



HAL
open science

The influence of cognitive, learning and social interaction skills of investors on the price formation mechanism : an analysis helped by the conception of an financial market simulator

Lucian Daniel Stanciu-Viziteu

► To cite this version:

Lucian Daniel Stanciu-Viziteu. The influence of cognitive, learning and social interaction skills of investors on the price formation mechanism : an analysis helped by the conception of an financial market simulator. Business administration. Université de Grenoble, 2013. English. NNT : 2013GRENG008 . tel-00996394

HAL Id: tel-00996394

<https://theses.hal.science/tel-00996394v1>

Submitted on 26 May 2014

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

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE GRENOBLE

Spécialité : **Sciences de Gestion**

Arrêté ministériel : 7 août 2006

Présentée par

Lucian Daniel STANCIU-VIZITEU

Thèse dirigée par **Ollivier TARAMASCO**

préparée au sein du **Laboratoire Centre d'Etudes et de
Recherches Appliquées à la Gestion**
dans l'**École Doctorale de Sciences de Gestion**

L'influence des processus cognitif, d'apprentissage et d'interaction sociaux des investisseurs sur le processus de formation des prix : une analyse grâce à la conception d'un simulateur de marché financier

Thèse soutenue publiquement le « **5 Juin 2013** »,
devant le jury composé de :

Monsieur Pascal LOUVET

Professeur à l'IAE de Grenoble (Président du jury)

Monsieur Olivier BRANDOUY

Professeur à l'Université Paris 1 (Rapporteur)

Monsieur Ollivier TARAMASCO

Professeur à l'ENSIMAG Grenoble (Directeur de Thèse)

Monsieur Ion STANCU

Professeur à ASE Bucarest, Roumanie (Rapporteur)

Monsieur Hervé ALEXANDRE

Professeur à l'Université Paris Dauphine (Examineur)



UNIVERSITY OF GRENOBLE
DOCTORAL SCHOOL OF MANAGEMENT
STUDIES

P H D T H E S I S

to obtain the title of

PhD of Science

of the University of Grenoble

Specialty : FINANCE

Defended by

Lucian Daniel STANCIU-VIZITEU

**The influence of investor biases on
the formation of prices in financial
markets**

Thesis Advisor: Ollivier TARAMASCO

prepared at CERAG Laboratory, FINANCE Team

defended on 5 June, 2013

Jury :

<i>Reviewers :</i>	Ion STANCU	- ASE (Roumanie,Bucharest)
	Ollivier BRANDOUY	- University Paris 1
<i>Advisor :</i>	Olliver TARAMASCO	- ENSIMAG (Grenoble)
<i>Examinators :</i>	Pascal LOUVET	- IAE (Grenoble)
	Herve ALEXANDRE	- University Paris Dauphine

Acknowledgments

Thank you!

First of all I would like to thank the people that took the time to read and offer hints for improving my work, especially those that accepted a position in the defense jury: Prof. Ollivier Brandouy, Prof. Ion Stancu, Prof. Herve Alexandre and Prof. Pascal Louvet.

It has been three (and a half) years since I started this work. That is really a long time (I now have some grey hair) and during this period I have met a lot of people which helped to shape my ideas about how finance works, and about the world in general. Here it goes: thank you Pascal Louvet for trusting and helping me discover the best in myself, thank you to the people who decided I was worth getting the doctoral scholarship, a big thanks to my office buddies (Frank, Ingrid, Laura, Ali, Sarah, Cyrielle, Jessica, Omar, Zakaria, the sharks, etc.) - it was a lot of fun. Thank you Charles, Marie-Christine, Sarah, Bernard, Florence and Brigitte for your precious help. Thanks to my fellow computer-geeks Maxime, Lionel and Francois - you understood me when no one in the finance department did (because they couldn't). A special thank you to Marco Galbiati with whom I worked at the ECB - you really made me enjoy research. Thank you to all the Romanian people I met in Grenoble - it really helped me feel closer to home. And thank you kind reader for taking the time to look through this document. I really hope it will offer you some interesting insights into the world of financial markets.

PS: oh yes, I kept the best for the end. Thank you Ollivier Taramasco - I could not have had a better advisor. I learned a lot of things from you and I hope you enjoyed this experience as much as I did. Thank you Maria for your support and for offering me a handful of invaluable life lessons. Lastly, thank you God for giving me strength and helping me finish this work.

Abstract

We construct an agent-based computer simulated financial market. Trading in this market is not continuous. The market price is formed using a limit-order book. The modelled investors receive biased information and they attempt to maximize their wealth. Different traders, from noise to chartist and informed, coexist in the same market. We show how stylized facts can be formed by the presence of chartists or a simple lag in investor information. Price bubbles can arise when market prices are dominated by technical traders. Interestingly we show that well informed investors can earn more if they adopt, in special situations, a technical strategy. Using our results we propose a new model for market dynamics called "sometimes efficient markets". Moreover, we define the concept of "strategy-strong efficient markets".

Keywords computational finance, stock markets, efficiency, multi-agent, simulation, bubbles, stylized facts, risk management

Contents

1	Resume de these	5
2	Introduction	11
2.1	Context and motivation	11
2.2	Important concepts	14
2.3	Behavior biases and the problems they cause	16
2.4	Behavioral finance vs. efficient markets	18
2.5	Research questions	20
2.6	Research Methodology	21
2.7	Assumptions, advantages and limitations of this study	23
2.8	Rationality and irrationality	25
2.8.1	What is an irrational investor?	26
2.8.2	Are irrational investors present in financial markets?	27
2.9	A very short history of financial markets research	29
2.10	Summary of chapter	30
3	Object of study	33
3.1	Information and market prices	34
3.2	Asymmetry and quality of information	36
3.3	Investors	38
3.4	Behavior biases	39
3.5	Price formation mechanism	41
3.6	Price anomalies and explanations	43
3.7	Classical approaches to studying market dynamics	47
3.8	Summary	48
4	Agent based artificial stock markets	51
4.1	Complex systems	51
4.2	Agent based simulation	56

4.2.1	Definition and components	56
4.2.2	Applications	56
4.2.3	Advantages and disadvantages of research using agent-based simulations in finance	59
4.3	Topology of financial market simulators	61
4.3.1	Notable simulators	62
4.3.2	Simulated market	68
4.3.3	Simulated time	69
4.3.4	Information	69
4.3.5	Bias in information processing	70
4.3.6	Risk and time preferences	71
4.3.7	Agent learning	71
4.3.8	Agent strategy bias	73
4.3.9	Price formation	73
4.4	Critics of agent-based artificial stock markets as a methodology for research	74
4.4.1	Equivalence between an agent-based artificial market and a mathematical analytical model.	76
5	Contribution - Our simulator (LUMA)	79
5.1	Simulator, environment and assumptions	79
5.2	Information and fundamental value	81
5.3	Orders and the price formation mechanism	84
5.3.1	Orders	84
5.3.2	Price formation mechanism	86
5.4	The agents, their biases and investment decisions	91
5.4.1	The agent	91
5.4.2	The bias	93
5.4.3	The investment decision	94
5.5	Review of configuration parameters	102
6	Answering the research questions	105
6.1	Proof of concept validations - Benchmark simulations	105

6.2	Biases and market prices	113
6.2.1	Biases involving lags in information interpretation	124
6.3	Volatility clustering and other "stylized" facts	129
6.4	Bubbles: how prices can grow above the asset's fundamental value	138
6.5	Survival and evolution of agents	147
7	The shark game: equilibrium with bounded rationality	173
7.1	Introduction	173
7.2	The model	175
7.3	Results	179
7.4	Conclusions	183
8	Conclusions	185
8.1	Our contribution	185
8.2	Main findings	187
A	Appendix	201
A.1	Are financial markets ethically efficient?	201
A.2	Central Limit Theory for number estimations	205
A.3	Signal to noise Ratio	205
	Bibliography	207

Resume de these

Entre Juin 2007 et Novembre 2008, les citoyens américains ont perdu, en moyenne, un quart de leurs avoirs nets. Le total des actifs des fonds de pension a chuté de 22% (perte de 1.3 milliards de dollars). La perte subie par les investisseurs internationaux est du même ordre de grandeur. Une vision cynique, voire simpliste, consiste à dire que cette récente crise est le résultat de l'explosion d'une bulle, en autres termes celle-ci étant liée à un retour des marchés à l'efficience. Dans notre travail, nous prouvons que l'idée d'efficience est incomplète pour décrire un marché financier. À partir de nos résultats nous proposons la théorie des *marchés parfois efficient*.

Un marché financier est composé de nombreuses composantes hétérogènes et complexes. Jusqu'à récemment, les études financières se sont appuyées sur des bases théoriques tels que l'homogénéité des espérances des investisseurs, leur rationalité ou l'accès uniforme à l'information. La réalité des marchés ainsi que l'évolution des outils d'analyse ont ouvert la possibilité d'étudier les marchés dans un cadre moins contraignant et plus réaliste. Pour modéliser un marché financier, nous remplaçons le langage mathématique par le langage informatique. Ainsi, nous avons simulé un marché financier avec tous ses composantes : information, investisseurs et formation des prix.

Dans notre marché simulé, nous avons un seul actif risqué qui a une valeur fondamentale V_t , avec des rentabilités normalement distribuées. Cette valeur fondamentale est un concept théorique qui nous aide à construire un flux d'information avec un sens financier. On définit l'information parfaite comme $I_t = P_t - V_{t+1}$. À partir de l'information I_t les investisseurs, via leurs ordres, vont créer le nouveau prix de marché P_{t+1} que l'on peut considérer comme l'estimation de marché de la valeur fondamentale V_{t+1} . Tenant compte de l'hétérogénéité des interprétations et des biais des investisseurs, chacun va recevoir une version modifiée de l'information parfaite. Ainsi l'investisseur \mathbf{x} va recevoir l'information $I_{x,t} = a_x * I_t + b_x$. Les paramètres a_x et b_x nous permettent de créer une large variété de comportements des investisseurs : $a_x > 1$ révèle des investisseurs qui ont la tendance d'exagérer l'information, $b_x > 0$ indique des investisseurs optimistes et $b_x < 0$ des investisseurs pessimistes.

Les investisseurs modélisés dans notre système peuvent utiliser une de ces deux stratégies: une stratégie fondamentale (informée) ou une stratégie chartiste (non-informée). La stratégie fondamentale implique que l'investisseur suppose que le prix est bon s'il est dans un intervalle autour de la valeur fondamentale. L'investisseur

va estimer la valeur fondamentale en utilisant le prix de marché et l'information reçue ainsi: $E_{x,t}(V_{t+1}) = P_t + I_{x,t}$. En passant des ordres, ces investisseurs esperent gagner une rentabilité minimum r_{min} , qui est le reflet de l'aversion au risque et de l'espérance des gains relative a chaque investisseur. Donc, l'investisseur x , qui utilise une stratégie fondamentale, va considérer que le prix de marché est bon quand il est dans l'intervalle $[L, H]$ ou

$$[L = E_{x,t}(V_{t+1}) * (1 - r_{min}); H = E_{x,t}(V_{t+1}) * (1 + r_{min})]$$

Ainsi l'investisseur fondamental va vendre ses actifs, au prix minimum H, quand $P_t > H$ et il va acheter des actifs, au prix maximum L, quand $P_t < L$. Un tel investisseur espere profiter des erreurs d'évaluation du marché. L'investisseur chartiste, non-informé, espere profite des tendances des prix. Une chartiste regarde les rentabilités passées et s'il observe la formation d'une tendance il va essayer d'en profiter. Par exemple, s'il observe que les trois dernieres rentabilités ont été positives ($R_t * R_{t-1} * R_{t-2} > 0$), le chartiste va acheter des actifs en espérant que la tendance positive va continuer. Tous les investisseurs chartistes ont deux regles de longueurs L, une d'achat et une de vente, qui sont dans des séries de type $[a_1, a_2, \dots, a_L]$ et $[v_1, \dots, v_L]$. Les parametres a_i et v_i peuvent prendre les valeurs {indifférente, positive, négative}. Chaque regle est comparé avec la rentabilité passée, a_i ou v_i est compare avec R_{t-i+1} . Quand une regle d'achat/vente est activée le chartiste va acheter/vendre ses actifs au prix $P_t * (1 \pm N(\mu_c, \sigma_c))$ sont des parametres qui caractérisent les espérances de gains des chartistes. Par exemple, la regle d'achat [**positive, indifférent, negatif**] va etre activée, et le chartiste va acheter, quand les rentabilités passé vont etre du type $[R_t > 0, R_{t-2} < 0]$.

Les investisseurs vont envoyer leurs ordres, pendant une période de trading, et celles-ci vont etre enregistrés dans un carnet d'ordre. A la fin de la période de trading le carnet d'ordre est fermé et le nouveaux prix de marché est découvert. Ce prix est trouvé a la suite d'une procédure qui consiste a déterminer le prix qui :

1. Maximise le volume des transactions résolu
2. Minimise le nombre des transactions non résolu a la vente ou a l'achat
3. Minimise la pression de marché soit a la vente soit a l'achat. Cette méthode assure qu'un maximum d'information est intégré dans le prix et cette derniere n'introduit pas des biais supplémentaires.

