4.60	5.09	5.54	5.96	6.35	5.29	5.45	5.61	5.76	5.89	6.01
μ^{EU} in %	3.49	4.36	5.24	6.12	6.99	7.85	3.65	4.38	5.11	5.84	6.58	7.32	4.27	4.66	5.04	5.42	5.80	6.15
μ_M	0.00	0.00	0.00	0.00	0.00	0.00	0.44	0.38	0.33	0.29	0.26	0.23	-0.59	-0.55	-0.52	-0.48	-0.45	-0.42
Taylor expansion																		
$\alpha^{EU,Taylor}$ in %	-0.31	-0.35	-0.38	-0.41	-0.42	-0.43	-0.19	-0.24	-0.28	-0.31	-0.34	-0.36	-0.15	-0.27	-0.39	-0.50	-0.61	-0.72
$\mu^{EU,Taylor}$ in %	3.26	4.01	4.73	5.44	6.13	6.79	3.38	3.99	4.60	5.19	5.77	6.34	3.89	4.21	4.52	4.82	5.11	5.39

The table analyzes the asset pricing implications of skewness preference. We consider two economies with two risky assets X and Y . The first economy is populated with a representative CPT investor, CPT , and the second economy is populated with a representative EU investor, EU . The table reports for both assets the CAPM alphas, α , and betas, β , and the equilibrium returns on the market μ_M . Subscripts specify the asset and superscripts the economy. The table also reports for each asset the (standardized) co-skewness with the market portfolio $\gamma_{\cdot,M}$, the volatility of the CAPM residual $\sigma_{\epsilon_{\cdot}}$, the return expectation in the 1% left and right tails (min and max). For asset X , it also reports its (standardized) residual co-skewness (i.e., the co-skewness of the CAPM residual with the market portfolio), $\gamma_{\epsilon_X,M}$, and the skewness of its CAPM residual γ_{ϵ_X} (we have $\gamma_{\epsilon_X,M} = -\gamma_{\epsilon_X,M}$ and $\gamma_{\epsilon_Y} = -\gamma_{\epsilon_X}$). The last two rows of the table report the alpha of asset X and the excess market return that are obtained when the market equilibrium is computed with a second-order Taylor expansion to marginal utility. The skewness of asset X is 0.6 and the skewness of asset Y is -0.6 . In Panel A both assets have a standard deviation of 15%, whereas in Panel B-C asset X has a standard deviation of 20% and asset Y of 10%. The supply of both assets (i.e., the market portfolio weights) is 0.5 in Panel A-B, and asset X and Y have respective supplies of 0.1 and 0.9 in Panel C. The preference coefficients for CPT are the same as in Table 3.5. The coefficient of risk aversion for the EU investor is 3.86 in Panel A, 3.78 in Panel B, and 4.85 in Panel C. These values are chosen such that risk premia in the two economies are the same when returns are normally distributed and $\rho = 0.4$.

Table 3.7: ΔCoVaR

Correlation	-0.2	0	0.2	0.4	0.6	0.8
Benchmark case: $\gamma_X = 0$ and $\gamma_Y = 0$						
$\Delta\text{CoVaR}_5^{Y X}$	0.329	0	-0.329	-0.658	-0.987	-1.316
Panel A: $\gamma_X = 0.6$ and $\gamma_Y = 0.6$						
$\Delta\text{CoVaR}_5^{Y X}$	0.279	0	-0.267	-0.490	-0.629	-0.729
Panel B: $\gamma_X = -0.6$ and $\gamma_Y = -0.6$						
$\Delta\text{CoVaR}_5^{Y X}$	0.381	0	-0.408	-0.848	-1.293	-1.670
Panel C: $\gamma_X = 0.6$ and $\gamma_Y = -0.6$						
$\Delta\text{CoVaR}_5^{Y X}$	0.471	0	-0.483	-0.956	-1.394	-1.786
Panel D: $\gamma_X = 0.6$ and $\gamma_Y = 0$						
$\Delta\text{CoVaR}_5^{Y X}$	0.381	0	-0.380	-0.754	-1.104	-1.393
Panel E: $\gamma_X = -0.6$ and $\gamma_Y = 0$						
$\Delta\text{CoVaR}_5^{Y X}$	0.312	0	-0.312	-0.630	-0.968	-1.344
Panel C': $\gamma_X = -0.6$ and $\gamma_Y = 0.6$						
$\Delta\text{CoVaR}_5^{Y X}$	0.222	0	-0.225	-0.449	-0.687	-0.960
Panel D': $\gamma_X = 0$ and $\gamma_Y = 0.6$						
$\Delta\text{CoVaR}_5^{Y X}$	0.235	0	-0.239	-0.473	-0.694	-0.868
Panel E': $\gamma_X = 0$ and $\gamma_Y = -0.6$						
$\Delta\text{CoVaR}_5^{Y X}$	0.411	0	-0.417	-0.833	-1.238	-1.624

The table contains the ΔCoVaR for the assets X and Y at the 5% probability. The SBN distribution is calculated to match the first three moments of the distribution of X and Y and a correlation that ranges from -0.2 to 0.8. In all panels, X and Y have zero means and unit variances. The skewnesses of X and Y vary through panels.

Table 3.8: CoES

Correlation	-0.2	0	0.2	0.4	0.6	0.8
Benchmark case: $\gamma_X = 0$ and $\gamma_Y = 0$						
$\text{CoES}_5^{Y X}$	0.413	0	-0.413	-0.825	-1.238	-1.650
Panel A: $\gamma_X = 0.6$ and $\gamma_Y = 0.6$						
$\text{CoES}_5^{Y X}$	0.472	0	-0.468	-0.916	-1.294	-1.579
Panel B: $\gamma_X = -0.6$ and $\gamma_Y = -0.6$						
$\text{CoES}_5^{Y X}$	0.403	0	-0.418	-0.882	-1.424	-1.962
Panel C: $\gamma_X = 0.6$ and $\gamma_Y = -0.6$						
$\text{CoES}_5^{Y X}$	0.468	0	-0.472	-0.948	-1.437	-1.955
Panel D: $\gamma_X = 0.6$ and $\gamma_Y = 0$						
$\text{CoES}_5^{Y X}$	0.470	0	-0.470	-0.941	-1.408	-1.846
Panel E: $\gamma_X = -0.6$ and $\gamma_Y = 0$						
$\text{CoES}_5^{Y X}$	0.408	0	-0.408	-0.817	-1.236	-1.698
Panel C': $\gamma_X = -0.6$ and $\gamma_Y = 0.6$						
$\text{CoES}_5^{Y X}$	0.418	0	-0.403	-0.796	-1.162	-1.481
Panel D': $\gamma_X = 0$ and $\gamma_Y = 0.6$						
$\text{CoES}_5^{Y X}$	0.417	0	-0.411	-0.816	-1.201	-1.509
Panel E': $\gamma_X = 0$ and $\gamma_Y = -0.6$						
$\text{CoES}_5^{Y X}$	0.411	0	-0.416	-0.851	-1.330	-1.883

The table contains the CoES for the assets X and Y at the 5% probability and is constructed in the same way as Table 3.7.

Table 3.9: Realized vs. theoretical CoES

	Russell	BarcBond	GSCI	MSCI EM	SumSq
Panel A: $\text{CoES}_5^{Y X}$					
Empirical	-2.207	-0.112	-0.762	-1.873	NA
SBN	-1.932	-0.216	-0.356	-1.606	0.323
BN	-1.661	-0.218	-0.359	-1.380	0.715
BSN	-2.045	-0.769	-0.808	-1.864	0.460
BST	-1.863	-0.221	-0.365	-1.534	0.403
Panel B: $\text{CoES}_{10}^{Y X}$					
Empirical	-1.699	-0.056	-0.348	-1.509	NA
SBN	-1.586	-0.182	-0.301	-1.310	0.071
BN	-1.413	-0.185	-0.306	-1.174	0.213
BSN	-1.679	-0.582	-0.628	-1.514	0.356
BST	-1.531	-0.186	-0.310	-1.262	0.108
Panel C: $\text{CoES}_{25}^{Y X}$					
Empirical	-1.081	-0.181	-0.165	-0.866	NA
SBN	-1.067	-0.131	-0.216	-0.877	0.005
BN	-1.023	-0.134	-0.222	-0.850	0.009
BSN	-1.127	-0.300	-0.355	-0.986	0.067
BST	-1.053	-0.136	-0.224	-0.871	0.006

The table contains the standardized empirical CoES at the 5%, 10%, and 25% probability and the corresponding theoretical expected shortfall that assumes index excess returns follow the indicated distribution. Following the definition in Section 3.4.2, the S&P 500 TR Index is the conditioning variable X and the other indices are the variable of interest Y . The last column contains the sum of squared differences between the empirical and theoretical CoES for the four index pairs for each distribution and each probability.

Substantial summary in French

Les modèles d'évaluation d'actifs ont pour rôle d'attribuer une valeur aux différents actifs financiers et de décrire leurs caractéristiques sur les marchés; cela permet alors aux agents d'élaborer leurs stratégies d'investissement. Afin de s'assurer de la pertinence de leurs conclusions, il est donc primordial que les hypothèses sous-jacentes à ces modèles soient convenablement définies : elles doivent être à la fois suffisamment réalistes, pour que les résultats théoriques en découlant reproduisent fidèlement les observations empiriques, et suffisamment simplificatrices, pour que l'on puisse procéder à une analyse théorique précise et, notamment, interpréter l'origine des résultats obtenus.

Dans cette thèse, constituée de trois chapitres indépendants, j'étudie la remise en question d'hypothèses simplificatrices qui se retrouvent dans un grand nombre de modèles d'évaluation d'actifs. Plus précisément, je cherche à analyser les impacts du relâchement de l'une ou l'autre des hypothèses suivantes : celle d'homogénéité des agents financiers et celle de normalité des rendements des actifs. Les deux premiers chapitres de cette thèse s'intéressent au relâchement de la première hypothèse et le troisième chapitre s'intéresse au relâchement de la seconde. J'étudie donc des modèles d'évaluation d'actifs plus réalistes, avec des agents hétérogènes ou avec des actifs dont les rendements ne suivent pas une loi normale.

Agents hétérogènes

L'hypothèse d'homogénéité des agents implique que les différents acteurs des marchés financiers sont identiques. En particulier, cette hypothèse présume que les agents (également

appelés de manière équivalente « investisseurs » dans le reste de la thèse) sont rationnels et ont les mêmes croyances concernant l'avenir. Ici, « rationnels » veut dire que les agents ne sont ni optimistes, ni pessimistes, et que leurs croyances sont les bonnes : s'il existe un processus économique de richesse ayant une tendance donnée, les investisseurs s'accorderont tous sur la vraie valeur de cette tendance et aucun ne la surestimera (comme un optimiste le ferait), ni ne la sous-estimera (ce qui serait le cas d'un pessimiste). L'avantage d'une telle hypothèse est qu'il est alors aisé d'agrèger tous les agents financiers en un unique agent représentatif, facilitant ainsi la résolution des modèles sous-jacents. Différents arguments ont longtemps été avancés afin de justifier une telle hypothèse simplificatrice.

Premièrement, dans l'esprit des travaux de Friedman (1953), il a été soutenu que, même s'il était possible que certains investisseurs soient optimistes ou pessimistes, ils ne devaient pas être pris en compte, car leur irrationalité les menait à faire de mauvais choix et, à plus long terme, menait à leur disparition (voir, par exemple, Sandroni, 2000). Toutefois, comme le soulignent Kogan et al. (2006), la survie des agents et leur impact sur le marché sont deux concepts différents qui doivent être étudiés séparément. Yan (2008) montre par exemple que l'élimination de ces agents irrationnels prend beaucoup de temps (au moins plusieurs dizaines d'années, voire plusieurs centaines) et qu'il faut donc tenir compte de leurs impacts sur les marchés. Une étude récente de Bottazzi et al. (2018) indique également que les agents hétérogènes ne sont pas nécessairement éliminés à long terme et que la non-optimalité du portefeuille d'un agent peut corriger l'inexactitude de sa croyance, ce qui entraîne sa survie.

Un second argument important en faveur de l'hypothèse d'homogénéité des croyances soutient que, comme il n'y a aucune raison d'avoir un biais optimiste ou pessimiste dans l'économie, les agents devraient être rationnels en moyenne et les effets induits par les investisseurs biaisés devraient donc s'annuler. Jouini and Napp (2011) montrent dans un modèle théorique que ce n'est pas le cas et qu'une situation où des agents hétérogènes sont rationnels en moyenne n'est pas équivalente à une situation où tous les agents sont rationnels.

Enfin, de façon similaire au concept des croyances pragmatiques de Hvide (2002), il

existe un argument en faveur de l'hypothèse d'homogénéité des croyances qui stipule que les agents irrationnels (c'est-à-dire optimistes ou pessimistes), observant que les agents rationnels prennent de meilleures décisions, devraient adopter la même croyance afin de connaître la même réussite. Remettant en cause ce troisième argument, Jouini and Napp (2016) montrent que, sous certaines conditions, des agents irrationnels peuvent faire mieux que des agents rationnels et qu'ils n'ont donc pas de véritables raisons de les imiter.

Étant donné que chacun des arguments avancés pour justifier cette hypothèse a été remis en cause, et étant observée la grande diversité d'actions des différents acteurs des marchés financiers, cette hypothèse d'homogénéité et de rationalité des croyances apparaît peu crédible. Il est ainsi pertinent d'étudier l'hétérogénéité des croyances des investisseurs et les potentiels impacts qu'elle peut avoir sur les marchés financiers ; cela fait l'objet des deux premiers chapitres de ma thèse.

Premier chapitre

Dans le premier chapitre, intitulé *Live fast, die young: equilibrium and survival in large economies*, paru dans la revue *Economic Theory* en avril 2021 et coécrit avec Elyès Jouini, nous étudions cette question d'un point de vue théorique. Nous considérons une économie peuplée par une infinité d'agents qui, en plus d'avoir des croyances différentes sur l'avenir, ont également des taux de préférences temporelles différents. Concrètement cela veut dire que, pour un niveau d'optimisme ou de pessimisme donné, les investisseurs seront également plus ou moins patients ou impatients, dans le sens où ils auront un taux d'actualisation temporelle plus ou moins important.

Bien que moins largement abordé, ce second type d'hétérogénéité a également été examiné dans la littérature. À la suite des travaux fondateurs de Samuelson (1937), il a été admis qu'un taux d'actualisation unique pouvait être utilisé pour condenser les choix intertemporels. Cependant, des études empiriques (voir, entre autres, Frederick et al., 2002) montrent que cette hypothèse n'est pas réaliste et qu'il existe en fait une grande variété de

taux de préférences temporelles parmi les investisseurs. En étudiant plusieurs pays, Wang et al. (2016) mettent par exemple en évidence cette hétérogénéité et montrent que des facteurs économiques et culturels peuvent l'expliquer.

De plus, nous envisageons dans notre étude la possibilité que ces deux types d'hétérogénéité soient corrélés. Plus concrètement, en supposant une corrélation négative, cela veut par exemple dire que nous considérons le cas où les agents financiers les plus optimistes sont également les plus patients—ce que des études empiriques semblent effectivement démontrer. À notre connaissance, nous sommes les premiers à proposer une étude des implications d'une corrélation entre ces deux types d'hétérogénéité, et cela constitue donc l'une des contributions principales de notre papier.

Une autre hypothèse importante sur laquelle repose le modèle développé dans l'article est l'hypothèse de la présence d'un continuum d'investisseurs sur le marché. En fait, l'intérêt de considérer un nombre infini d'agents est triple. Tout d'abord, une telle hypothèse permet de tenir compte de l'intégralité des croyances potentielles, ainsi que de tous les degrés de patience possibles. De plus, cela implique qu'il existera toujours, et quelles que soient les circonstances, une hétérogénéité parmi les investisseurs. En effet, considérer un nombre fini d'agents implique nécessairement de définir l'agent « le plus optimiste de tous ». Dans les états du monde les plus positifs, cet agent concentrerait toute la richesse de l'économie et il n'y aurait alors plus d'hétérogénéité sur le marché (puisque'il n'y aurait plus qu'un investisseur). Enfin, d'un point de vue technique, cette hypothèse nous autorise à décrire les différentes croyances et les différentes préférences temporelles des agents à l'aide de lois statistiques. Cela permet à la fois de limiter le nombre de paramètres à utiliser (nous avons seulement besoin d'identifier deux paramètres par loi statistique) et d'exprimer nos résultats en fonction de ces paramètres. Comme précédemment évoqué, cette utilisation de lois statistiques facilite également la possibilité d'étudier des hétérogénéités corrélées en introduisant un coefficient de corrélation entre les deux lois décrivant les hétérogénéités.

Dans un premier temps, nous calculons l'équilibre de notre modèle. Nous étudions en-

suite l'impact de l'hétérogénéité des investisseurs sur certaines caractéristiques d'équilibre du marché et, en particulier, nous analysons la manière dont la corrélation entre les deux types d'hétérogénéité les affecte. Enfin, en nous focalisant sur une économie où les hétérogénéités ne sont pas corrélées, nous déterminons les caractéristiques de l'agent qui survit à long terme et celles des agents qui maximisent leur utilité attendue (à la fois d'un point de vue ex-ante, c'est-à-dire avant que tous les événements à venir se réalisent et d'un point de vue ex-post, après leur réalisation).