Dans chaque simulation les données d'entrée sont représentées par quelques parametres caractéristiques :

1. le nombre et la distribution des investisseurs entre les differentes types
2. le nombre de jours de simulation
3. la distribution de la valeur fondamentale de l'actif $R_{F_t} \sim N(\mu_V, \sigma_V)$

4. l'intervalle des biais possibles des investisseurs fondamentaux $a_x \in [a_{min}, a_{max}]$ et $b_x \in [b_{min}, b_{max}]$
5. l'intervalle des rentabilité minimum espérée $r_{x,min} \in [r_{min}, r_{max}]$
6. la distribution des espérances des gains des chartistes avec moyenne $\mu_c \in [\mu_{min}, \mu_{max}]$ et variance $\sigma_c \in [\sigma_{min}, \sigma_{max}]$.

Notre outil de recherche, le simulateur de marchés financiers dénommé LUMA, nous sert à expliquer les liens de causalité entre les différentes stratégies des investisseurs et la formation des prix de marché.

Question de recherche #1 : *Quelle est la distribution des stratégies d'investissement qui peuvent produire, pour de longues périodes, des prix biaisés?*

Cette première question de recherche porte sur la possibilité qu'un marché peut rester dans un état inefficace pendant des longues périodes, même avec la présence d'investisseurs rationnels et parfaitement informés. Pour prévenir les éventuelles critiques d'efficacité, nous nous sommes également demandé si un investisseur biaisé peut survivre dans un marché financier. Si l'on trouve qu'il peut survivre, très longtemps, alors nous avons les bases pour discuter d'inefficacité des marchés.

Nous avons découvert que les investisseurs biaisés peuvent créer un marché avec des prix biaisés pendant de longues périodes de temps. En fonction de la distribution des biais des investisseurs, le marché peut avoir des prix sur ou sous évalués avec des magnitudes plus ou moins importantes. La condition nécessaire, afin d'avoir un marché avec des prix biaisés est que les investisseurs avec des biais peuvent détenir la plupart des actifs en circulation sur le marché.

Question de recherche #2 : *Un investisseur biaisé peut-il survivre dans un marché financier?*

Nos résultats montrent que les investisseurs biaisés peuvent survivre, pendant longtemps, dans les marchés. Il est cependant nécessaire que le marché ne retrouve, trop souvent, son état d'efficacité.

Question de recherche #3 : *Les investisseurs biaisés peuvent-ils gagner davantage que les investisseurs non biaisés ?*

Si l'investisseur biaisé peut survivre et les marchés peuvent être biaisés alors on peut se demander si des investisseurs bien informés ont intérêt à payer pour une information correcte. Si on trouve que l'information correcte n'est toujours pas rentable alors comment peut-on espérer qu'un marché peut être, autrement que par accident, efficace?

En effet, nous observons que lorsque les actifs sont surévalués alors les investisseurs bien informés vont avoir tendance à les vendre. Si les prix restent élevés, les investisseurs bien informés vont perdre les rentabilités qu'ils pouvaient avoir en gardant l'actif.

Question de recherche #4 : Quelles sont les comportements des investisseurs a l'origine de la formation de bulles et des effets stylisés ?

Dans un marché financier nous observons, peut-etre trop souvent, des phenomenes impossibles a expliquer avec les concepts theoriques classiques. Parmi eux, nous regardons de pres les effets des autocorrelations non lineaires des volatilités (ou la longue memoire des marches) et les bulles des prix. Est-ce que ces phenomenes sont des accidents ou ont-ils leurs racines causales dans les différentes strategies des investisseurs ?

On observe que les faits stylisés peuvent etre créés dans un marché avec des différentes proportions d'investisseurs bien informés et d'investisseurs biaisés. Mais ces effets sont aussi visibles quand dans un marché il y a que des investisseurs bien informés, mais avec des delais d'information.

Question de recherche #5 : Est-ce que le SITH peut battre YODA ? Qui va survivre le plus ?

Cette derniere question de recherche ressort naturellement a partir des premiers résultats. Si les marchés financiers restent biaisés pendant de longues durées alors la bonne information économique n'a pas forcément apportée de la rentabilité. Si des bulles spéculatives peuvent etre créées, meme en présence d'investisseurs bien informés, alors ces derniers peuvent-ils en profiter ? Un investisseur bien informé va, normalement, vendre ses actifs risqués des qu'ils sont surévalués. Ainsi, il peut perdre la possibilité de spéculer sur une bulle. Nous avons conçu un investisseur, bien informé, qui peut volontairement changer entre une stratégie fondamentale (appelé YODA) et une stratégie non-informée visant a obtenir le maximum de profit. On appelle cet investisseur SITH car, meme s'il sait qu'un actif est surévalué, il peut meme alimenter un trend de prix croissant. Son seul intérêt est le profit et l'efficience de marché n'est pas une de ses priorités.

Nos résultats montrent que le SITH peut tirer plus de profit que l'investisseur YODA. Cette différence vient du fait que le SITH va prendre, de temps en temps, de l'argent provenant des investisseurs biaisés ou chartistes.

Notre travail nous amene a conclure que les marchés financiers, tels qu'ils sont aujourd'hui, sont des concours de beauté. Cette conclusion est supportée par le fait qu'un investisseur avec l'information parfaite peut gagner moins qu'un autre investisseur, avec la meme information, qu'il va essayer d'imiter d'autres strategies (quand celles-ci vont dominer la marché). Ainsi on voit qu'il n'y a pas toujours d'intérêt a spéculer en faveur d'un prix de marché correct (relativement aux fondamentaux économiques de l'actif sous-jacent). Pour tirer le plus de profit, les investisseurs doivent utiliser et profiter des strategies dominantes.

Nous postulons la théorie de **marchés financiers parfois efficaces**. Selon notre théorie, les marchés financiers se trouvent dans un mouvement continu entre deux états d'équilibre instable : un état d'efficience et un état d'inefficience complete (le prix de marché a alors aucune relation objective avec les fondamentaux d'actif

sous-jacent). Ce mouvement est causé par les investisseurs cherchant des profits supplémentaires et adaptant leurs stratégies (plus ou moins informées en fonction du contexte de marché). Suite à cette théorie, nous proposons comme nouvelle mesure de risque le niveau d'efficacité d'un marché. Ce niveau d'efficacité de marché peut être mesuré, indirectement et avec l'aide des simulateurs de marché, en estimant les proportions relatives des différentes stratégies utilisées dans le marché.

Introduction

Contents

2.1	Context and motivation	11
2.2	Important concepts	14
2.3	Behavior biases and the problems they cause	16
2.4	Behavioral finance vs. efficient markets	18
2.5	Research questions	20
2.6	Research Methodology	21
2.7	Assumptions, advantages and limitations of this study	23
2.8	Rationality and irrationality	25
2.8.1	What is an irrational investor?	26
2.8.2	Are irrational investors present in financial markets?	27
2.9	A very short history of financial markets research	29
2.10	Summary of chapter	30

2.1 Context and motivation

The period 2008-2012 offers a landscape of financial crisis: from debt crisis to bank crisis ending with state debt crisis. These events were encouraged and sustained by political, legal and economic factors but they were directly caused by the inefficient financial markets.

The epic fall of Lehman Brothers bank¹, in August 2007, revealed the systemic fragility of the global financial system. The global banking system suffered important write-downs and lead financial institutions on the brink of bankruptcy. Banks that took too much risk and lost were expected to file for bankruptcy. In most countries, the bankruptcy of their main banks implied a temporary systemic distress to the economy - due to the dependency of companies on the revolving lines of credits offered by the banking sector. 'Too big to fail' banks are saved from bankruptcy with the aid of state capital. State sponsorship has ended the banking

¹Find in this link references to a complete judicial report on the bankruptcy of the Lehman Brothers bank - <http://blogs.wsj.com/deals/2010/03/11/lehman-brothers-heres-a-copy-of-the-court-examiners-report/>

crisis by a subtle transfer of risk: from banks to states. Since 2010 countries have suffered from a sovereign debt crisis that spilled its effects into national economies. Countries like Spain, Italy, Portugal, Ireland and Greece experience high level of unemployment and lack access to credits at affordable prices. The "risk-free" state asset is challenged and its social and economic effects are yet to unfold.

Acknowledging the gravity of the current events, researchers have a duty to understand and prevent the systemic causes that lead to such crisis. The expression 'subprime crisis' is advertised as a suitable culprit for the world's financial and economic events. This denomination refers to the effects of massive 'subprime' debt insolvency. Insolvency is common to every financial transaction and its cost is reflected in the asset prices. For subprime debt, these costs were reflected in losses and not in asset prices. The catalyst for the enormous growth of subprime debt portfolios can be found in US laws designed to enable the American dream to the average American. History offers countless examples of laws, designed with noble intentions and in good faith, which ultimately have predatory and counter-expected effects. The US Community Reinvestment Act provided the legal incentives for subprime debts. The US financial system exploited this new market opportunity. Rapid financial speculation led to a housing bubble and an amplification of subprime debt. Exporting this debt was the next step for the US financial system. Global contagion by subprime debt was enabled by two conditions:

1. The transformation of 'subprime' debt into sellable assets (ABS)
2. The acquisition of these asset based securities by investors from outside the United States, at overestimated prices.

Condition 1) is a national legal matter, called 'financial innovation', and a simple catalyst for contagion. Condition 2) is related to investor behavior and the efficiency of global financial markets. Acquiring any asset, no matter the value of its future cash flows, is a correct investment provided that the price paid justifies the asset's future cash flows. Global investors overpaid US debt securities and suffered devastating losses when these assets showed their real value. These bad investments decision were enabled by two factors:

- A Financial markets did not correctly price ABS securities
- B Institutional investors entered the ABS market and provided liquidity to further enable the existence of condition A.

We believe that condition A, the mispricing of asset backed securities, is the essential factor that triggered the global 2007 financial crisis and all of its consequences. Because of mispricing, international investors were eager to buy US mortgage backed ABS. This demand encouraged US financial institutions to increase their efforts towards the mortgage market. All participants in the US debt market

were interested to provide a house loan and to resell it for a commission. The US financial sector acted as an intermediary between the debt applicants (U.S. families) and the international debt sellers. Even if US financial companies knew these credits would go bad, they continued their actions because: a) US laws permits subprime debt (condition 1) and b) such debt were sold in international markets (condition B). A correct pricing of ABS debt would have stopped the global effects of the crash of the US housing market.

We believe risk mispricing was the fundamental cause of the global 'subprime crisis'. Yet, financial mispricing has different forms: return mispricing, risk mispricing, etc. The 2001 IT bubble, see (Morris and Alam, 2012) for details, which also originated in the United States, can be considered as a clear case of market future cash flow mispricing. Such events have long durations (measurable in years) and it would be hard to argue that they do not invalidate the idea of efficient markets. Due to the significant economic and social impact of these market inefficiencies we have devoted this work towards shedding more light about market dynamics.

In this thesis, we show how and when financial markets misprice assets. Because numerous studies look at aggregate level causes, like macro economy or policies, we focus on the links between individual investor behavior and asset prices. Our study draws its theoretical foundations from the field of behavioral finance. Behavioral finance creates a knowledge model that can explain the empirical anomalies observed in financial markets. As a tool for analysis, we design and use a agent-based financial market simulator. This tool provides the flexibility and freedom to make hypothesis on the individual investor level and to observe the effects at an aggregate level. As indicated by the title of our thesis, we look at when and how investor level biases can affect the efficient formation of market prices. Before we explain, in depth, the objects of our study we will first explain the notion of bias and how it relates to investors. A "bias" is defined as a systematic departure from a reference point. Therefore the concept of bias is relevant to the nature of the reference it points to. We can distinguish a great number of biases like cognitive bias, social bias, behavioral bias, measurement bias, risk bias, etc. We will focus on one in particular, the most relevant to financial research: behavior bias.

In finance, an investor has a behavior bias if his actions are not those of a perfectly rational investor. A perfectly rational investor can perfectly interpret all the available information and, using this information, he can deduce the correct rationally expected value of an asset. Using this correct expected value the investor tries to profit when markets misprice asset. In the chapter 2.3 we explain the main behavior biases of financial investors.

2.2 Important concepts

In the course of this thesis we use financial concepts that have different meanings, depending on the reader's assumptions and context. To mitigate the risk of misinterpretation we present and explain here our understanding of the most important concepts we use.

The **fundamental value** of an asset, or intrinsic value, refers to the present value of an asset. In a usual context, an asset (like a stock, investment project or real-estate) is supposed to provide cash-flows to its owner. In real financial markets, investors, who use 'fundamental analysis', try to estimate an asset's fundamental value using discounted cash flow methods (see a good introduction about fundamental analysis and discounted cash flows in (Graham and Dodd, 1951)). In our study we refer to the fundamental value of an asset as a **theoretical concept** that helps us build our model. This fundamental value helps us create an information flow. When we assume that an asset has a theoretical fundamental value we do not make any assumptions about market efficiency or investor behavior.

Efficiency theory is a highly debated subject between finance academics. In our work we consider the (Fama, 1970) definition of efficiency: "A market in which prices always "fully reflect" available information is called efficient". This definition, as it is often referred to in economic theory, also implies the existence of rational expectations. In our case, where we specify a model of the intrinsic² (or fundamental) value of an asset, we consider a market to be efficient when the market price will follow the fundamental value with a error factor (with zero mean). If not mentioned otherwise, we use the word "efficient" (when describing a market) as a short form for informational efficiency (as defined by (Fama, 1970)).