L'hétérogénéité des croyances et celle des taux de préférences temporelles ont tout d'abord un impact sur le taux sans risque et le prix de marché du risque. Plus précisément, contrairement au cas homogène (c'est-à-dire celui où tous les agents sont identiques), nous obtenons des prix de marché du risque anticycliques et des taux sans risque procycliques. Ces résultats, obtenus en considérant un continuum d'agents hétérogènes, sont en accord et complètent de précédentes études, comme celle de Jouini and Napp (2011) ou celle de Bhamra and Uppal (2014), qui considèrent des modèles avec seulement deux groupes d'investisseurs. Ils sont également conformes aux observations empiriques selon lesquelles « les primes de risque sur les actions semblent être plus élevées au creux des cycles économiques qu'au sommet » (Campbell and Cochrane, 1999) et que le taux à court terme est un indicateur procyclique de l'activité économique (voir, par exemple, Friedman, 1986). Nous trouvons que ces effets sont amplifiés lorsqu'il existe un biais optimiste moyen dans l'économie (c'est-à-dire quand la moyenne de la loi statistique utilisée pour décrire le biais des croyances des investisseurs est positive) et réduits lorsque les investisseurs sont en moyenne pessimistes. De plus, contrairement aux études précédentes, nous constatons que plus les différentes croyances des agents sont dispersées, plus ces effets sont importants. En étudiant la dynamique temporelle de ces résultats, nous constatons en outre que le prix de marché du risque asymptotique est égal à celui obtenu dans un cadre homogène, alors que le taux sans risque asymptotique est inférieur à celui du cas standard. Cela est dû au fait que seuls les investisseurs les plus patients, caractérisés par un faible taux de préférences temporelles, survivent à long terme.

L'étude de la corrélation entre les deux types d'hétérogénéité conduit à une conclusion supplémentaire relative à ces deux caractéristiques du marché : dans le cas où la corrélation est négative, et où les agents optimistes sont donc également les plus patients, le taux sans risque est réduit et le prix de marché du risque augmenté. Ce nouveau résultat est intéressant à la lumière des difficultés rencontrées par la littérature pour expliquer les valeurs observées de la prime de risque et du taux sans risque.

Dans la suite de notre analyse, nous considérons un actif financier dont le processus de dividendes est donné par la dotation totale de l'économie et étudions sa volatilité (étant donné qu'il n'y a qu'un actif dans l'économie que nous considérons, la volatilité de cet actif peut donc être vue comme étant la volatilité du marché). Notre principal résultat est alors de démontrer qu'une corrélation non nulle entre les deux types d'hétérogénéité a un impact sur la volatilité du marché. Plus précisément, si nous supposons une corrélation négative entre elles, nous observons une volatilité excessive sur le marché. De plus, cet effet est d'autant plus important que les agents sont en moyenne patients. Ce résultat nouveau fait écho à celui trouvé par Atmaz and Basak (2018), qui déduisent un résultat similaire de volatilité excessive dans le cadre d'un continuum d'agents qui ne sont hétérogènes qu'en termes de croyances. Dans leur travail, le principal facteur explicatif de la volatilité excessive est la dispersion des croyances, alors qu'elle ne joue aucun rôle dans notre analyse. En effet, nous considérons des agents ayant des fonctions d'utilité logarithmiques, qui sont connues pour atténuer l'effet de la dispersion des croyances, et c'est la corrélation entre les deux types d'hétérogénéité qui est à l'origine de notre résultat. Dans une économie à horizon fini peuplée d'un agent patient et d'un agent impatient, Li (2007) obtient un résultat similaire : il constate que la volatilité du marché est supérieure à la volatilité de la dotation globale si l'agent patient est plus optimiste que l'agent impatient. Par conséquent, notre résultat étend sa constatation au cas d'un continuum d'investisseurs, ce qui nous permet d'étudier plus en profondeur le rôle de la corrélation.

Une autre caractéristique de marché que nous étudions est le volume d'échange qui,

contrairement au cas où les agents sont homogènes et agissent donc de manière identique sans pouvoir trouver de contrepartie à leurs échanges, est non nul. En considérant l'économie dans son ensemble et en supposant une absence de corrélation, nous remarquons que ce volume diminue avec le temps du fait de la disparition progressive de certains investisseurs. Il dépend à l'inverse positivement de la dispersion des croyances, car des agents plus hétérogènes impliquent davantage de possibilités d'échanges. De même, à l'échelle individuelle, nous constatons que plus un investisseur est patient, plus son volume d'échange est important. Ce résultat est dû à un effet de richesse et se renforce lorsque la corrélation des hétérogénéités est élevée en valeur absolue. En revanche, il n'y a pas de relation claire entre le volume des transactions d'agents partageant le même taux de préférences temporelles et le biais de leur croyance.

Dans la dernière partie de ce premier chapitre, nous étudions les caractéristiques de certains agents spécifiques dans le cas où la corrélation des hétérogénéités est nulle. Plus précisément, nous déterminons les caractéristiques de l'investisseur survivant, c'est-à-dire celui dont la part de consommation de richesse reste non nulle asymptotiquement, et celles des investisseurs qui ont le plus haut niveau d'utilité tout au long de leur vie. Comme Yan (2008), nous obtenons tout d'abord que l'agent survivant de l'économie est l'agent le plus patient des agents dont le biais de croyance est nul. De même, en formant des groupes d'investisseurs partageant le même taux de préférences temporelles, nous constatons que l'agent rationnel de chaque groupe est celui qui survit à long terme. Néanmoins, nous montrons également que, pour chacun des groupes, cet agent rationnel n'est pas nécessairement celui qui a le plus haut niveau d'utilité (à la fois ex-ante et ex-post). En particulier, si, par exemple, l'économie dans son ensemble présente en moyenne un biais en faveur de l'optimisme, l'agent qui maximisera son utilité ex-post sera un agent caractérisé par un biais de croyance positif non nul (mais inférieur au biais moyen de l'économie). Ce résultat s'inscrit dans la lignée de celui trouvé par Jouini and Napp (2016) qui montrent que « même s'ils sont éliminés à long terme, les agents irrationnels pourraient rationnellement rester irrationnels ». En d'autres termes, cela

correspond au fait que, dans certaines circonstances, une vie plus courte (c'est-à-dire associée à un taux de survie plus faible) peut être plus gratifiante en termes d'utilité qu'une longue vie. Ainsi, notre résultat vient enrichir la littérature déjà existante car nous analysons un continuum d'agents qui sont donc preneurs de prix. En effet, dans leur étude, Jouini and Napp (2016) considèrent un modèle à deux agents où chaque agent a donc un impact sur le prix d'équilibre ; il est alors naturel que le biais d'un agent puisse lui profiter via son impact sur le prix. Nous constatons que ce résultat est toujours vérifié dans une économie peuplée d'une infinité d'investisseurs, où l'impact sur les prix de chacun d'eux est nul.

Enfin, à l'instar de Detemple and Murthy (1994), nous observons que notre économie se caractérise par des vagues d'optimisme et de pessimisme, puisque la part de consommation agrégée des agents optimistes augmente dans les bons états du monde et diminue, au profit de celles des agents pessimistes, dans les mauvais états du monde.

Deuxième chapitre

Le deuxième chapitre de ma thèse, intitulé *Disagreeing forever: a testable model with non-vanishing belief heterogeneity*, poursuit l'étude de la question de l'impact de l'hétérogénéité des croyances sur les caractéristiques du marché. L'une des contributions théoriques principales de ce travail est que je considère un modèle stationnaire où l'hétérogénéité des croyances ne disparaît pas à mesure que le temps passe et où, en conséquence, les effets de cette hétérogénéité sur les caractéristiques du marché sont persistants. L'avantage d'une telle propriété de stationnarité, qui ne se retrouve pas dans le papier de Atmaz and Basak (2018) qui sont les plus proches de mon travail, est double. Tout d'abord, empiriquement, il apparaît que les agents ne s'accordent pas plus sur les rendements futurs du marché de nos jours qu'ils ne le faisaient dans le passé ; la persistance de divergence des croyances traduit donc la réalité de façon plus pertinente. De plus, la stationnarité du modèle permet d'en tester les principales implications à l'aide de données réelles sur une longue période. Ce chapitre comporte ainsi à la fois une partie théorique et une partie empirique, qui confirme les

résultats obtenus théoriquement. Il est également important de souligner que la stationnarité du modèle permet d'en déduire de nouvelles implications. En effet, cette propriété implique que les relations obtenues sont vérifiées quelque soit l'horizon considéré, ce que j'examine et confirme empiriquement.