In this work we refer to **speculative** investment behavior. We consider as speculative any investment strategy that implies actions which go against the available information. For example, consider a investor with information showing an asset is overpriced. If this investor buys the asset we refer to his action as speculative. We do not imply that such actions are not rational or not profit-driven. The inverse of a speculative action will be, in this context, an arbitration action. We refer to a investor as being **fundamentalist**, or as having an arbitration strategy, if he acts in accordance to the information he holds. In the example before, a fundamentalist investor will sell an asset which he considers, from his information, to be overvalued.

The concepts of **risk premium** and **risk estimation** are clearly defined in the financial literature, see (Valdez, 2007). Risk premium is a subjective weight that an investor places on his expected returns from a risky asset. A risk estimation is an investor's estimation of an asset's risk. To simply our model we propose the concept of **minimal expected return**, r_{min} , which is an investor subjective value.

²This fundamental value is not known directly by investors. Hints of the fundamental value are revealed through a biased information process.

A minimal expected return of an investor combines his risk premium as well as his estimation of an asset's risk. When considering an investment action, an investor will use his minimal expected return to estimate a good price for the investment.

A **discount rate** is used to compute the present value of future payments. In financial theory, there are a number of methods to estimate the discount rate of (expected) future dividends of equity, such as the CAPM (Sharpe, 1964). We consider that a model where investors have to estimate both future dividends and discount rates is not parsimonious. Therefore, we propose an information process from which investors can directly compute **expectations of future discounted dividends**. In this way, we avoid the problems that are inherent to methods for separately estimating dividends and discount rates. Because our informational process is biased (due to investor subjective characteristics) our approach does not necessarily imply investors have the same estimations of discount rates nor of future dividends.

2.3 Behavior biases and the problems they cause

The main purpose of financial markets is the efficient allocation of capital. Investors use financial markets as guides to invest their capital in the most profitable ventures. Referring to stock markets, the fundamental value of a stock, in theory, is considered to be the discounted sum of the cash that will flow to investors from the stock. A stock's theoretical (fundamental or intrinsic) value can be expressed as:

$$F_t = \sum_{i=1}^n \frac{CF_{t+i}}{(1+r_{t+i})^{t+i}} \quad (2.1)$$

where CF is a future cash flow (either dividend or resell price). An investor has to know the value of all future cash-flows and appropriate rates of return to compute this fundamental value of the stock. In reality, investors make estimations of what future cash-flows and rates of return will be. For simplicity (but without restraining the breadth of our results) let us suppose that the rate of return is constant in all time period,

$$r_i = r, \forall i \quad (2.2)$$

In this context, an investor looks only for the information needed to estimate the future cash flows from the stock, as in the formula 2.3.

$$E_{investor}(F_t) = \sum_{i=1}^n \frac{E_{investor}(CF_{t+i}, Information)}{(1+r)^{t+i}} \quad (2.3)$$

An investor's ability to correctly compute future cash flows is limited only by his intelligence and understanding of the factors that govern the performance of a company, assuming equal access to information. Naturally, investors have different levels of financial knowledge. Therefore their estimations of a stock's fundamental value are different. Using their value estimations, investors come to the markets and try to sell or buy for a profit. Irrelevant of the market price formation mechanism it is correct to say the market prices are a function of the investors' combined expectations of the stock's value (see equation below where f_{pm} stands for price formation mechanism and m is the number of investors who trade that stock at moment t)

$$P_t = f_{pm}(E_1(F_t), E_2(F_t), \dots, E_m(F_t)) \quad (2.4)$$

Equation 2.4 and its resolution is actually the main debate point in modern finance. According to the market efficiency theory (Fama, 1970) and the rational expectations theory (Muth, 1961) a market will quickly and correctly integrate information into prices so they accurately reflect fundamental values. This line of theory does not imply that each investor has a good estimation of value, rather that the market can somehow transform all of the individual estimations into a good, unbiased, rational expectation.

$$P_t = F_t + \varepsilon_t \quad E(\varepsilon_t) = 0 \quad (2.5)$$

Contrary to the efficient markets theoretical foundation, behavioral finance asserts that equation 2.5 is not true and aggregate prices can be persistently far from fundamental values. The basic assertion of behavioral finance states that investors are not perfectly rationality, and this fact can explain markets inefficiencies. We refer you to (De Bondt and Thaler, 1987) or (Coval and Shumway, 2005) for examples and discussions about behavioral traits affecting market prices.

In our theoretical context, were we suppose the apriori existence of fundamental market ³, we can consider that behavioral finance states that market prices can be persistently biased (in relation to fundamental values). This idea can be expressed using the equations 2.6 (asset prices are not at all related to their fundamental value) and 2.7 (asset prices are persistently biased in regards to their fundamental value) :

$$P_t \neq F_t + \varepsilon_t \quad (2.6)$$

$$P_t = F_t + \varepsilon_t \quad E(\varepsilon_t) \neq 0 \quad (2.7)$$

Looking again at equation 2.4 we see the market price is formed, through a deductive procedure (called price discovery method), by combining the expectations of all the trading agents. Therefore two natural questions come to mind:

1. Can (biased) investor expectations create a biased price?
2. If 1) is true, in what conditions do these biases persist?

To describe the importance of price biases we have to think from the perspective of the client of a financial market: **the investor**. When using the services of a financial market, most investors expect to buy an asset at (at most) it's fair (fundamental) value. Moreover, the investor expects the market price will quickly follow when the asset's fundamental value goes up or down. In a financial market with biased prices, the investors bare the risk of overpaying for assets and/or selling assets at less than their fundamental value. For some investors, such risks are unacceptable since they only trade to preserve their capital. Consequently biased financial markets exhibit decrease volumes or no trading at all. Paradoxically, some investors strive to enter markets when assets are undervalued and exit when assets become correctly valued (or even better, overvalued). Ideally, some investors would prefer markets with two cyclical stages: a biased stage (where assets are under or overvalued) and an efficient stage (where prices would reflect fundamental values). In such a market, investors make returns by profiting from biased prices and closing

³Please refer to section 2.2 for a clear explanation of what we assume by a fundamental value.

their positions when assets are correctly priced. In markets with consistently biased prices a rational investor cannot make excessive profits using his insights into economic fundamentals.

2.4 Behavioral finance vs. efficient markets

According to efficient market theory, any pricing bias poses no problems to the market and its effect is only temporary. The theory states that smart and adaptable investors quickly detect the new profit opportunity, created by the pricing bias, and put in place a strategy that will permit to profit out of this inefficiency. When enough money is invested in the new strategy the new profit possibilities disappear. Through this process of rapid speculation, the market biases quickly disappear. This theoretical assertion is true but incomplete. The efficient markets theory ignores the way in which efficiency is actually restored. The theory implicitly assumes there are investors with excellent information which quickly reveals market inefficiencies. This assumption also implies the knowledgeable investors consistently monitor markets and look to exploit any inefficiency. In fact, they would not invest their capital as long as they don't detect any market inefficiency. While not investing, such investors would always incur the costs for their excellent information. So, these investors would require every new profit opportunity (arising from market inefficiencies) to provide enough gains to cover: the cost of the strategy, the cost of all the previous information bought and a sufficient risk premium for engaging their capital against the market's blunder. Therefore, these investors never engage their capital to correct small market inefficiencies since such actions would not be profitable. Instead, they wait until the market develops a profitable enough inefficiency (depending on the three factors mentioned above) and only then do they take action. Such rational behavior, from the part of well informed investors, can generate (by inaction) the appearance of significant market mispricing.

In their famous article, '*On the Impossibility of Informationally Efficient Markets*', (Grossman and Stiglitz, 1980) provide a theoretical basis for this paradox and conclude by saying that "*because information is costly, prices cannot perfectly reflect the information which is available, since if it did, those who spent resources to obtain it would receive no compensation*". We add to the conclusions of (Grossman and Stiglitz, 1980) and say that an investor has the possibility of using his capital to either go against an incorrect (relative to the investor's personal beliefs) market trend or to speculate and go with the market (even when his information says the asset is mispriced). In both cases the investor makes a gamble: either the market "wakes up" and he profits from the price or the market continues in the wrong direction and the investor profits from going with the trend. In this view, we assume that there are situations when knowledgeable investors decide to follow the 'incorrect' market trend. This '*informed speculative*' behavior amplifies the momentum for the trend, which increase even more the inefficiency. When this happens,

prices become disconnected from their fundamental values and knowledgeable investors start a classic game of cat and mouse. Investors with open positions look for signs that other investors will turn against the market by revealing information that encourages others to close positions and speculate on the correction of prices. Investment decision are not driven anymore by economic information (strictly related to the asset's real value) but rather by indicators of how many investors are willing to ride (or continue to) the market trend. This idea is not new since it resembles a Keynesian beauty contest where investors try to guess each others intentions. When the market moves away from an efficient state, into a state where the market is controlled by ad-hoc beliefs and momentum actions, prices are biased and can easily develop into bubbles. In these situations, risk can greatly increase since both informed and non-informed investors will have higher expectations.

Some natural questions arise:

- How often can these mispricing events happen?
- How long does it take to return to efficiency?
- Can investors profit from these inefficiencies by not acting against them?

These are some of the questions we answer with this thesis.

Aside from the academic interest, the answers to these questions have an important regulatory impact. By understanding the systemic causes of market inefficiencies, financial markets organizers and regulators can impose new operating rules and regulations that limit or stop these biases. Identifying such factors provides new ways to:

1. Accurately communicate to markets the possibility of market inefficiencies (better information)
2. Develop new risk measures to asses potential losses (especially especially for investors who cannot support much risk)
3. Change the view of "efficient markets" into "*sometimes efficient markets*". This can help institutional investors make better risk assessments on behalf of their clients.
4. Develop better ways to discover and signal market inefficiencies (improve the attractiveness and future efficiency of markets).

As argued before, we observe through-out our study that market "inefficiencies" are profitable for some types of investors. Therefore it is reasonable to assume that it is in the interest of such investors to preserve financial markets in their actual state: *sometimes efficient* and occasionally *very inefficient*.

2.5 Research questions

We explained that the most important characteristic of a financial market is its ability to correctly price assets. Mispricing of assets is due to a number of factors like: lack of liquidity, lack of information or investor behavior. In regards to the first two factors (liquidity and information) the recent ten years have provided financial markets with technological tools (completely electronic markets, algorithmic trading, etc) that enable, as never seen before, the fast and reliable access to information and liquidity. In this study we focus on the relationship between individual investor behavior. We assume access to liquidity and information is constant and equal for all investors.

In our study of the relationships between investor behavior and market inefficiencies we will be answering five research questions.

Research Question #1: *What mix of biased investor behaviors produces consistently biased market prices?*

By answering this question we show how particular biases can affect aggregate market prices. It is intuitive to say, for example, that if all investors are optimists than the market price will be higher than its fundamental value. Therefore we also show if it is realistic that such a mix of biases can actually appear in a financial market. Because most behavior biases are common to all investors we seek those biased investors who survive in financial markets.

Research Question #2: *Can a biased investor(s) survive in a financial market?*

If the answer to question 2 is yes than we have proven that biased behavior can persist in a financial market. We are still left with the issue of how biased behavior can appear in markets. One can argue investors undergo extensive preparations, before entering financial markets, to discover and eliminate their biases. Moreover we can suppose that all investors constantly look to eliminate their own biases. We can also assume investors constantly adjust their interpretations of information in order to increase their profits. In view of such pure profit objective, one can argue an investor can voluntarily choose to bias his beliefs or actions if he can make a bigger profit. In general, stock markets provide an approximately 7% return per year, on the long term (in spite of years with huge financial crises). Yet stock prices sometimes rise very fast and very high. Such moments offer profit opportunities much higher than 7%. Therefore, on a pure profit basis, investors have incentives to 'ride' and cash in on abnormal, short-term, market opportunities. To earn such returns, not explained by the economic fundamentals of the traded assets, investors often use strategies not solely based on fundamentals (hence the existence of technical trading).

Research Question #3: *Can biased investors earn more than non-biased investors? If yes, in what conditions?*

If a biased investor survives and earns good profits in a financial market, we expect other rational participants to mimic such profitable behaviors. The more biased participants enter a market the more the market price will be disconnected from its fundamental information flow. But all investors want to make profits so they expect prices to rise. Lacking (or not using) the possibilities of 'economic reality checks' prices can increase without any economic backing. Increased prices can ultimately lead to moments where no 'bigger fool' exists. In such cases, trading between biased investors can end and market prices will start being formed by other groups of investors (usually rational non-biased investors).

Research Question #4: What are the market micro conditions for the emergence of price bubbles?

Lastly, we study the market impact of two special kinds of investors: Yoda and Sith. Both investors have amazing capacities for analysing economic information and for computing the correct fundamental value of any asset. Yoda tries to profit from investment actions that close the gap between the asset's market price and fundamental value. In opposition, Sith looks for profit opportunities without regarding the relation between prices and fundamental value. In this work we show which of these two types of investors, Yoda or Sith, survives in a competitive financial market. If the Sith investor survives then we can reinforce the idea that market prices can be disconnected from economic realities. Research question number 1 seeks the conditions, at an aggregate level, for the disconnection between prices and fundamentals. Research question number 5 looks for profitable individual behaviors that can generate biased market prices.

Research Question #5: Can Yoda beat Sith? How can live longer, Yoda or Sith?

Answering these five research questions helps us understand how different investor biases can have an impact on market prices and indirectly on the efficient allocation of capital in our economies.

2.6 Research Methodology

Our research goal is to explain the links between aggregate financial phenomena the behavior of individual investors. To answer the research questions we will be using a methodology that called a 'generative approach' (Epstein, 2006). We first define a hypothesis asserting a connection between a specification of individual investor behavior and an aggregate financial observation. The hypotheses we make all follow the same general structure:

H: Suppose that conditions $X_1 \dots X_n$ are respected (at the investor level of a financial market). Then conditions $Y_1 \dots Y_m$ will emerge at the aggregate level of that financial market.

For example, we look at the financial 'puzzle' called "volatility clustering". In this case the analysed aggregate propriety (Y) is the volatility of returns. To determine the behavioral causes/correlations of volatility clustering we test hypothesis such as:

H1: In a market with 30% optimistic investors (X_1) and 70% rational investors (X_2) we will observe clustering of volatility (Y_1).

Solving this hypothesis will immediately show a correlation effect between our premises and the desired effects. Because our analysis environment is deterministic and completely observable we can next make claims about causality.

To test such a hypothesis we use a tool named 'agent-based simulator'. An agent based simulator is a computer program that simulates the activity of all the components of financial market: the evolution of information, investors who interpret information and send orders, the clearing house and the mechanism of market prices discovery. Starting for a complete micro level specification of a financial market (assumptions $X_1 \dots X_n$) the simulator computes, in a deductive manner, the time series of relevant macroscopic variables $Y_1 \dots Y_m$ (like prices, transaction volumes, returns, etc).

We underline that the transformation of the input data X_i into output data Y_i is a deterministic process. An agent based simulator is equivalent to a mathematical function like: $f(X_1 \dots X_n) = (Y_1 \dots Y_m)$

Because a financial market is a highly complex system it is extremely hard to create a mathematical model capable of describing it completely. For this reason, instead of a mathematical language we use computer program language. This language has the ability to describe complex system in a natural and intuitive fashion. We propose a simple problem that highlights the advantages of an agent-based simulator:

- *Given a tree leaf (like in the figure below) please discover what species of tree it belongs to.*



Figure 2.1: Global view of methodology

Using a classical approach, a research would analyze the characteristics of the leaf and find patterns similar to other known-origin leaves. This approach would identify leaves based on correlations and not causality: if two leaves look the same they belong to the same type of tree. Using an agent-based simulation approach, we search for different tree seeds and grow entire trees from their seeds. When the grown trees are mature we observe directly their leaves and compare them to our unknown leaf. In this way we show a causal link between a tree seed and a leaf. In a financial market the 'leaves' represent unexplained phenomena (like the ones presented in chapter) and the seeds are the investor level specifications of a financial market (like the DNA strings of a tree seed).

2.7 Assumptions, advantages and limitations of this study

This chapter explains the main assumptions, advantages and limitations of our study. Each of the assumptions is explained in more detail in the next chapters.

Main assumptions of our study:

- Market prices are only influenced by the orders of investors. Perfect information exists and it can be discovered by investors. Implicitly, we assume that all other markets are efficient (bond or derivatives markets) and do not affect the assets in the market we study.
- The price formation mechanism does not produce biased prices. We acknowledge that different market design features, like tick size or order life-time, can

impact the dissipation of information into prices, and can have some unwanted influences on market price (as reported by (Chiarella and Iori, 2002)).

- Investors, as a group, learn passively through an evolutionary process: bad strategies lose money while good strategies earn money.
- A theoretical fundamental economic value for a company exists. We model and use such a value to generate a perfect information series.
- The only information describing our traded asset is based on the asset's fundamental value. Other events like market shocks, rumors, or macroeconomic evolutions are all included in this fundamental value calculation (which is assumed perfect).
- Investors do not separately estimate discount rates and future dividends. They directly compute expected future discounted dividends. This assumption does not imply investors use the same discount rates or expected future dividends.
- We simulate and analyze a financial market with a single stock and no other interlinked markets (options market, bond market, etc) because we assume the following:
 1. Investors are not affected by events from other markets
 2. Investors focus their decisions only on the simulated stock
 3. Investors have equal access to information and liquidity

One of the main advantages of our study is that we propose financial theories that are sustained by a deductive proof. This provides the scientific framework for clear and precise causal explanations.

Working with a tool, agent based simulator, that is completely controllable we can also determine correlations between different market factors. In a classical approach we assume a condition of 'ceteris paribus' but it is hard to verify such an assumption. With a simulation approach we have the choice to stop or allow the variation of any state variable. Thus we can make observations and inferences about variable correlations using a *veritas ceteris paribus* method.

When testing a theory, our method proves more than the 'truth' property of a theorem. A classic financial research method implicitly assumes that if a phenomenon can be generated by certain conditions than we can also observe this phenomenon in real markets. For example, the Capital Asset Pricing Model provides the theoretical basis for measuring risk and estimating returns. In theory the CAPM is true yet empirical tests have shown that its consequences are not observed in real markets. Thus, one of the reasons why CAPM is not validated in reality is because the financial market states it describes are not attainable in reality.

We believe that the most important advantage of our study is the **possibility of proving the attainability of true theories**. Moreover, for attainability to be

relevant in a financial system a market state has to be achieved in a reasonable time frame. If a financial market integrates new information in a matter of years than the theoretical true state of 'efficiency' is not attainable in market. This does not necessarily imply that the market is never efficient, rather markets have a quality of 'noisy efficiency' as described by Prof. Herve Alexandre in this doctoral thesis (Herve, 1994).

Our study also has a few disadvantages. From a scientific point of view there are two problematic issues with the methodology we use: the accuracy of our basic assumptions and the validity of our agent based simulator. Most of our assumptions are references to patterns of investor behavior as well to the relative proportions of these behaviors. The investor behavior characteristics we model have a scientific validity in psychological studies. Yet, our models of these behaviors are an idealistic reflection of the way real investors make decisions. Consequently, if some of our assumptions are flawed our results can also be false.

Agent based simulators have been used successfully to solve financial problems. We mention its use by the American NASDAQ market in 1997. The management of NASDAQ wanted to reduce the market tick size from 1/8 to 1/16. Intuitively, it seems correct that if prices can move smoother than market quality can be improved. To avoid negative market reactions, NASDAQ commissioned the creation of an agent-based simulation of their market where this regulation could be tested. The results were surprisingly counter-intuitive. The simulation showed that a decrease in tick size would hinder the market's ability to discover prices and would inevitably increase bid-ask spreads. The NASDAQ board decided not to decrease the tick size, probably saving a lot of money in the process.

2.8 Rationality and irrationality

Throughout this work we refer to rational and irrational behavior. The purpose of this chapter is to offer insights for the answers to these two questions:

1. What is an irrational investor?
2. Are irrational investors present in financial markets?

It is now well accepted that some traces of possible 'irrationality' can be seen in financial markets and that there are more and more studies on this topic. It is important to notice that even though we can see an increasing volume of recent literature that focuses on 'biases' and 'irrationality' there is not clear and common definition of what actually irrationality means and also to whom it refers. In (Rubinstein, 2001) the author argues that a, vague, definition of a rational investor is of an investor who follows the axioms of (Savage, 1954). This view, actually shared by most researchers, implies that a rational investor tries to maximize his expected

utility of wealth using unbiased subjective probabilities. Because the key element of this view on rationality is that of unbiased subjective probabilities, irrationality is immediately thought of as the behavior that maximizes expected utility using biased subjective probabilities. This line of thought, pioneered by the works (Kahneman, 1973) and (Kahneman and Tversky, 1984), has created a new field in finance called behavioral finance. This theoretical research field focuses on the discovery and analysis of biases of the subjective probabilities and also on the deviations from the assumption of maximization of expected utility. Due to these two ways, biases and the deviations from expected utility maximization, of measuring and discovering 'investor irrationality' we infer that rationality can be viewed in the relation to the objectives of the investor but also in relation to the means of achieving these objectives.

2.8.1 What is an irrational investor?

A rational investor can be defined as an investor that takes actions consistent with achieving his immediate objective. In financial literature, it is often implied that the objective of the investor is the maximization of the expected utility of wealth and rationality is measured in relation to the actions that are done to reach this goal. This important limitation, of a single financial objective, can be the cause of many false-positive results in judging the rationality of investors. As there are professional athletes, who do sports for a living, and amateur athletes, who do sports for pleasure, there are professional investors, with clear monetary objectives, and also novice investors, that are learning to invest or trade just for the thrill of seeing the up's in the market. From this point of view, we argue that it is not easy to find an investor that is irrational in relation to his objective. Let us assume that an investor, with no financial education, buys stocks. His investment decision is based on the price charts (or names) and the purchased quantities depend on the stocks' price quotes. In relation to the objective of maximizing expected utility of wealth, such an investor is considered irrational. When asked, most of these kinds of investors say that have never heard about expected utility and probably never used this measure an objective. Most likely these investors wanted to learn about financial markets and maybe gain some wealth. In view of this objective, we say that this type of investor is rational because he is doing actions that help him learn more about markets and that can sometime offer a return. In extremis, we can say that someone that commits suicide is rational in relation with his objective (of not living anymore). With this in mind, it is hard to find someone that is irrational in relation with his personal objectives. In financial markets, there are investors that want to make steady returns, others that want not to lose money, a growing number of investors that want to be ethical or in line with their religious beliefs, some (if not many) that want to get very rich (as quick as possible) and still many other investors with different objectives. Some of these objectives cannot even be quantified and measured and can be hard to rank in relation to the normative optimal objective of

utility maximization. From this point of view, the vast majority of investors make rational actions in relation to their own subjective objectives. In an attempt to explain and justify the dynamics of markets one should keep in mind that aggregate behavior can sometimes be better explained by a mix of investors that have different objectives (among which utility maximization can be found) rather than by investors doing unpredictable irrational actions whilst all having the same single normative objective.

We argued that an investor can be regarded as rational in relation to his personal objective, which can be different than the normative objective. Therefore, the action of taking on more risk can be rational for an investor that has the objective of making a lot of money very fast (while disregarding risk). If an investor's actions are in the same direction with his objective we say that the investor is rational in relation to his objective. It is intuitive that investors don't always do the optimal actions in relation to their goals. For example, an investor that wants to maintain the purchasing power of his wealth should invest in a bond or other security that guarantees a return equal to at least the annual inflation rate. Accordingly, two investors that buy bonds with a 5% return and respectively a 3% return are rational objective-wise yet what can we say about their actions-wise rationality? Provided that the investors' country has shown a very stable inflation rate at around 2% per year, we say that both investors are action-wise rational. If we compare these two investors with another one that buys a government bond indexed on the official inflation-rate, what can we say about the rationality of their actions? In our view, we can say that all three investors are action-wise rational, though arguments can be made about which one is better - which action is more "rational".

We have defined investor rationality as being related to objectives - objective-wise rationality, and also rationality when related to actions specific to an objective - action-wise rationality. One can see that an objective-wise rational investor will make actions that can be ranked as being 'more' or 'less' close to the optimal objective action.

2.8.2 Are irrational investors present in financial markets?

The efficient markets theory (EMT), first introduced by (Fama, 1970), says that in an efficient market all new information is immediately and correctly integrated into price. In other words, the price is correct and an average investor cannot make abnormal profits (more than the evolution of the market price). This theory does not imply that markets are populated by some specific types of investors, but instead implies that efficient markets can behave 'rationally' as a whole. When referring to individual investor rationality, the EMT defenders believe that irrational investors cannot survive in a market since, as pointed out initially in (Friedman, 1953), they will exhaust their wealth because of their bad decisions, especially by buying at high prices and selling at low prices. Intuitively this assumption makes immediate sense:

repeatedly buying and selling at a loss will exhaust one's wealth. At a second look, we say that it is not very probable to see an investor that will always buy and sell at a loss. It is sometimes likely that after buying an asset, even if it is overvalued at the moment of buying, its price will go up even more. A rising price gives an irrational investor the opportunity of closing his position with a profit - that could help decrease possible earlier losses. It is not likely that all irrational investors would not close out a profitable position or that an overvalued asset would not become even more overvalued. Indeed, there are numerous articles that build and test models with irrational investors which preserve their wealth and can sometimes control markets. In (De Long et al., 1991) we see the results of a model of portfolio allocation by noise traders with incorrect expectations about return variances. The authors show that these irrational action-wise traders (in this case called 'noise' traders) can earn higher expected returns than rational investors with similar risk aversion and they can even arrive to dominate the market.

Using an agent-based model of investors' decisions, (Evans D. and Nettle, 2003) argues that a biased behavior can actually be an evolutionary design feature rather than a flaw of human decision-making. This study shows that under certain conditions unbiased agents using classical expected utility maximization will be outperformed by agents with human biases that have been documented in empirical studies. The idea of 'irrational' agents surviving in a competitive market is counter-intuitive and it immediately raises critics. We think that most of the critic for the survival of irrationality is due to confusion in terms. We look first at the origins of evolutionary rational markets and than explain why these concept don't apply in real financial markets.

Classical finance 'rationality' has been defined as describing the actions of an investor in relation to the optimal actions one should take whilst trying to achieve the goal of maximizing the expected utility of his wealth. The pioneering work of (Muth, 1961), (Radner, 1972), (Lucas, Robert E, 1978) have put the notion of rational expectations in the center of academic finance. In order for investors to portray and exploit rational expectations they must be able to discover the true distributions of variables in the economy - otherwise, as pointed by (Gilboa et al., 2009) there would not much rationality behind using a Bayesian rule in making a decision (footnote - better not decide at all). To defend this idea, of investors knowing (or getting to know) the true probabilities of economic variables, (Alchian, 1950), (Enke, 1951), and (Friedman, 1953) have introduced and argued the idea of natural selection (or the market selection hypothesis). In a similar way, (Cootner, 1964) and (Fama, 1965) argued that in financial markets, investors with 'bad' belief (or with irrational expectations) will be driven out of markets (they will loose their money and will not be able to get new ones) by investors with rational expectations.