Dans la partie théorique de ce chapitre, je développe, dans un cadre d'équilibre général dynamique, un modèle à générations imbriquées où chacune des générations est constituée d'un continuum d'agents, eux-mêmes membres de différentes familles ayant des croyances hétérogènes. De générations en générations, je fais en effet l'hypothèse que chaque investisseur donne naissance à un investisseur ayant les mêmes croyances que lui, ce qui est l'un des mécanismes permettant à l'hétérogénéité des croyances de persister dans le temps. Comme dans le modèle présenté dans le chapitre 1, je me sers en outre d'une distribution statistique pour décrire la répartition de la richesse des agents d'une génération donnée. En particulier, la résolution du modèle mène à une distribution ayant un écart-type constant (ce qui garantit la persistance de l'hétérogénéité des croyances) et une moyenne qui dépend positivement de l'état du monde : dans les bons états du monde, les agents sont plus optimistes en moyenne et dans les mauvais, plus pessimistes.

Il est également utile d'expliquer la dynamique temporelle de ce modèle. Plus précisément, la chronologie de la vie de chacun des agents d'une génération donnée est la suivante. À leur naissance (par exemple à la date t), ils planifient les futures actions qu'ils réaliseront à la fin de leur vie d'une durée T afin de maximiser leur utilité (il est d'ailleurs important de préciser que je fais l'hypothèse que tous les agents sont dotés d'une fonction d'utilité CRRA, ce qui diffère de la fonction d'utilité logarithmique utilisée dans le chapitre 1 qui est plus restrictive mais qui présente des avantages calculatoires pour tenir compte de la corrélation avec l'hétérogénéité en préférences temporelles). Puis, à la date suivante (à la date $t + dt$), les agents donnent naissance à la nouvelle génération d'investisseurs qui ont également une durée de vie T . Enfin, à la fin de leur vie (à la date $t + T$), ils réalisent ce qu'ils avaient planifié. Plus particulièrement, les investisseurs consomment une partie de leur richesse et

lèguent la partie restante à la génération suivante. Cette nouvelle génération d'agents hérite donc de la part de richesse non consommée et peut à son tour répéter les mêmes actions : juste après avoir hérité (c'est-à-dire à la date $t + dt + T$), ils vont réaliser les actions qu'ils avaient planifiées à leur naissance (à la date $t + dt$) et notamment léguer une partie de leur héritage à leurs successeurs (nés à la date $t + dt + dt$), qui le consommeront à la date suivante (à la date $t + dt + dt + T$), etc. Cette dynamique résulte ainsi en une économie avec consommation effective à chaque date, ce qui n'est pas le cas du modèle à horizon fini proposé par Atmaz and Basak (2018) où la seule date de consommation intervient à la date terminale du modèle (à laquelle l'hétérogénéité des croyances a totalement disparu).

Par ailleurs, dans ce chapitre, je développe également un modèle alternatif, succinctement décrit dans le paragraphe suivant, menant à des résultats équivalents et facilitant la construction d'un agent représentatif intertemporel.

Ce modèle alternatif repose sur la présence d'un continuum d'investisseurs ayant une durée de vie infinie, qui révisent continuellement leurs plans de consommation et décalent continuellement leur date de consommation, conduisant ainsi à un modèle à horizon glissant (et sans consommation effective). Dans ce cadre, chacun des agents correspond approximativement à une famille entière du modèle à générations hétérogènes imbriquées et les révisions successives des plans de consommation coïncident avec les différents plans établis par les générations successives. Cette approche alternative fait écho au travail séminal de Lindahl (1939) qui observe que « les plans des sujets économiques à un instant donné ne sont pas totalement cohérents ni entre eux ni avec les conditions extérieures et doivent donc être révisés successivement ». Ceci est également conforme à la théorie plus générale de l'équilibre général temporaire de Grandmont (1977, 2008).

La résolution de ces modèles permet de caractériser l'équilibre et d'en déduire les résultats théoriques (stationnaires) de l'étude, qui sont décrits ci-dessous.

Tout d'abord, en examinant le prix de l'actif, je trouve qu'il dépend positivement du biais moyen des croyances, ce qui est conforme aux études de Jouini and Napp (2007) et de

Kurz and Motolese (2011). Je déduis également que l'impact de la dispersion des croyances est positif pour les états du monde suffisamment bons et négatif pour les états suffisamment mauvais. Enfin, comme d'autres études théoriques (Veronesi, 1999, Xu, 2007, Atmaz and Basak, 2018), je trouve une relation convexe entre le prix de l'actif et la production attendue de l'économie. Cette convexité des prix implique que le cours de l'actif réagit davantage aux bonnes nouvelles qu'aux mauvaises, et que la réaction du cours de l'actif à tout type de nouvelles est plus forte dans les états relativement bons. Basu (1997) et Nagel (2005) confirment empiriquement la première prédiction et, en accord avec la seconde, Conrad et al. (2002) montrent que le marché réagit plus fortement aux mauvaises nouvelles en période de prospérité qu'en période de crise.

J'étudie également la relation entre l'hétérogénéité des croyances et les rendements moyens de l'actif et observe que plus l'hétérogénéité est élevée, plus les rendements sont importants. La relation positive que je documente contribue ainsi à la littérature portant sur ce sujet, qui trouve des résultats contradictoires, et est conforme à la conjecture de Williams (1977) selon laquelle une plus grande dispersion des opinions représente un plus grand risque, et entraîne donc une rémunération attendue des agents plus élevée. Banerjee and Kremer (2010) confirment cette relation positive dans un modèle dynamique dans lequel les investisseurs sont en désaccord sur l'interprétation de l'information publique et Buraschi and Jiltsov (2006) obtiennent un résultat similaire en liant l'hétérogénéité des croyances aux positions ouvertes des options. À l'inverse, un autre courant de la littérature, fondé sur le travail précurseur de Miller (1977), documente un lien négatif qui dépend de manière critique de la présence de frictions sur le marché. De plus, je trouve que le rendement moyen de l'action diminue avec l'aversion au risque. En effet, dans une économie avec agents hétérogènes, les investisseurs plus averses au risque spéculent moins agressivement et obtiennent donc des rendements plus faibles.

Pour finir, j'observe que la volatilité de l'actif augmente de manière monotone avec la dispersion des croyances et qu'elle est toujours supérieure à la volatilité du processus de

production. Comme l'indiquent Atmaz and Basak (2018), cela s'explique par le fait que des fluctuations plus importantes du biais moyen des croyances se traduisent par des fluctuations supplémentaires du prix de l'actif et augmentent donc sa volatilité. Cette relation positive monotone est bien documentée sur le plan théorique (voir, par exemple, Shalen, 1993, dans un modèle à anticipations rationnelles à deux périodes, Scheinkman and Xiong, 2003, dans un modèle avec des contraintes de vente à découvert, Buraschi and Jiltsov, 2006, dans un modèle avec des agents rationnels ayant des informations incomplètes et hétérogènes, Andrei et al., 2019, dans un modèle avec désaccord sur la durée des cycles économiques). Ici, ma contribution est d'obtenir une formule stationnaire où les effets de l'hétérogénéité sur la volatilité restent persistants dans le temps.

Puisque les formules sont stationnaires, je peux alors traduire les principales implications théoriques du modèle en hypothèses testables et, dans la partie empirique du chapitre, examiner si ces hypothèses sont effectivement validées en utilisant des données de marché réelles et en effectuant des régressions par moindres carrés ordinaires. Plus précisément, je vérifie si une plus grande dispersion des croyances à l'échelle du marché prédit des rendements de marché plus élevés et une plus grande volatilité du marché. Alors que la plupart des études empiriques dans cette littérature se concentrent sur des données mensuelles, j'étudie également la validité de ces relations pour des données calculées sur des horizons plus longs.

Pour construire mes données de dispersion des croyances, j'utilise les prévisions mensuelles des analystes sur le taux de croissance à long terme du bénéfice par action d'actions individuelles de la base de données IBES Unadjusted Summary de janvier 1982 à décembre 2019. En adoptant une méthodologie similaire à celle de Yu (2011), je les agrège dans le temps (pour des horizons allant d'un mois à deux ans) et sur l'ensemble des actifs pour obtenir des données relatives au marché dans son ensemble. J'utilise en outre des données de la base CRSP pour construire mes variables de rendements de marché à différents horizons. En particulier, il est important de noter que mon analyse empirique diffère de celles existantes car j'étudie spécifiquement les rendements d'un indice de marché qui est construit à partir

de tous les actifs utilisés dans la construction de la variable de dispersion des croyances. Cela me permet de capturer plus directement le lien entre les caractéristiques du marché et les croyances des investisseurs. De même, je construis les données de volatilité du marché à partir des rendements quotidiens de cet indice de marché.