A number of researchers have showed that this hypothesis of natural selection (or market selection) does not always hold. (De Long et al., 1991) has showed that using plausible biases, noise traders can survive as a group, earn higher returns

than rational investors and even be able to dominate the market, in terms of wealth, in the long run. Similarly, (Blume and Easley, 2001) argues that in an incomplete market **irrationality** can survive. This result, as argued by the authors is important because market incompleteness is natural.

As we pointed out in the first part of the article, investors can have different objectives in markets and their actions can be rational in relation to their personal objectives and not in relation to a normative objective. So, investors with apparently 'irrational' behaviors persist in markets even if their goals are not those of utility maximization or their actions are not optimal relative to the normative objective. In real financial markets, professional managers and funds are benchmarked on market indexes or relative to other similar managers or funds. Therefore, money controlled by 'less-performing' managers flow to 'more-performing' managers. When particular assets are overpriced, some funds can make excess returns and attract more capital which is lost by funds that may have had the normative objective. Moreover, in (Yan, 2008) it is argued that even when investors or funds have the normative objective accompanied by sub-optimal actions, their evolutionary disappearance from markets will take a lot of time - long enough for new generations of similar investors to enter the markets. Having argued that some forms of irrationality can persist in markets, we will next discuss about modelling this assumptions into asset pricing studies.

2.9 A very short history of financial markets research

Our study uses the tools of computational finance and is based on the theory of behavioral finance. This chapter presents a short history of financial market research and highlights the place of our work between the modern academic financial theories.

One of the earliest modern financial markets was created in Antwerp, Belgium in 1531. Individual lenders and brokers meet there to exchange government, company and also individual debt notes. Since then financial markets have constantly grown and evolved. The total value of the world's financial stock, comprising equity market capitalization and outstanding bonds and loans have increased from *"\$175 trillion in 2008 to \$212 trillion at the end of 2010, surpassing the previous 2007 peak"* according to a report by the McKinsey Global Institute from August 2011. If we take into account the fact that the world's GDP in 2010 was \$74.54 trillion we can understand the vital importance of these markets. Because of their magnitude in economic activity, financial markets have become a critical infrastructure of our society. Consequently, financial markets are and have been for a long time the object of academic research.

Because of the nature of academic research, which focuses mainly on past observations, financial research has always looked to analyse, interpret and understand financial developments. Before 1950, most of the work done in financial research was descriptive. There are some notable researchers, before 1950, which engaged in more

axiomatic approaches to research in finance. Among these we can mention (Fisher, 1925) which explains a theory of interest rates and of the internal rate of return, (Williams, 1938) that modelled theoretical stock prices as functions of discounted future cash payments (dividends) or (Bachelier, 1900) that published in this doctoral thesis a theory on random stock price movements and also set the framework of a mathematical abstraction called the Wiener Process.

After 1950 the neoclassical era of finance began. The beginning of this period was marked by the works of (Markowitz, 1952) on Portfolio Theory continued by (Tobin, 1956) with his theoretical extensions on the efficient frontier and the capital market line. Although revolutionary at their time, these theories proposed methods based on very strict hypothesis that are hard to be meet in financial markets. Other researchers based their work on a new founding theory in finance called the Capital Asset Pricing Method. The CAPM, created separately by (Sharpe, 1964), (Lintner, 1965) and (Mossin, 1966), offers a method for investors to value securities. Even today the CAPM is debated and numerous studies offer proof for or against the validity of this theory. Because of the non-axiomatic quality of social sciences in general, and finance in particular, a theory can be argued both ways without anyone being able to declare a clear winning argument. Even so, the CAPM has at its core a set of strong assumptions like: all investors are rational and risk-averse and are broadly diversified across a range of investments. These assumptions, and especially the one implying the rationality of financial investors, have been recently scrutinized by a series of studies in the fields of psychology and sociology. Since the '70 years, researchers like Daniel Kahneman and Amos Tversky have been researching the cognitive mechanisms and the errors made by humans in the process of risky decision making. In view of the Nobel Prize for Economy awarded in 2002 for (Kahneman and Tversky, 1979) work on "*Prospect Theory*", we see that the academic community acknowledged the inconsistencies of some of the modern finance founding theories' assumptions with the realities of human understanding and behavior in risk-related situations. From these discoveries a new field in finance research was created, namely behavioral finance. For a complete literature review on behavioral finance we recommend the writings of (Hirshleifer and Teoh, 2003), (Barberis and Thaler, 2003), (Subrahmanyam, 2008).

2.10 Summary of chapter

In this chapter we presented the social and financial context that motivates our study. As indicated by our research questions, we explore the causal links between investor behavior and the aggregate series of financial markets. We explained the methodology we use, detailing its assumptions, advantages and disadvantages. We end this first chapter with a short description of objective and action-wise rationality and a short history of financial research. The organisation of the rest of the document is as follows.

In Chapter 2 we explain the characteristics of the object we study: how financial markets form prices and how investors construct their expectations and send orders. We continue with empirical proof of the financial 'puzzles' we seek to explain. In the last part of the chapter we discuss the theoretical foundations of our research.

In Chapter 3 we give in-depth details about what is an agent based financial market simulator and we review the relevant literature. Chapter 4 presents the simulator we built and discusses the methodology used for testing and calibration.

Chapter 5 presents simulation runs which provide answers to the research questions. Chapter 7 presents a case-study application of our simulator as a tool for research in finance and economics. Chapter 8 draws the conclusions, restates our findings and offers hints to future work.

Object of study

Contents

3.1	Information and market prices	34
3.2	Asymmetry and quality of information	36
3.3	Investors	38
3.4	Behavior biases	39
3.5	Price formation mechanism	41
3.6	Price anomalies and explanations	43
3.7	Classical approaches to studying market dynamics	47
3.8	Summary	48

In this chapter we explain the different proprieties that describe the system we study: the modern financial market. We view a financial market like a system composed of different types of elements: information sources, investors, orders, price formation mechanisms and transmission media. Looking only at price formation, we can view financial markets as a heterogeneous system that aggregates different sources of information, via human interpreters (the investors), into an output called the 'market price'.

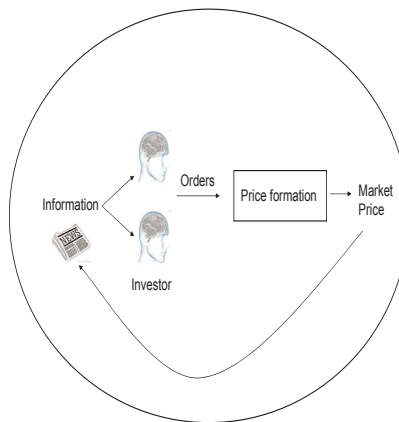


Figure 3.1: Global view of methodology

In this work we focus on a single part of this system: the investor. Therefore, we have made assumptions about how other parts of a financial market works. The following paragraphs describe our assumptions. These descriptions constitute the basis of a theoretical model of investor behavior that influences price formation.

3.1 Information and market prices

The basic ingredient that fuels a financial market is information. We consider only the information regarding one particular asset from a financial market (ex: a stock). Our model can be extended to support multiple assets and information sources.

A company's stock fundamental value is computed from the actualised future cash flows that will be released to stockholders from the company's assets (see equation 2.1). These future cash flows are estimated, by investors, using information related to the different aspects of the company:

- *Economic information*: what does the company produce, at what costs (fixed and/or variable)
- *Market information*: what will be the future market demand for the company's products
- *Financial information*: how is the company financed, is it in debt, what is the probability it will go bankrupt if something unexpected happens
- *Management information*: how do managers use the company's assets? Are they doing their best to maximize the company's net profits? Are they selling their private stocks (insider trading)?

Knowledge of all these types of information is used to make estimations of future cash flows. The sheer volume of information, necessary to understand one company, is impressive and it is not easy for an investor to understand and keep up with all the relevant details needed for correct estimations. Therefore, some investors rely on alternative sources of 'higher-level' information like, see (Abreu and Mendes, 2011) or (Bebczuk, 2003) for reviews:

- Analyst's reports and recommendations
- Peer opinions
- Technical information

Financial analysts, as reported by (Barry and Jennings, 1992), with solid financial knowledge, provide regular estimates of companies' stock values. Their job is to

digest and interpret information and to write investment reports and recommendations. Investors often use these recommendations to make their investment decisions and update their portfolios. The accuracy of analyst recommendations can be misleading and is a catalyst for the distortion of market prices. Numerous research studies focused on the influence of such recommendations and they showed that analyst opinions can have good as well as bad effects on market efficiency. For a good review of research related to financial analysts' recommendations we suggested reading (Ramnath et al., 2006).

As an alternative to analyst recommendations, investors use information provided by their peers, see (Chang et al., 2000). In financial research literature this behavior, called 'mimetic', implies that an investor ignores other sources of information and bases his decisions on the interpretation of another investor. Such behavior is justified by the fact that some investors (ex: portfolio managers that disclose their results) show consistently good performance. It seems natural that investors seek to imitate the decisions of constant winners (and hope to achieve the same results). This behavior can lead, in extremis, to prices which reflect the opinion of a single group of investors and therefore are disconnected from the underlying economic fundamentals. This was the case with the 2001 IT bubble, when investors and analysts were bidding up prices for technology start-ups when those business did not respect certain basic accounting and audit requirements. Some analysts and investors defended their views by quoting theories on how to make 'new' evaluations that justified the high prices. In the case of the 'dot-com' bubble, financial markets did not manage to correctly integrate all the available information into the market prices. Instead, as analyzed by (Sharma, 2009), investors focused on a particular source of information that gave a short-term profitable evaluation of stocks but a wrong evaluation of long-term value.

Even when information of quality is abundant, investors can still (and do) choose to ignore information about fundamentals and engage in 'technical analysis'. 'Technical analysis', as thoroughly explained by (Lo et al., 2000), consists of a set of methods that help investors compute an expected value of the future stock price based on historical price information.

The underlying assumption of market efficiency is that prices reflect perfectly and completely all economic information. Thus, looking at fundamentals is not useful for making excess profits. Technical traders do not believe markets are efficient and they focus on discovering trends and patterns in prices. Using statistical methods they try to make profits out of repeating patterns. If a 'trend' or pattern is detected by a sufficient number of technicians then, through sharing of information or independently, there is a possibility that their combined actions will actually create a trend in price - realizing what is called a 'self-fulfilling prophecy'. (Fernandez et al., 2009) and (Jordan, 2006) provide good accounts of such prophecies.

3.2 Asymmetry and quality of information

From the previous examples we saw that investors are faced with a large volume of information. If we assume that investors can, on average, extract the relevant information we should also consider the possibility that information can sometimes lack in quality or precision.

One of the theories standing behind the idea that information can be correctly inferred from a variety of sources is the central limit theory, see (Dudley, 1999) for an exhaustive description of central limit fundamentals and applications. To see how this theory works with information we propose the following example: Imagine a hundred people that can observe a transparent jar filled with an unknown number of plastic balls. Each individual is asked to look at the jar (without opening it) and give an estimation of the number of balls that are inside.




Figure 3.2: Estimation of correct information with central limit theory

The observers estimate the number of balls using different approaches: counting, mathematical analysis using volumes, guesses, thumb rules or other methods. Therefore these individuals estimations are biased. The idea of the central limit theory is that if we have a sufficient number of independent and identically distributed estimations, implying that people do not exchange ideas, than the distribution of the **average estimation** will be normally distributed and it's mean will be the true number of balls in the jar (N). In other words, according to the central limit theory, we expect to have a similar number of people that will undervalue N as people that overvalue N . Moreover, for this theory to hold, it is crucial that each evaluation is independent. In Appendix A.2 we provide a programming code that can prove the above facts.

We know that due to mimetic behavior, see (Cipriani and Guarino, 2009),

(Chang et al., 2000) for literature reviews, investors don't always make independent interpretations about asset proprieties. For this reason, the basic assumptions for the validity of central limit theory are not met in financial markets. Economic studies, as reported in (Kirman, 1992), usually use the concept of a representative individual. The above consideration casts doubt on such simplifications of real systems.

The mimetic and asymmetrical biasing of information or its interpretation is ubiquitous in financial systems. This behavior is found with company managers as well as financial analysts. Company managers, as reported by (Kyle, 1985) or (Brandouy et al., 2000), have the incentives to abuse insider information or deform financial and accounting information. Through the process known as "window-dressing", company managers can report inflated numbers concerning the performance and well-being of their companies. These practices, also called "creative" accounting, have led to mispriced stocks and artificially enabled the rise and fall of numerous companies like ENRON¹, Lehman Brothers or WorldCom. A good review of such occurrences, as well as their ethical context, can be read in (Amat and Gowthorpe, 2004) or (Beatrice and Grosanu, 2011).

We have to consider that investors understand and can, sometimes measure, the quality of the information they receive. Moreover they understand the fact that there is an asymmetry of information between investors as well as between investors and managers. Therefore, the (perceived) levels of information quality and asymmetry can represent in themselves sources of information.

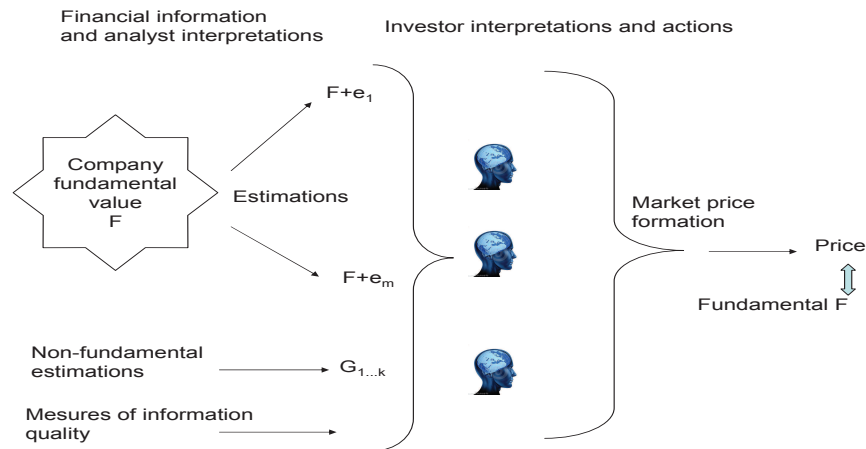


Figure 3.3: Transformation of information into market prices

As in Figure 3.3, investors are faced with multiple sources of information that

¹Read (Barrevel, 2002) for detailed explanations about the creative accounting methods used in the case of ENRON

can shape their estimations. It is easy to imagine a scenario were a bubble can arise due to poor information. To avoid such wrong interpretations, investors usually attach a measure of risk to investing in a certain stock. This perceived risk has a number of explanatory components:

- *fundamentals of the company*
- *quality of information*
- *quality of market*

By estimating risk factors, an investor can compensate for the lack in the quality of information. Like with economic information, risk-related information can also be biased and lacking in quality.

The subprime crisis is a good example of a problem that arises from the lack of information about risk factors, see (Landier et al., 2011), as well as the dot-com bubble² is an example for bad market estimations of future cash-flows. In the first case, new financial innovation permitted the creation of mortgage-backed assets that were so opaque and complicated that investors could not make good estimations about their respective fundamental values. Yet, even if these assets would have been correctly price it is very possible that the markets would have missed an important risk factor: the systemic risk. The collapse of the Lehman Brothers investment bank affected not only the mortgage-backed securities market but the entire American financial system. Soon after, a chain of events unfolded and researchers and regulators called for more transparency on financial institutions as well as for more information collection in regards to 'systemic risk'.

Systemic risk refers to the possibility of collapse of an entire financial system or of an entire market. It should not be confused with a well-known risk called 'systematic risk' which is not diversifiable and refers to the vulnerability to events that affect aggregate outcomes of markets.

We saw how information about a specific asset comes in high volumes and can be wrongly interpreted. The next phase, in the process of transformation of information into prices, is represented by *the investors' interpretations and actions*.

3.3 Investors

Investors are people with savings (or managing the savings of others) who are looking for opportunities to invest their capital. They have access to a wide range of information sources and by compiling this information they make their investment

²See (Ofek and Richardson, 2001) for a complete description of the rise and fall of Internet stock prices

decisions. The process of transforming information into an aggregate measure, which reflects the attractiveness of a certain asset, is difficult and partially subjective.

Classic financial and economic theories often make hard assumptions about how investors interpret and use information. Until recently, these assumptions included that investors interpret information correctly and react only to economic information (meaning they don't have biases and are not persuaded by rumors). The founding studies of (Kahneman and Tversky, 1979) and (Kahneman et al., 1981) showed that investors display a set of behavior biases that affect both the information they perceive as well as their investment actions. Next we describe the most common behavior biases that affect investors in financial market.

3.4 Behavior biases

According to the modern financial theory, all investors are perfectly rational, risk-averse, correctly interpret information and come up with the same future expectations. It is possible that such investors exist but it is not likely that all investors have these characteristics. To explain 'non-rational' behavior cognitive psychologists have worked to understand how people actually interpret information. We underline the idea that the following human cognitive traits are defined as biases in reference to a theoretically correct rational expected behavior. Recent studies, see (Dupre D., Girerd-Potin I. and P., 2006) and (S. and P., 2009), have showed that investors can use non-economic based measures, like ethical or social-responsible measures, to adjust their portfolios. Therefore we cannot attach a character of 'wrong' or 'right' to such human qualities. Alongside the definition of a behavioral bias we explain their possible impacts in financial markets.

The *cognitive dissonance* bias is representative for a person that ignores information which is contradictory to his current beliefs. Imagine an investor that has bought a stock and formed the belief that the stock's value will rise (which would give him a profit). Having a 'cognitive dissonance' bias the investor can sometimes ignore relevant information that indicates the contrary to his beliefs. Such a bias can be a simple explanation for price bubbles, where even if abundant information exists about prices being overvalued investors still ignore them and continue trading at high prices. Another implication of this bias is seen when investors tend to hold on for too much time to losing stocks or when they do not cash in in on winning stocks.

Another famous and important bias is called the *hindsight bias*. This bias affects investors that alter their beliefs in face of new information. This implies that new and revealing information can drive investors to instantly change their beliefs and forget their older perspectives. A trivial way in which investors exhibit such a bias is when they say: "I knew it all along that a crash was due" (even though they didn't exit the market before the crash). Because investors have a

tendency to view the realization of risky events as being 'already predicted' they tend to become overconfident with their investment strategy and in general with their ability to predict markets. As mentioned by (Shiller, 1980) "hindsight bias encourages a view of the world as more predictable than it really is", Even in face of important losses, investors can, due to hindsight, find ways to defend their strategies and beliefs. Because investors sometimes need to be right (more than they need to rest objective) it is plausible to say that some will persist in markets even when they suffer important losses. In this context, it is possible that incorrect strategies or price expectations can persist, for long periods, in financial markets.

Closely related to the hindsight bias, *availability bias* is also observed in the behavior of investors. The 'availability bias' implies investors tend to overestimate the probability of an event when examples of realizations of the event are presently available. For example, when seeing a rise in stock price for a few periods and investor can make a biased judgment and draw the conclusion that prices will continue to grow. If prior information provided a different picture (not supporting growth) than an investor can still change his beliefs (in favor a raising price trend) if he is also affect by the 'hindsight' bias (which will distort his recollection of older information that negates the new raising trend belief). In a more general context, the 'availability' bias is often triggered by news bulletins that focus on a particular type of event (ex: shark attacks, car accidents, etc.). If a news bulletins is repeated a few times, a considerable number of viewers will estimate the probability of occurrence of the events displayed (probability of getting attacked by a shark) as much higher than it actually is in reality. Thus the 'availability' of a few repeated realisations of an event will distort people's estimation of the likelihood of the event. The 'availability' and the 'hindsight' bias have been described in (Kahneman et al., 1981).

In opposition to the previous bias, some people tend to exhibit an *overconfidence* bias, or *conservatism* bias as first proposed by (Edwards, 1968). Because of this bias, investors tend to underestimate the weight of new information and overestimate the importance of their older information. In the he famous study about car drivers, (Svenson, 1981), the researchers asked people to rate their driving skills relative to the average driver. The results, mathematically impossible, showed that a high majority of drivers considered themselves as having 'about-average' driving skills.

The *representativeness* bias affects human cognition by allowing for overvalued beliefs about the probability of an event provided we observe events that have similar proprieties with the event we are looking for. Such a biased interpretation will trick us into making generalisations about the statistical proprieties of an entire population based on the proprieties of a single sample. If we prompt a person to estimate the probability of a 'special' coin falling on its face (1) most would probably answer 50%. In contrast, if we show a realisation of a series of tosses, like 1101011, most people will wrongly assume that the probability of the coin showing 1 is more than 50%. The sequence 1101011, of realisations of a coin toss, has the

same probability of appearing as does the sequence 111111 (or all zeros) and these samples come from the toss of a perfectly balanced coin. In a financial market, such a cognitive bias can be responsible for trend creation: investors who observe a close sequence of positive returns can infer that it is more likely that positive, rather than negative, returns will follow. In a large time-frame, this equates to the belief that 2,3 consecutive years with positive returns implies that the next year's returns are more likely to be positive than negative (fact that is statistically incorrect). In (Griffin et al., 1993), the authors propose two characteristics for information: force (intensity of change in the watched variable) and statistical weight (probability of such event occurring). They observe that information with high force is considered as pertinent even if its statistical weight is insignificant.

The five types of biases presented can, in isolation or combined, have multiple effects on the way investors perceive and interpret information of different types. Besides a biased interpretation of information, investors can diverge from the 'rationally expected behavior' because of their investment objectives which can be different from the maximization of utility of wealth.

After investors have interpreted the available information they send their investment orders to the market. These orders are gathered and introduced in a price formation mechanism that creates the market prices. While this mechanism is completely deductive (unbiased and predictable) some of its variations can have properties that influence the quality of markets and the ways through which markets become more efficient. Therefore we describe in the next chapter different methods for the formation of market prices.

3.5 Price formation mechanism

After receiving and interpreting information, investors send orders to the market in relation to their investment objectives. These orders can be divided in two main categories:

- *Market orders*: orders that specify the direction of the action (buy/sell) and the amount of money/stocks that are involved in the trade. These orders will be executed, if possible, at the best available market price. Market orders usually have higher priority than limit orders.
- *Limit orders*: orders that have the same specification as a market order with extra information describing the maximum/minimum price the investor is willing to accept in order to buy/sell assets. There are several variations of limit orders (ex: all-or-cancel, stop-losses, time-limits) which help investors take very specific actions in a market.

These orders are received and introduced into the market's execution system. This system consists of a set of procedures that discover market prices and match

orders of buyers and sellers. Depending on the market, the execution system has different types of sessions: call market sessions or continuous sessions.

With *call market sessions* trading occurs at well-specified times. All orders are gathered in an order book and at specified moments times a market price is discovered based on certain requirements, usually volume maximization (see chapter for a detailed presentation of such a price discovery method). In this type of market, the trade execution system is called '*single price auctions*'. All matching trades are executed at the discovered unique market price (as long as this market price is between the buy and the sell prices).

During *continuous sessions* trading occurs at any time during the session. A market order is executed immediately using the best counterparty available. A limit order is executed when another limit order is suitable as counterparty (or a market order arrives). Orders that are not resolved are placed in a limit-order book for later processing. This type of session allows traders to execute their orders faster (and therefore quickly profit on their information). The execution procedure for trades in this case is called '*continuous two-sided auctions*'.

Having reviewed the types of orders available as well as the trading sessions, we discuss the different types of execution systems available:

1. Quote-driven markets: investors don't have the possibility of trading between each other. A specialized financial partner, called "market maker", places quotes at which he is willing to buy and sell assets. The market maker offers liquidity and is the counterparty for every investor.

2. Order-driven markets: Investors can trade directly between themselves. The orders are stored in a common place and are solved according to different rules. There are numerous variations of order-driven markets that dictate the way in which trades are matched:

a) Oral auctions: investors publicly express their offers for buying and selling actions. In opposition, sealed auctions give investors the possibility to make private offers (usually found in real-estate markets).

b) Rule-based system for matching orders: according to specific rules the prices for trading are discover. Different rules are used for different types of markets. From these we describe one of the variations which we didn't previously mention (like single price or continuous two-side auctions): *Crossing networks*. In crossing networks trades are executed at prices that are not produced by the orders from this market. The trading prices are set elsewhere, usually another market. These markets are very specialized and are considered as '*alternative trading systems*'.

A good review about the different microstructure aspects of market can be found in (Harris, 2002). It is important to understand that each market price formation mechanism serves specific objectives. Yet, these systems are all conceived to "correctly" (unbiased) represent the beliefs of the trading investors.

3.6 Price anomalies and explanations

We saw the elements that assure the formation of prices in financial markets. Now we define and explain some of the puzzles that are found in the time series of real financial markets. We will be reproducing these "puzzles" using our model of a financial market.

If we consider any price/return series from a financial market, it would be interesting to compare such price series with the asset's fundamental values. Because such fundamentals rarely exists (like audited periodic financial reports, which have too few observations) we are left with looking at the statistical proprieties of these series. According to the consequences of the efficient markets theory, assets returns should have normal distributions. Instead, real market return series exhibit distributions with 'heavy tails', also called *leptokurtic*. This feature implies the following: there are too few average returns and too many extreme returns as opposed to what the theoretical model proposes. To illustrate this 'financial oddity' we can look at the return series of the French CAC40 index in the period 1990-2012.




Figure 3.4: (black) Distribution of CAC40 returns, (green) Normal distribution

We observe that the CAC40 has more extreme returns (the average returns \pm three times or more the standard deviation) than expected. Similarly, the return distribution has less than expected medium returns (average return \pm two times the standard deviation) and more than expected average returns. One other return anomaly called 'day-of-the-week', (Louvet and O, 1990) and (Louvet and Dubois, 1996), implies that returns are consistently higher in some days of the week than in others. Discovery of this inefficiency led the way to its disappearance and it is now rarely observed.

Another interesting financial puzzle has the name of '*volatility clustering*'. As first explained in (Mandelbrot, 1963), the puzzle implies that *large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes*. This phenomenon is put into evidence by the existence of a slow-decaying autocorrelation function between the absolute or squared returns of asset

prices.