Les tests empiriques confirment la relation positive prédite entre la dispersion des croyances et les rendements du marché pour la plupart des horizons étudiés. Ainsi, en considérant un cadre avec une dispersion des croyances persistante et en utilisant des données plus étroitement liées, je montre de nouveaux résultats sur les impacts à long terme. Enfin, les résultats concernant l'impact sur la volatilité du marché sont plus mitigés. En effet, bien que je trouve une majorité de coefficients positifs (en contrôlant la volatilité passée), ils ne sont pas statistiquement significatifs. Des régressions à fenêtre glissante montrent en outre que le signe et l'intensité de l'impact de la dispersion des croyances varient dans le temps, ce qui pourrait expliquer la faiblesse des résultats. D'autres travaux empiriques étudient cette relation et confirment la relation positive prédite par le modèle (voir, par exemple, Ajinkya and Gift, 1985 qui utilisent des données sur une période de dix mois, Anderson et al., 2005 qui utilisent des données mensuelles sur une période de sept ans, Banerjee, 2011 dans un cadre d'analyse en coupe transversale).

Rendements non normaux

L'hypothèse de normalité des rendements est également une hypothèse courante dans les modèles théoriques d'évaluation d'actifs. S'appuyer sur les propriétés de la loi normale permet en effet de grandement faciliter la résolution de nombreux calculs et d'exprimer les résultats de manière simple. De plus, il arrive que les distributions empiriques des rendements d'actifs se rapprochent d'une distribution normale. De prime abord, une telle hypothèse apparaît donc appropriée, ce qui explique qu'elle soit largement utilisée.

La loi normale présente néanmoins un défaut fondamental : étant donné qu'il s'agit d'une

distribution symétrique, son coefficient d'asymétrie (également appelé de manière équivalente « skewness » dans le reste de la thèse) est nul. Autrement dit, à moyenne et variance données, faire l'hypothèse que les rendements d'un actif suivent une loi normale revient à supposer que les réalisations positives des rendements de cet actif ont autant de chances de survenir que celles négatives. Or, de nombreuses études montrent que ce n'est pas le cas. En considérant le marché des actions américaines dans son ensemble, Albuquerque (2012) montre par exemple qu'un krach est plus susceptible de survenir qu'un boom, ce qui revient à supposer que la distribution de ces rendements a un skewness négatif. Ainsi, l'utilisation de la loi normale ne permet pas de tenir compte des effets (potentiellement importants) du skewness des rendements des actifs et n'est donc pas adaptée pour décrire leur distribution. À l'inverse, bien que moins intuitives, l'usage de distributions asymétriques permet une étude plus complète.

Dans le troisième chapitre de cette thèse, intitulé *Two skewed risks* et coécrit avec Paul Karehnke, nous définissons et utilisons une distribution asymétrique afin d'étudier, dans un cadre bivarié, les impacts du skewness sur différentes métriques. Plus précisément, nous nous intéressons aux questions de l'impact de l'asymétrie des rendements sur les caractéristiques des actifs, mais aussi sur les choix de portefeuille et sur certaines mesures utilisées en gestion des risques. La principale contribution de ce chapitre est donc de déduire des enseignements économiques sur les effets de l'asymétrie dans un cadre bivarié simple et intuitif. L'utilisation d'un tel cadre complète l'étude de De Roon and Karehnke (2017) qui s'intéressent à un cadre univarié, et permet également de mieux cerner les interactions entre skewness et corrélation. En outre, comme chacun des deux actifs étudiés (ou, dans la partie relative à la gestion des risques, des deux risques) peut lui-même être considéré comme issu de l'agrégation de plusieurs actifs (ou risques), notre étude permet de proposer un cadre d'analyse général. Là encore, le choix d'un cadre bivarié et d'une distribution spécifique pour les rendements répond à l'impératif de trouver un juste milieu entre traçabilité et simplification d'une part, et réalisme et pertinence de l'autre.

Revenons tout d’abord sur la construction de la distribution que nous utilisons pour notre analyse. Cette distribution, que nous appelons la distribution normale fractionnée bivariée (*“split bivariate normal distribution”*), est une extension bivariée de la distribution normale fractionnée (*“split normal distribution”*) de Fechner (1897), qui est elle-même un cas particulier de la distribution t biaisée (*“skewed t distribution”*) de Hansen (1994), plus largement utilisée. Notre version bivariée est en outre très similaire aux développements plus généraux de Geweke (1989), Bauwens and Laurent (2005), et Villani and Larsson (2006), et elle coïncide avec ces distributions lorsque les deux risques sont indépendants.

Les principaux avantages de notre distribution sont sa simplicité et son intuition qu’elle hérite de la loi normale usuelle, ce qui permet de nombreux calculs explicites, tout en ayant la capacité d’avoir un skewness non nul. La distribution normale fractionnée bivariée généralise en effet la distribution normale bivariée et introduit de l’asymétrie en permettant aux paramètres d’écart-type de chacune des variables considérées de différer au-dessus et en dessous du mode. Ainsi, elle est particulièrement intéressante pour les applications financières, car elle tient compte du fait que les moyennes et les variances des rendements des actifs sont liées empiriquement (Duffee, 1995). Nous nous assurons en outre de la continuité de notre distribution en utilisant des facteurs d’ajustement pour chacun de ses quadrants. Enfin, une évaluation de la pertinence empirique de la distribution normale fractionnée bivariée nous permet de confirmer qu’elle fournit une description raisonnable des rendements des actifs. Plus précisément, nous analysons les rendements excédentaires d’indices de plusieurs catégories d’actifs et proposons différents tests de notre distribution ainsi que de trois lois connexes : la distribution normale bivariée, la distribution bivariée skew-normale (*“bivariate skew-normal distribution”*) d’Azzalini and Dalla-Valle (1996), et une version bivariée de la distribution t biaisée construite avec une copule normale. Les tests de Kolmogorov-Smirnov bivariés, une comparaison des co-skewness implicites et empiriques et une comparaison des niveaux théoriques et empiriques du déficit attendu conditionnel (que nous définirons plus en détail par la suite) montrent notamment la supériorité de la distribution normale frac-

tionnée bivariée sur la loi normale bivariée usuelle. Notre distribution offre également une meilleure description des données dans les queues de distributions que les deux autres lois asymétriques considérées, ce qui est particulièrement utile dans une optique de l'étude de la gestion des risques. La distribution normale fractionnée bivariée apparaît donc comme une loi pertinente pour analyser théoriquement l'impact du skewness dans un cadre bivarié.

Dans un premier temps, nous utilisons cet outil pour étudier les choix de portefeuille. Pour ce faire, nous considérons deux actifs dont les rendements suivent la distribution et étudions les poids optimaux à investir dans chacun d'eux en fonction des caractéristiques de leurs rendements. La moyenne et la variance de ces rendements sont fixes et nous faisons varier leur niveau d'asymétrie ainsi que leur corrélation afin d'étudier les implications de telles variations. En nous intéressant aux rendements certains équivalents, nous montrons que les gains d'utilité liés à la prise en compte de l'asymétrie ont tendance à diminuer avec la corrélation lorsque les asymétries sont égales. En revanche, lorsque le niveau de skewness des rendements de chacun des actifs diffère, la relation prend une forme de U : les gains sont plus importants pour les actifs fortement corrélés que pour les actifs modérément corrélés. En outre, l'étude des poids d'investissement optimaux souligne qu'il devient alors optimal de ne pas diversifier autant que le cadre moyenne-variance classique le suggérerait, en particulier pour les niveaux élevés de corrélation. Cette sous-diversification va de pair avec une hausse du skewness des portefeuilles associés. Notre analyse des choix de portefeuille souligne donc le rôle de la corrélation dans l'arbitrage entre diversification et asymétrie et complète ainsi le travail de Mitton and Vorkink (2007) qui ne considèrent qu'un seul actif asymétrique.

Nous étudions également l'impact du skewness des rendements dans un modèle d'évaluation d'actifs où les deux actifs constituent l'intégralité du marché. Dans une économie d'échanges pure, nous montrons que les alphas du modèle d'évaluation des actifs financiers (MEDAF) sont non nuls lorsque les actifs présentent des niveaux d'asymétrie différents, et qu'il y a donc un écart avec les résultats induits par le MEDAF. La littérature empirique s'est intéressée à cet écart et différentes explications potentielles ont été apportées : l'écart a été relié au skew-

ness systématique (Harvey and Siddique, 2000), au skewness idiosyncratique (Boyer et al., 2010) et aux rendements maximaux (Bali et al., 2011). Notre analyse théorique permet de justifier de tels résultats empiriques. En effet, nous montrons que l’actif qui sous-performe (celui dont l’alpha du MEDAF est négatif) présente également un skewness systématique, un skewness idiosyncratique et des rendements maximaux plus élevés que ceux de l’autre actif. Si l’on fait en outre l’hypothèse que l’actif ayant une asymétrie plus positive est également plus volatile (Kumar, 2009 montre que cette hypothèse est vérifiée empiriquement), notre cadre d’analyse permet de rendre compte de l’anomalie liée aux faibles rendements des actifs ayant des bêtas importants (Frazzini and Pedersen, 2014, Bali et al., 2017). Il est important de préciser l’intuition derrière de tels résultats. Ici, les différences de rendements maximaux découlent directement de la distribution asymétrique sous-jacente des rendements, et les différences de skewness idiosyncrasiques et systématiques résultent des conditions d’équilibre du marché. La préférence de l’investisseur représentatif en matière d’asymétrie conduit alors à des prix d’actifs qui s’écartent du MEDAF. En outre, notre analyse montre que l’actif à haut rendement maximal est à la fois une loterie (Bali et al., 2011) et une assurance (Barinov, 2018). Plus précisément, si la sous-performance de l’actif est liée à des rendements maximaux élevés et à un fort skewness idiosyncratique, elle est principalement due à un fort skewness systématique. Ceci est empiriquement confirmé par Barinov (2018) et Langlois (2020).

Afin de compléter notre étude, nous nous intéressons en outre à une seconde théorie comportementale qui accorde une importance plus élevée à l’asymétrie : la théorie cumulative des perspectives aléatoires (*“cumulative prospect theory”*) de Tversky and Kahneman (1992). Les résultats relatifs au choix de portefeuille et à l’évaluation d’actifs décrits jusqu’à maintenant sont en effet obtenus en supposant que l’agent représentatif effectue ses choix en suivant la théorie de l’utilité attendue. Cette approche présente l’avantage de permettre la résolution explicite d’un grand nombre de calculs mais, comme montré récemment par Ebert and Karehnke (2020), n’accorde qu’une faible importance—de troisième ordre—au skewness. À l’inverse, la théorie cumulative des perspectives aléatoires lui accorde une importance de

premier ordre, bien qu'il ne soit pas possible d'obtenir les résultats directement par le calcul et que nous soyons obligés de procéder à des simulations numériques. Comme attendu, nous montrons que ces préférences alternatives amplifient l'importance de l'asymétrie pour les poids de portefeuille, les équivalents certains et les écarts par rapport au MEDAF.

Enfin, nous nous intéressons également à l'impact joint du skewness et de la corrélation dans le cadre de la gestion des risques. En effet, comme la distribution normale fractionnée bivariée présente des similitudes avec la loi normale bivariée usuelle, elle permet une généralisation simple des formules de certaines mesures de risque conditionnel au cas des rendements asymétriques. Plus précisément, Adrian and Brunnermeier (2016) proposent une mesure appelée ΔCoVaR , définie comme la différence entre la valeur à risque (*"value-at-risk"*) d'un actif donné lorsqu'un risque conditionnel connaît un événement très négatif et cette valeur à risque lorsque le risque conditionnel est dans son état médian. Acharya et al. (2017) proposent également une mesure de risque conditionnel, ici nommée déficit attendu conditionnel (*"conditional expected shortfall"*), qui est définie comme la valeur attendue d'un actif risqué conditionnellement à un autre risque se trouvant dans la queue gauche de sa distribution. Nous fournissons des expressions simples pour chacune de ces mesures sous l'hypothèse que les risques concernés suivent la distribution normale fractionnée bivariée. Ainsi, lorsque la corrélation entre les deux risques considérés est positive, les deux mesures sont plus importantes en valeur absolue lorsque le risque principal a un skewness négatif. Enfin, nous montrons que l'impact quantitatif du skewness est important : pour, par exemple, une corrélation de 0,6, une probabilité de 5%, un risque conditionnel symétrique et une asymétrie de $-0,6$ au lieu de zéro pour le risque principal, ΔCoVaR est environ 30% plus élevée en valeur absolue. L'impact de l'asymétrie du risque conditionnel sur chacune des mesures est en revanche très limité, ce qui corrobore la conclusion empirique de Adrian and Brunnermeier (2016) que ΔCoVaR est faiblement corrélée avec la valeur à risque du risque conditionnel.

Bibliography

- V. V. Acharya, L. H. Pedersen, T. Philippon, and M. Richardson. Measuring systemic risk. *Review of Financial Studies*, 30(1):2–47, 2017.
- T. Adrian and M. K. Brunnermeier. Covar. *American Economic Review*, 106(7):1705–1741, 2016.
- B. B. Ajinkya and M. J. Gift. Dispersion of financial analysts’ earnings forecasts and the (option model) implied standard deviations of stock returns. *The Journal of Finance*, 40(5):1353–1365, 1985.
- R. Albuquerque. Skewness in stock returns: reconciling the evidence on firm versus aggregate returns. *Review of Financial Studies*, 25:1630–1673, 2012.
- T. G. Andersen, T. Bollerslev, F. X. Diebold, and H. Ebens. The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1):43–76, 2001.
- E. W. Anderson, E. Ghysels, and J. L. Juergens. Do heterogeneous beliefs matter for asset pricing? *Review of Financial Studies*, 18(3):875–924, 2005.
- D. Andrei, B. Carlin, and M. Hasler. Asset pricing with disagreement and uncertainty about the length of business cycles. *Management Science*, 65(6):2900–2923, 2019.
- A. Ang and J. Chen. Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443–494, 2002.

BIBLIOGRAPHY

- A. Ang, R. J. Hodrick, Y. Xing, and X. Zhang. The cross-section of volatility and expected returns. *The Journal of Finance*, 61(1):259–299, 2006.
- A. Atmaz and S. Basak. Belief dispersion in the stock market. *Journal of Finance*, 73(3):1225–1279, 2018.
- D. Avramov, T. Chordia, G. Jostova, and A. Philipov. Dispersion in analysts’ earnings forecasts and credit rating. *Journal of Financial Economics*, 91(1):83–101, 2009.
- A. Azzalini and A. Dalla-Valle. The multivariate skew-normal distribution. *Biometrika*, 83(4):715–726, 1996.
- S. Babić, C. Ley, and D. Veredas. Comparison and classification of flexible distributions for multivariate skew and heavy-tailed data. *Symmetry*, 11(10):1216, 2019.
- K. Back, A. D. Crane, and K. Crotty. Skewness consequences of seeking alpha. *The Review of Financial Studies*, 31(12):4720–4761, 2018.
- T. G. Bali, N. Cakici, and R. F. Whitelaw. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99:427–446, 2011.
- T. G. Bali, S. J. Brown, S. Murray, and Y. Tang. A lottery-demand-based explanation of the beta anomaly. *Journal of Financial and Quantitative Analysis*, 52(6):2369–2397, 2017.
- S. Banerjee. Learning from prices and the dispersion in beliefs. *Review of Financial Studies*, 24(9):3025–3068, 2011.
- S. Banerjee and I. Kremer. Disagreement and learning: Dynamic patterns of trade. *The Journal of Finance*, 65(4):1269–1302, 2010.
- N. Barberis and M. Huang. Stocks as lotteries: the implications of probability weighting for security prices. *American Economic Review*, 98:2066–2100, 2008.

- A. Barinov. Stocks with extreme past returns: Lotteries or insurance? *Journal of Financial Economics*, 129(3):458 – 478, 2018.
- S. Basak. Asset pricing with heterogeneous beliefs. *Journal of Banking and Finance*, 29(11): 2849–2881, 2005.
- S. Basu. The conservatism principle and the asymmetric timeliness of earnings. *Journal of Accounting and Economics*, 24(1):3–37, 1997.
- L. Bauwens and S. Laurent. A new class of multivariate skew densities, with applications to generalized autoregressive conditional heteroscedasticity models. *Journal of Business & Economic Statistics*, 23(3):346–354, 2005.
- A. Beber, A. Buraschi, and F. Breedon. Difference in beliefs and currency risk premia. *Journal of Financial Economics*, 98:415–438, 2010.
- G. S. Becker and C. B. Mulligan. The endogenous determination of time preference. *The Quarterly Journal of Economics*, 112(3):205–223, 1997.
- A. Beddock and E. Jouini. Live fast, die young: equilibrium and survival in large economies. *Economic Theory*, 71(3):961–996, 2021.
- S. Benoit, J.-E. Colliard, C. Hurlin, and C. Pérignon. Where the risks lie: A survey on systemic risk. *Review of Finance*, 21(1):109–152, 2017.
- T. Berrada, J. Detemple, and M. Rindisbacher. Asset pricing with beliefs-dependent risk aversion and learning. *Journal of Financial Economics*, 128:504–534, 2018.
- H. Bhamra and R. Uppal. Asset prices with heterogeneity in preferences and beliefs. *The Review of Financial Studies*, 27:519–580, 2014.
- M. Bianchi, R.-A. Dana, and E. Jouini. Shareholder heterogeneity, asymmetric information, and the equilibrium manager. *Economic Theory*, Forthcoming, 2021.