Figure 3.5: Autocorrelogram of absolute CAC40 returns

As documented by (Louv et al., 1993) and (Louv et al., 1997), if we look at the same market indicator, CAC40, we observe that its simple returns are not autocorrelated, therefore confirming the weak-form efficiency theory. Looking at Figure 3.5 we observe that absolute returns are correlated to each other with a strength that decreases slowly and lasts for almost 200 days. This propriety is also called 'volatility clustering', a tendency that returns can have persistent levels of low or high volatility levels. Because high/low volatility implies, with statistical significance, a future high/low volatility level we say that the behavior of the market at time $t+1$ can be, in a certain degree, estimated from the market data at time t . Thus this propriety is also called a long-memory of volatility, see (Ding et al., 1993) for a mathematical model. As shown in (Brandouy et al., 2012a), stylized facts like the ones presented above and below, have been already reproduced using statistical models or market microstructure constrains which are unfortunately unrealistic in terms of investor behavior. Different methods for measuring market efficiency have been used, from the ones mentioned above to others that use complicated statistical models like (Herve, 1992). In our study we investigate the possible behavioral characteristics that cause such empirical financial puzzles.

Price bubbles are one of the most important 'puzzles' of modern financial markets. Even the definition of such an event is still debated and there is no widely accepted definition. In (Stiglitz, 1990) the author offers this definition of the puzzle: "*the basic intuition is straightforward: if the reason that the price is high today is only because investors believe that the selling price will be high tomorrow-when "fundamental" factors do not seem to justify such a price-then a bubble exists.*" If we disregard its actual causes, the fact remains that we can recognize bubbles when prices raise and then drop dramatically.




Figure 3.6: US residential real-estate price index (Source: S&P Case-Shiller Home Price Indices)

One of the most recent important price bubbles was seen in the US real-estate market, as explained by (Levitin and Wachter, 2012), where prices kept rising for the last ten years, escalating in 2007 with an abrupt growth and crash (see Figure 3.6).

Besides these classical 'financial puzzles' we draw the attention towards two market quality measures: *volatility of returns* and *signal to noise ratio*. The *volatility* of returns is today at the kernel of all generally accepted investment risk measurements. The higher the volatility of an asset the higher are its risk measures. Investors compute risk using the information they receive. If investors do not agree on risk levels, prices fluctuate until a consensus is found. So, the more precise and clear the information the easier can investors interpret and estimate risk. With the advent of electronic markets, high-speed internet and digital media it is clear that information is available in abundance and on demand. Therefore we can imagine that increases in volatility, due to information asymmetry or lack of quality, should have disappeared. The empirical evidence shows the contrary. In the middle graph we can observe that the volatility level has been steadily increasing the in last 15 years.

The signal to noise ratio is computed, see Annexe A.3 for code, using the formula $S = \frac{\mu_R}{\sigma_R}$. This measure is usually used in engineering to measure the ratio between a desired level of a signal and the surrounding noise. In finance we know that returns are linked to volatility via risk. The higher risks an asset has, the higher should be

its returns as well as its volatility. We measure S , see Figure 3.7, for the returns of the CAC40 index and observe that it is decreasing. In windows of 250 days, we compute the volatility and mean of returns of the CAC40 market index.

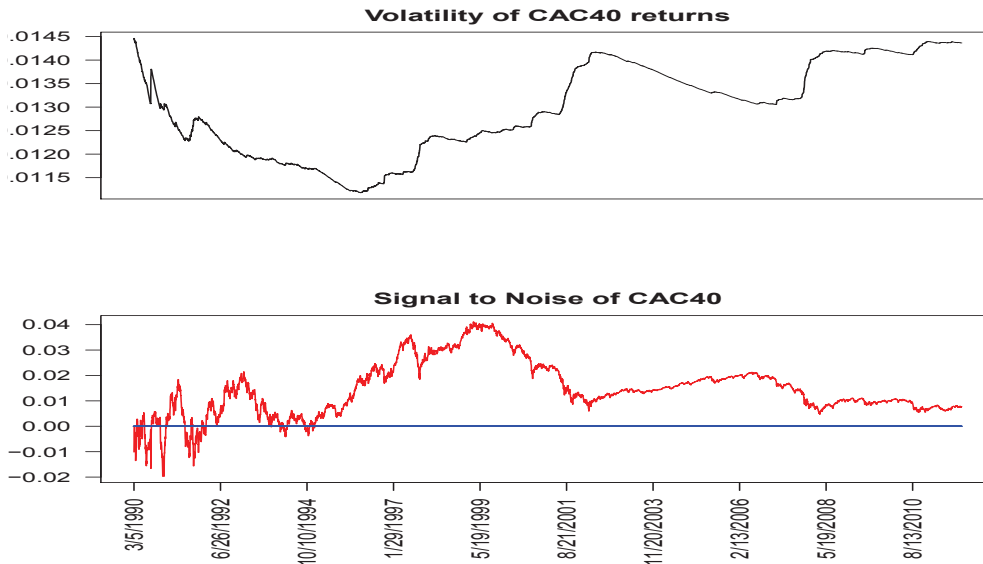


Figure 3.7: In 250 days time-periods: (above) Volatility of returns of the CAC40, (below) Signal to Noise Ratio of CAC40 returns

We observe that the signal-to-noise ratio is decreasing, which implies that volatility (a measure of risk) is growing much faster than returns thus rendering prices less and less informative. The authors of (Silver, 2012) explain in detail, and with a great number of examples, how can signal to noise ratios explain the failure of most investing strategies. From an informational view, we are left wondering if the new technologies, like high frequency trading (see study of Litzenberger et al. (2012)), have more likely hindered the quality of information that investors receive. Another explanation could be that instead of the changing quality of information, it is the changing structure of economies (more interconnected, more fragile, with unpredictable political influences) that have provoked a general increase in the perception of risk and consequently in the growth of volatility in financial markets.

All these empirical financial 'puzzles' are proof that the predominant financial research theories are not entirely reflected in real financial markets. As we mention in the first chapter, most modern financial theories are theoretically correct but the market states they describe are usually normative. Real market have not yet been reached such states. Researchers, using classical approaches, are trying to explain the anomalies by making slight theoretical alternations to accommodate these empirical realities.

3.7 Classical approaches to studying market dynamics

Market dynamics refers to the study of the aggregate behavior of markets. The theoretical foundations for the market models, used to describe such dynamics, are based on a few important paradigms:

- Portfolio theory developed by (Markowitz, 1952)
- Capital asset pricing theory, CAPM, developed by (Sharpe, 1964)
- The expected utility theory initiated by (Bernoulli, 1954)
- Efficient market hypothesis proposed by (Fama, 1970)

As described in the previous chapter, the ensemble of these theories has been found not to coincide with some empirical proof (also called 'stylized facts'). Good reviews of empirical puzzles can be seen in (Turner et al., 2012) and (Cont, 2001). The assumptions that rest at the foundations of these theories are rather strict and are often quoted as the source of the discrepancy between the theoretical description and the market reality. The aforementioned theories share a number of fundamental assumptions:

- *All traders are rational*: In the view of the founding theories, rational implies that the trader has perfect information and maximizes the expected value of his utility function.
- *All traders have homogenous expectations*: This assumption implies that traders share the same expectations about returns and risks of assets. This assumption could be valid in two ways:
 1. All traders share the same information and since they are all perfectly rational they arrive at the same conclusion
 2. Traders don't have the same information but they exchange between them information on their expectations of risk and return
- *Trading mechanisms don't affect price formation*: Different methods of price formation allow for the same market prices. This statement also assumes that individual trades don't have any effect on prices.
- *The market is efficient*: This assumption has many implications and acts as a safety-net for the previous assumptions. Firstly, this idea implies that all less than rational investors are supposed to be driven out of the market because the other rational investors will profit off of them. Secondly, rational agents with less than perfect information will be driven out of the market by agents with superior information. Thirdly, because the market will have only rational investors with perfect information it will behave as an efficient market.

Some of the first empirical studies of the dynamics of financial markets intended to support these market theories. New and more powerful tests have revealed many more inconsistencies in real markets. (Un) fortunately some of the stylized facts first discovered (like the 'January effect' or the 'Monday effect') have faded away (probably due to arbitrage). The disappearance of such empirical facts has sparked the debate about whether stylized facts are just phenomena that will be eliminated by the markets (because of their imminent efficiency) or that implies the need for revising the founding theoretical models of financial markets.

Researchers have focused their attention on the verification of all of the main assumptions of the modern financial theories. Microstructure studies focus on the links between price formation and the price formation mechanism. Evidence, like (O'Hara, 1998), exists that the price mechanism can influence the efficiency of markets. Moreover the distribution of information throughout a market is an important factor for market efficiency (Glosten and Milgrom, 1983).

Behavioral finance, an emerging research trend, investigates the assumptions of investor rationality and homogeneity. The two main methodologies used in behavioral finance are experimental studies and agent-based simulations. Experimental studies consist of either surveys or laboratory experiments. These studies are useful because they can control, to a certain extent, the environmental settings of the studied individuals and observe their behavior when face with investment related decisions. The aim of this type of studies is to discover how investors make decisions under risk. Ground-breaking advances using these research methods have been made by Nobel Prize winner Daniel Kahneman (in association with Amos Tversky). Their 'prospect theory', (Kahneman and Tversky, 1979), shows how investors have biased perceptions of information depending on environmental conditions (e.g. losses, wealth). A good article explaining the advantages of behavioral finance is (Mitroi Adrian, 2007).

Microstructure and experimental studies have provided proof that most assumptions, used by the founding financial theoretical models, are not encountered in reality. Still, the implications of the market efficiency theory state that these empirical findings are a source of profit for rational investors and are sure to be eliminated. To test such assumptions, a new research methodology has been devised, called 'agent-based artificial market simulation'.

3.8 Summary

In this chapter we presented the components of a generic financial market, from the point of view of the transformation of information into prices. Firstly, we showed how information quality can affect the formation of prices. Secondly, we explained the main categories of biases, of perception, behavior or cognition, which real investors suffer from. Thirdly, we described the main types of price

formation mechanisms that exist in real markets. We gave references to studies that show how these mechanisms affect the efficiency of markets. Next we described the 'stylized facts' and other market quality measures that show there are inconsistencies between market empirical realities and the implications of current financial theories. In the conclusion, we presented the main classical approaches used in analyzing the dynamics of financial markets.

Agent based artificial stock markets

Contents

4.1	Complex systems	51
4.2	Agent based simulation	56
4.2.1	Definition and components	56
4.2.2	Applications	56
4.2.3	Advantages and disadvantages of research using agent-based simulations in finance	59
4.3	Topology of financial market simulators	61
4.3.1	Notable simulators	62
4.3.2	Simulated market	68
4.3.3	Simulated time	69
4.3.4	Information	69
4.3.5	Bias in information processing	70
4.3.6	Risk and time preferences	71
4.3.7	Agent learning	71
4.3.8	Agent strategy bias	73
4.3.9	Price formation	73
4.4	Critics of agent-based artificial stock markets as a method- ology for research	74
4.4.1	Equivalence between an agent-based artificial market and a mathematical analytical model.	76

4.1 Complex systems

Before we define and understand the purpose and meaning of an agent-based simulator we will explain the theoretical basis of the problems it solves.

In the world of academics, the hypothesis and methodology of reductionism, as noted by (Nagel, 1961),(Anderson, 1972), is widely accepted as a valid and productive method of research. Methodological reductionism can be described, loosely, as

the idea that a system, be it a human, object or phenomena, can be fully explained by understanding the laws that govern the inner basic particles of that system. Reductionism can be view as theoretical, when one theory can completely explain another theory, methodological or ontological, for example when we view the world of being composed of a single building part (like the atom). We will focus on the methodological aspects of reductionism.

In (Anderson, 1972) the author confronts this hypothesis and shows a few examples when reality points that reductionism is not enough. We propose a few examples of systems, closer to our domain of expertise, that cannot be fully understood using a reductionist methodology. These examples will guide the reader towards understanding complex systems, their problems and how can one study them.

A system is defined by a set of components, and the interactions between them, located within a boundary that defines the surroundings of the system. For example, a pool table can be considered as a system (see Figure 4.1). The balls and the sides of the table represent the system's components and the outer sides of the table act as a natural boundary. The balls and the table sides can interact and are guided by a mix of forces: gravity, friction (with air and table) and elastic deformation (when balls collide with each other or with the table sides).




Figure 4.1: Representation of generic complex system

Because such a system, a pool table, follows known rules we are able to construct an analytical model that predicts that state of the system (position of the balls, etc.). If we interact with the system, by hitting a ball with a cue, we can also easily predict the trajectory of the ball and that of the first balls that are also touched. In fact, professional pool players do these kinds of models and predictions in order to win pool games. A similar system is represented by a car on a road. We know all

the components of a car and how they work and interact with each other. If we start the engine and push the right pedals the car will start moving forwards (or backwards) in a direction that depends on the position of the steering wheel and the gear selected in the gear-switch. So far, a car looks like a simple system. We can define a "simple system" as a system that has, as a whole, properties (trajectory, states, etc..) which are determined from knowing the properties of all the individual components. A pool table and a car are considered as "simple systems".