BIBLIOGRAPHY

- G. Bottazzi, P. Dindo, and D. Giachini. Long-run heterogeneity in an exchange economy with fixed-mix traders. *Economic Theory*, 66:407–447, 2018.
- B. Boyer, T. Mitton, and K. Vorkink. Expected idiosyncratic skewness. *Review of Financial Studies*, 23:169–202, 2010.
- R. G. Brandon and S. Wang. Earnings belief risk and the cross-section of stock returns. *Review of Finance*, 24(5):1107–1158, 2020.
- E. Britton, P. Fisher, and J. Whitley. The inflation report projections: understanding the fan chart. *Bank of England Quarterly Bulletin*, 38:30–37, 1998.
- M. K. Brunnermeier, C. Gollier, and J. A. Parker. Optimal beliefs, asset prices, and the preference for skewed returns. *American Economic Review*, 97:159–165, 2007.
- A. Buraschi and A. Jiltsov. Model uncertainty and option markets with heterogeneous agents. *Journal of Finance*, 61:2841–2897, 2006.
- A. Buraschi, F. Trojani, and A. Vedolin. Economic uncertainty, disagreement, and credit markets. *Management Science*, 60(5):1281–1296, 2014.
- J. Campbell and J. Cochrane. By force of habit: a consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.
- J. Y. Campbell. Asset pricing at the millennium. *The Journal of Finance*, 55(4):1515–1567, 2000.
- F. Chabi-Yo, D. P. Leisen, and E. Renault. Aggregation of preferences for skewed asset returns. *Journal of Economic Theory*, 154:453–489, 2014.
- Y. Chan and L. Kogan. Catching up with the joneses: heterogeneous preferences and the dynamics of asset prices. *Journal of Political Economy*, 110:1255–1285, 2002.

- J. Chen, H. Hong, and J. C. Stein. Breadth of ownership and stock returns. *Journal of Financial Economics*, 66(2-3):171–205, 2002.
- C. Conrad and A. Glas. 'Déjà vol' revisited: Survey forecasts of macroeconomic variables predict volatility in the cross-section of industry portfolios. SSRN Working paper, 2018. URL <https://ssrn.com/abstract=3186567>.
- J. Conrad, B. Cornell, and W. R. Landsman. When is bad news really bad news? *The Journal of Finance*, 57(6):2507–2532, 2002.
- J. Cvitanic and S. Malamud. Price impact and portfolio impact. *Journal of Financial Economics*, 100:201–225, 2011.
- J. Cvitanic, E. Jouini, S. Malamud, and C. Napp. Financial markets equilibrium with heterogeneous agents. *Review of Finance*, 16:285–321, 2012.
- Z. Da. Cash flow, consumption risk, and the cross-section of stock returns. *The Journal of Finance*, 64(2):923–956, 2009.
- M. Dahlquist, A. Farago, and R. Tédongap. Asymmetries and portfolio choice. *The Review of Financial Studies*, 30(2):667–702, 11 2016.
- R.-A. Dana. An extension of milleron, mitjushin and polterovich's result. *Journal of Mathematical Economics*, 24(3):259–269, 1995.
- S. Das, C. M. Kuhnen, and S. Nagel. Socioeconomic status and macroeconomic expectations. *Review of Financial Studies*, 33(1):395–432, 2020.
- A. David. Heterogeneous beliefs, speculation, and the equity premium. *Journal of Finance*, 63:41–83, 2008.
- L. I. De Castro and N. C. Yannelis. An interpretation of ellberg's paradox based on information and incompleteness. *Economic Theory Bulletin*, 1(2):139–144, 2013.

BIBLIOGRAPHY

- F. De Roon and P. Karehnke. A simple skewed distribution with asset pricing applications. *Review of Finance*, 21(6):2169–2197, 2017.
- L. Del Viva, E. Kasanen, and L. Trigeorgis. Real options, idiosyncratic skewness, and diversification. *Journal of Financial and Quantitative Analysis*, 52(1):215–241, 2017.
- J. Detemple and S. Murthy. Intertemporal asset pricing with heterogeneous beliefs. *Journal of Economic Theory*, 62:294–320, 1994.
- K. B. Diether, C. J. Malloy, and A. Scherbina. Differences of opinion and the cross section of stock returns. *The Journal of Finance*, 57(5):2113–2141, 2002.
- J. A. Doukas, C. F. Kim, and C. Pantzalis. Divergence of opinion and equity returns. *Journal of Financial and Quantitative Analysis*, 41(3):573–606, 2006.
- R. Duchin and M. Levy. Disagreement, portfolio optimization, and excess volatility. *Journal of Financial and Quantitative Analysis*, 45(3):623–640, 2010.
- G. R. Duffee. Stock returns and volatility a firm-level analysis. *Journal of Financial Economics*, 37:399–420, 1995.
- S. Ebert and P. Karehnke. Skewness preferences in choice under risk. Working Paper SSRN, 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3480519.
- S. Ebert and P. Strack. Until the bitter end: on prospect theory in a dynamic context. *American Economic Review*, 105(4):1618–33, 2015.
- S. Ebert and D. Wiesen. Testing for prudence and skewness seeking. *Management Science*, 57(7):1334–1349, 2011.
- G. T. Fechner. *Kollektivmasslehre*. Engelmann, 1897.
- B. Feunou, M. R. Jahan-Parvar, and R. Tédongap. Modeling market downside volatility. *Review of Finance*, 17(1):443–481, 2013.

- A. Frazzini and L. H. Pedersen. Betting against beta. *Journal of Financial Economics*, 111(1):1–25, 2014.
- S. Frederick, G. Loewenstein, and T. O’Donoghue. Time discounting and time preference: a critical review. *Journal of Economic Literature*, 40:351–401, 2002.
- K. R. French, G. W. Schwert, and R. F. Stambaugh. Expected stock returns and volatility. *Journal of Financial Economics*, 19(1):3–29, 1987.
- B. Friedman. Money, credit and interest rates in the business cycle. In *Robert Gordon, ed. The American Business Cycle*, University of Chicago Press:123–182, 1986.
- M. Friedman. The case for flexible exchange rates. *Essays in Positive Economics*, University of Chicago Press, 1953.
- J. Geweke. Bayesian inference in econometric models using monte carlo integration. *Econometrica*, 57(6):1317–1339, 1989.
- A. Ghosh and G. Roussellet. Identifying beliefs from asset prices. Working Paper SSRN, 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3400005.
- S. Giglio, M. Maggiori, J. Stroebl, and S. Utkus. Five facts about beliefs and portfolios. SSRN Working paper, 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3336400.
- E. D. Giorgi and T. Hens. Making prospect theory fit for finance. *Financial markets and portfolio management*, 20(3):339–360, 2006.
- C. Gollier and R. Zeckhauser. Aggregation of heterogeneous time preferences. *Journal of Political Economy*, 113(4):878–896, 2005.
- J. M. Grandmont. Temporary general equilibrium theory. *Econometrica*, 45(3):535, 1977.

BIBLIOGRAPHY

- J.-M. Grandmont. *Temporary Equilibrium*, chapter Temporary Equilibrium, pages 6579–6583. Palgrave Macmillan UK, London, 2008.
- S. Grossman and J. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70:393–408, 1980.
- B. E. Hansen. Autoregressive conditional density estimation. *International Economic Review*, 35(3):705–730, 1994.
- C. R. Harvey and A. Siddique. Conditional skewness in asset pricing tests. *Journal of Finance*, 55(3):1263–1295, 2000.
- F. Holzmeister, J. Huber, M. Kirchler, F. Lindner, U. Weitzel, and S. Zeisberger. What drives risk perception? A global survey with financial professionals and lay people. *Management Science*, 66(9):3977–4002, 2020.
- H. Hvide. Pragmatic beliefs and overconfidence. *Journal of Economic Behavior & Organization*, 48:15–28, 2002.
- J. E. Ingersoll. *Theory of Financial Decision Making*. Rowman & Littlefield Publishers, 1987.
- N. L. Johnson and S. Kotz. *Distributions in statistics: Continuous multivariate distributions*. Wiley, New York, 1972.
- E. Jondeau, S.-H. Poon, and R. M. *Financial modeling under non-Gaussian distributions*. Springer Finance. Springer London, London, 2007.
- E. Jouini and C. Napp. Consensus consumer and intertemporal asset pricing with heterogeneous beliefs. *Review of Economic Studies*, 74:1149–1174, 2007.
- E. Jouini and C. Napp. Unbiased disagreement in financial markets, waves of pessimism and the risk-return trade-off. *Review of Finance*, 15(3):575–601, 2011.

- E. Jouini and C. Napp. Live fast, die young. *Economic Theory*, 62(1-2):265–278, 2016.
- A. Kane. Skewness preference and portfolio choice. *Journal of Financial and Quantitative Analysis*, 17:15–25, 1982.
- P. Karehnke and F. de Roon. Spanning tests for assets with option-like payoffs: The case of hedge funds. *Management Science*, 66(12):5969–5989, 2020.
- E. K. Kelley and P. C. Tetlock. Why do investors trade? Working Paper SSRN, 2013. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2064144.
- L. Kogan, S. Ross, J. Wang, and M. Westerfield. The price impact and survival of irrational traders. *Journal of Finance*, 61:195–229, 2006.
- A. Kraus and R. H. Litzenberger. Skewness preference and the valuation of risk assets. *The Journal of Finance*, 31:1085–1100, 1976.
- A. Kumar. Who gambles in the stock market? *The Journal of Finance*, 64(4):1889–1933, 2009.
- M. Kurz and M. Motolese. Diverse beliefs and time variability of risk premia. *Economic Theory*, 47(2-3):293–335, 2011.
- A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.
- H. Langlois. Asset pricing with return asymmetries: Theory and tests. Working Paper SSRN, 2013. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2362944.
- H. Langlois. Measuring skewness premia. *Journal of Financial Economics*, 135(2):399–424, 2020.
- T. Li. Heterogeneous beliefs, asset prices, and volatility in a pure exchange economy. *Journal of Economic Dynamics & Controls*, 31:1697–1727, 2007.
- E. Lindahl. *Studies in the theory of money and capital*. London, 1939.

BIBLIOGRAPHY

- F. A. Longstaff and J. Wang. Asset pricing and the credit market. *The Review of Financial Studies*, 25(11):3169–3215, 2012.
- M. Meeuwis, J. Parker, A. Schoar, and D. Simester. Belief disagreement and portfolio choice. 2019. doi: 10.3386/w25108.
- R. Merton. On estimating the expected return on the market: an exploratory investigation. *Journal of Financial Economics*, 8:323–362, 1980.
- E. M. Miller. Risk, uncertainty, and divergence of opinion. *The Journal of Finance*, 32(4):1151–1168, 1977.
- T. Mitton and K. Vorkink. Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies*, 20:1255–1288, 2007.
- S. Nagel. Short sales, institutional investors and the cross-section of stock returns. *Journal of Financial Economics*, 78(2):277–309, 2005.
- W. K. Newey and K. D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.
- T. Odean. Do investors trade too much? *The American Economic Review*, 89(5):1279–1298, 1999.
- B. S. Paye. 'Déjà vol': Predictive regressions for aggregate stock market volatility using macroeconomic variables. *Journal of Financial Economics*, 106(3):527–546, 2012.
- J. A. Peacock. Two-dimensional goodness-of-fit testing in astronomy. *Monthly Notices of the Royal Astronomical Society*, 202(3):615–627, Mar. 1983.
- P. Samuelson. A note on measurement of utility. *Review of Economic Studies*, 4:155–161, 1937.

- A. Sandroni. Do markets favor agents able to make accurate prediction? *Econometrica*, 68: 1303–1341, 2000.
- J. A. Scheinkman and W. Xiong. Overconfidence and speculative bubbles. *Journal of Political Economy*, 111(6):1183–1220, 2003.
- C. Schneider and O. Spalt. Conglomerate investment, skewness, and the CEO long-shot bias. *The Journal of Finance*, 71(2):635–671, 2016.
- P. Schneider. An anatomy of the market return. *Journal of Financial Economics*, 132(2): 325–350, 2019.
- P. Schneider, C. Wagner, and J. Zechner. Low-risk anomalies? *The Journal of Finance*, 75 (5):2673–2718, 2020.
- G. W. Schwert. Why does stock market volatility change over time? *The Journal of Finance*, 44(5):1115–1153, 1989.
- C. T. Shalen. Volume, volatility, and the dispersion of beliefs. *Review of Financial Studies*, 6(2):405–434, 1993.
- Y. Simaan. Portfolio selection and asset pricing-three-parameter framework. *Management Science*, 39(5):568–577, 1993.
- A. Stuart and J. K. Ord. *Kendall’s advanced theory of statistics Volume 1 Distribution theory*, volume 1. Arnold, 1994.
- P. Theodossiou. Financial data and the skewed generalized t distribution. *Management Science*, 44(12-part-1):1650–1661, 1998.
- A. Tversky and D. Kahneman. Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:297–323, 1992.

BIBLIOGRAPHY

- P. Veronesi. Stock market overreactions to bad news in good times: A rational expectations equilibrium model. *Review of Financial Studies*, 12(5):975–1007, 1999.
- M. Villani and R. Larsson. The multivariate split normal distribution and asymmetric principal components analysis. *Communications in Statistics - Theory and Methods*, 35(6):1123–1140, 2006.
- T. Vuolteenaho. Understanding the aggregate book-to-market ratio. 1999.
- K. F. Wallis. The two-piece normal, binormal, or double gaussian distribution: Its origin and rediscoveries. *Statistical Science*, 29(1):106–112, 2014.
- M. Wang, M. O. Rieger, and T. Hens. How time preferences differ: Evidence from 53 countries. *Journal of Economic Psychology*, 52:115–135, 2016.
- M. Weitzman. Gamma discounting. *American Economic Review*, 91(1):260–271, 2001.
- J. T. Williams. Capital asset prices with heterogeneous beliefs. *Journal of Financial Economics*, 5(2):219–239, 1977.
- D. Won, G. Hahn, and N. C. Yannelis. Capital market equilibrium without riskless assets: heterogeneous expectations. *Annals of Finance*, 4:183–195, 2008.
- W. Xiong and H. Yan. Heterogeneous expectations and bond markets. *The Review of Financial Studies*, 23(4):1433–1466, 2010.
- C. Xiouros and F. Zapatero. The representative agent of an economy with external habit-formation and heterogeneous risk-aversion. *The Review of Financial Studies*, 23:3017–3047, 2010.
- J. Xu. Price convexity and skewness. *The Journal of Finance*, 62(5):2521–2552, 2007.
- H. Yan. Natural selection in financial markets: does it work? *Management Science*, 54(11):1935–1950, 2008.

- J. Yu. Disagreement and return predictability of stock portfolios. *Journal of Financial Economics*, 99:162–183, 2011.
- D. Zhu and J. W. Galbraith. A generalized asymmetric student- distribution with application to financial econometrics. *Journal of Econometrics*, 157(2):297–305, 2010.
- A. Ziegler. Why does implied risk aversion smile? *The Review of Financial Studies*, 20: 859–904, 2007.

RÉSUMÉ

Cette thèse examine les effets du relâchement d'hypothèses simplificatrices souvent formulées dans les modèles d'évaluation. Considérant un nombre infini d'investisseurs, les deux premiers chapitres étudient l'hétérogénéité des agents. Plus précisément, le premier explore les impacts de la corrélation des hétérogénéités. Il y a par exemple un excès de volatilité de marché quand les agents les plus optimistes sont aussi les plus patients. Le deuxième chapitre développe un modèle stationnaire où l'hétérogénéité des croyances persiste. Cela permet de tester empiriquement l'effet positif de la dispersion des croyances sur les rendements et la volatilité. Enfin, le troisième chapitre s'intéresse à la non-normalité des rendements : définissant une distribution bivariée asymétrique intuitive, le chapitre analyse, dans un cadre simple à deux actifs, comment le skewness et son interaction avec la corrélation affectent les choix de portefeuille, les prix des actifs et les mesures de risque.

MOTS CLÉS

Evaluation d'Actifs, Croyance, Hétérogénéité, Asymétrie, Corrélation

ABSTRACT

This thesis consists of three chapters and studies the consequences of releasing simplifying unrealistic assumptions often made in asset pricing models. Specifically, the first two chapters focus on agent heterogeneity and deal with models populated by a continuum of investors who agree to disagree. The first chapter highlights the impacts of correlated heterogeneities. In particular there is some excess market volatility when the most optimistic agents are also the most patient ones. The second chapter considers a stationary model with non-vanishing belief heterogeneity and allows for an empirical test of the model-implied positive effect of belief dispersion on returns and volatility. Lastly, the third chapter is interested in the non-normality of asset return distributions. Defining and using the split bivariate normal distribution, it analyzes in a simple two-asset framework how skewness and its interaction with correlation affect portfolio choice, asset prices, and risk metrics.

KEYWORDS

Asset Pricing, Belief, Heterogeneity, Skewness, Correlation