What happens if we decide to introduce a human driver into the car's system? Are we still able to predict the car's position? We can, provided we know the laws that a human driver obeys. Unfortunately, there are no good models that can predict the way in which a human driver acts. Consequently we make assumptions about how a driver will act, based on certain observations. We assume that a driver reduces speed to avoid obstacles or stays on the correct side of the road. Using this model we make predictions about the road system's state in the future. Moreover, knowing that policemen place sanctions and fines on drivers that misbehave, we can even postulate an 'efficient roads' theorem:

'A transportation system is called efficient if the trips of all cars are fast and on time'. To ensure such efficiency, regulators have created traffic controllers (e.g. traffic police, traffic surveillance equipment) that observe and sanction irregularities. If empirical irregularities are observed (queues, roadblocks, etc.), traffic controllers intervene (with sanctions and fines) thus forcing the 'non-rational' drivers to behave well in traffic. A corollary to this theorem can be that all drivers are rational and obey perfectly the traffic laws. Those that are not rational or don't obey rules are quickly fined and punished until they do become rational drivers or they lose the right to drive.

The result of such a model is that we would miss out on a number of system states (like traveling only in reverse, speeding while approaching an obstacle or changing lanes, flipping the car upside down because of excessive speed in a curve, etc.). Moreover, even if traffic control has been in place for a long time the result is we still have a considerable number of misbehaving drivers that often provoke traffic problems. Even when these human factors are eliminated, traffic is still not efficient since we often see road jams. In real life drivers are often very "creative" and surprising in their actions, and are not at all orderly and respectful of traffic laws (as our "efficient roads" theory predicts). Because of the non-deterministic nature of human thinking and actions, it is almost impossible to create an analytical model of a human. Therefore we cannot easily predict the state of a system which has as a component a human factor. These kind of systems, among others, have been coined "complex systems". See (Licata and Sakaji, 2008) and (Kaneko and Tsuda, 2001) for good introductions to the theory of complex systems.

Because "complex systems" represent a new object of study, they don't have a standard definition. Here are a set of definitions that may better express the idea of a complex system:

- *A complex system is a highly structured system, which shows structure with variations* (N. Goldenfeld and Kadanoff)
- *A complex system is one whose evolution is very sensitive to initial conditions or to small perturbations, one in which the number of independent interacting components is large, or one in which there are multiple pathways by which the system can evolve* (Whitesides and Ismagilov)
- *A complex system is one that by design or function or both is difficult to understand and verify* (Weng, Bhalla and Iyengar)
- *Complex systems are systems in process that constantly evolve and unfold over time* (W. Brian Arthur).
- *Complex systems are composed of interacting units. They exhibit emergent properties, that is, properties arising from the interactions of the units that are not properties of the individual units themselves* (Flake G. W.)

All these definitions have a common idea behind them: a complex system's properties are hard to predict by knowing the individual parts and applying a classic analytical model. There are different subtypes of complex systems. Chaotic systems are very sensitive to input conditions. Depending on the initial state the system can develop (non)deterministic non-evident future states (like the Brownian movement or a double pendulum). Adaptive complex systems have components that evolve and are able to adapt their behavior according to the past experiences (like stock markets, human immunity system, bird flocks, etc.).

Modern science has, until recently, viewed most systems as being simple and has modeled them using few variables and a simple dynamics structure. This method has worked especially well in natural sciences which deal with universal laws (gravity, motion, entropy, etc.). With such models, scientists discovered nuclear energy, gas motors and electricity. Using the same reductionist point of view we immediately notice that two natural systems: diamonds and graffiti (used in pencils). Both systems have the same basic structural components here carbon atoms) yet can have extremely different overall properties when the structure of the connections between atoms is slightly different. Extending this idea, we see that everything around us is made of the same atoms but the patterns of their interconnections can lead to very different objects (e.g. ink, egg shells, hair, petrol, etc...).

If we regard a system consisting of a chicken egg we know that if we don't expose this egg to heat (or to too much heat) it will go bad (rot). Instead, if the egg is properly heated (constant temperature on all the sides of the egg) it will develop into a living chick that will break the egg shell. While it looks simple, from a modeling point of view it is a difficult problem because the slight variation of a single factor leads to completely different results (decomposition vs. the development of a new life form). The complex system representing an egg acts like a chaotic system since

a variation in a single input variable (e.g. temperature) can generate very different results in the system outputs.

Acknowledging that a financial market has, as participants, many heterogeneous individuals, each one of them being a non-deterministic system, we safely say that a financial market is a complex system. Professor Herve Alexandre mentioned, see (Herve, 1994), the phrase "deterministic chaos" when describing stock markets. In the traditional view of modern finance, financial markets are viewed as very simple systems with homogenous investors that compete and create excellent market conditions. This way of viewing and studying systems, from the perspective of an idealized framework, was borrowed from natural sciences. In our opinion, a theory like CAPM is perfectly suited for the study of financial market with "*identical perfectly-spherical investors that trade non-stop using infinite resources in a vacuumed thermodynamically stable environment*". The problem with this approach is the assumption that the developed theory can be adapted and maintained as valid even when some of its' assumptions are relaxed. As we have seen, relaxing variables is not an acceptable method for the case of complex systems where slight changes of the initial system states (heated or not heated egg) can generate completely different future aggregate states. Inside a financial system, a slight change affecting a variable, for example a new law like 'compulsory subprime loans on bank's assets', can create extreme changes in the behavior of participants (real-estate speculation, excessive consumption, higher appetite for risk) or inside the market's structure (debt securitization, asset packaging and repackaging). Such systemic changes, triggered by small variation of inputs, cannot be captured inside abstract dynamic (stochastic) general equilibrium homogenous models. Because they adapt to new inputs and interactions with new trading opportunities, investors can react and create new non-evident dynamics for every new input condition. Moreover, because of this constant evolution of the participants in a financial market an equilibrium state (e.g. efficient market) may not even be reachable (or at least not in a time scale of interest). (Brandouy, 2005a) explains how complexity governs and should be taken into consideration when looking at critical phenomena in the world of finance.

Due to the homogeneous and reactive nature of the components inside a complex system, classical analytical models of research are often not very useful. The concept of multi-agent simulation emerged from this need of making better models of complex systems. By creating a simulated system where the components have relatively simple rules and can interact with each other we can recreate system-wide properties that resemble those of real systems.

An agent-based simulation is a research tool that can model the components of a real system using simple rules and interaction methods. It is important to note that an agent-based system can have the same system qualities of a real system. For example, using this method researchers have discovered that birds need to follow three simple rules to form the V-shaped flock arrangements we observe during migration periods.

In financial markets, the goal of agent based simulators is to find sets of micro specifications and interaction rules and observe how these rules affect the main characteristics of the market: prices, volumes, microstructure quality and informational efficiency. Classical approaches for studying financial markets can be regarded as equation-based models. Such models usually make hard assumptions like that all investors are perfectly informed, that they are rational (as opposed to bounded rationality which we explain in this chapter). Agent-based modeling provides a rigorous method for relaxing and testing these assumptions (e.g. limited information, bounded rationality or observing far from equilibrium states).

4.2 Agent based simulation

4.2.1 Definition and components

An agent based simulator is a research tool that shows the emergent proprieties of a system starting from the behavioral definitions of the system's components. By using this tool, we describe the fundamental components of a system, as well as their interactions, and can simulate the interactions between these components at any level of aggregation. From a broader point of view, agent based simulators allows us to create in-vitro experiments of large systems (like reproducing a city's transportation infrastructure). In this manner we observe the correlation and causality between inputs, sets of behavioral rules and the aggregate proprieties of the system.

Because agent based simulation are a generic research tool we will look at its applications related to problems faced by social science, and especially finance. In academic finance, we can distinguish a particular branch of research, called agent based computation finance, that try to explain emergent phenomena using multi-agent tools. Good introductions about the concepts, approaches and methods of agent based computational finance, are available from (Brandouy et al., 2006b), (Epstein, 2006),(Veryzhenko, 2012), (Derveeuw et al., 2007a), and (Beaufils et al., 2009).

4.2.2 Applications

Agent based simulation may be used in applications where emergent phenomena are found. When choosing such a research tool, we have to search for the problematic characteristics of the systems that need to be modeled:

1. Simple averaging statistical methods do not work. Systems exhibit the most interesting particularities in extreme conditions and averaging may hide such events.
2. Individual behaviors cannot be described by simple functions or differential equations. Instead, these behaviors are best modeled using condition-action

rules (ex: thresholds for changing the pattern of behavior, specific input - specific action, rare behaviors or borrowing of strategies from other individuals). Moreover, individual behavior exhibit non-linear variations and can have memory. (Huang et al., 2010) and (Brandouy, 2005b) make an attempt at explaining financial dynamics using investor behavior and cognitive heterogeneity.

3. Components of the systems interact locally, based on non-homogenous topologies (like a scale-free topology - see Figure 4.2). Slight differences in the system's topology leads to very different aggregated behaviors even if the individuals are the same.
4. The emergent proprieties of the system are not evident if we look at its individual components and interactions.




Figure 4.2: (A) Centralized (B) Scale-free (C) Closely distributed - Graphs with different topologies of connections between different entities of a system

In Figure 4.2 we observe three types of topologies that can be used to model interaction patterns between the components of a system. The model of interaction chosen is important when explaining aggregate behaviors. For example, we can consider a national banking sector as being organized using a centralized model (see Figure 4.2 A). In this case, the central node of the banking sector can represent the central bank. With this representation it would be easy to consider that the only source of financial contagion (or cascading bankruptcy) is through the central bank. For mathematical simplicity, models like the Erdos Reyni¹ random graph-model can be used. Even if these random models are more simple to handle they do

¹see (Erdos and Renyi, 1960) for the original description of the Erdos-Reyni random graph model

not manage to capture the complexity of real interaction patterns between banks. In real banking sectors, as reported by (Galbiati and Soramaki, 2008) and (Galbiati and Stanciu-Viziteu, 2013), bank interactions are better described by, see Figure 4.2 B, scale-free networks (some banks tend to have much more connections than other banks thus creating more points of contagion). Of course, scale-free networks are very particular and assume that the components of a systems (here the banks) have a behavioral trait of preferential attachment. Preferential attachment can also be observed in social relationships, where people tend to cluster around a certain individual (like the father of the family). Other examples about the importance of topology and interaction models and methods are provided in (Smaldino et al., 2012) and (Derveeuw et al., 2007b).

Agent based simulators have been in use since the development of modern computers. Some of these simulators' earliest uses have been in the field of automobile traffic analysis. See (Shen et al., 2011) or (Meister et al., 2010) for examples of real application of such simulators. Through simulations, city officials observe the impact of changes in stop-light timing and positioning, roundabouts, priority or one-way lanes. These simulations show traffic behavior in normal states (normal driving hours) but also in unstable states (peak hours, when people go/return from work). Such a simulation tool permits city planners to test and adapt street policies before implementing them on the streets. In medicine, agent-based simulations are used to model complex systems such as the human immune system, inflammation mechanisms or the spread of bacteria and viruses (see (Paranjape and Sadanand, 2010) for a good review). In commercial applications, agents are used to model logistic and supply chains or consumer behavior. In civil constructions, agent simulations are used to observe the movement of crowds inside buildings. These results will improve evacuation safety, as explained by (Bo et al., 2009), for future buildings (see Figure 4.3).




Figure 4.3: Simulation of the flow of people evacuating a building

One of the more recent applications of multi-agent simulation is in the field of social sciences and especially in economics and finance. Oil and other natural resources production processes can be modeled in such simulators. In finance, agent simulators are now one of the main tools supporting behavioral finance in its path towards recreating the manner in which researchers think and view financial markets. Economists are already using agent-based simulators for different purposes: observe the formation of certain equilibrium states, measure impact of different macroeconomic policies and recently to measure systemic risks. This approach in economics, usually including game theory, is called artificial economics. (Feichtinger, 1991), (Beaufils and Brandouy, 2005) and (Epstein and Axtell, 1996) explain the complex system view of economy, the agent-based methods that should be used and how these can explain puzzles that were not yet solve using traditional research tools.

In academic financial research, agent-based simulators are used to observe the conditions that can enable (or not) market efficiency, see (Brandouy et al., 2006a), the polarization of investor wealth, see (Bottazzi and Anoufrieu, 2004) or the aggregation of information into prices. For an excellent review we recommend (Tsfatsion, 2006b), (Epstein, 2006) and (Veryzhenko, 2012). (B. et al., 2008) makes a more general assessment of the complexity aspects of modern financial systems that can be better understood with agent-based tools.

4.2.3 Advantages and disadvantages of research using agent-based simulations in finance

Classical methods that test market efficiency have the methodological issue called *the joint hypothesis problem*. The underlying idea of efficiency is that market prices

reflect fundamental values and this assumption can be tested only if we actually know the fundamental value model of assets. Therefore, researchers propose a model for the fundamental value (which they assume is correct) and then resolve to see if the market price converges to the theoretically correct values or not. The validity of the efficiency test is linked with the validity of the joint fundamental value model. In an agent based artificial market, the method used to generate the stock's fundamental value is axiomatically correct. Therefore the researcher has to observe the ways in which individual investor characteristics can influence an aggregate output (like the price) even when they know the fundamental value model.

We list and explain some of the most important advantages of agent-based simulator research:

- *Attainability*: As described in chapter 2.7, a theoretical efficiency state in a market may be true but not necessarily attainable (at least not on time scales of interest). Using an agent-based model we can show how, starting from a plausible specification of a market state, a market can reach (step by step) an efficient state (if possible).
- *Causality*: Because an agent based simulation is deterministic and completely observable, we can make accurate 'ceteris paribus' assumptions and observe the causal links between certain micro specifications and their desired aggregate states.
- *Flexibility*: Because of the language used in describing a financial market (computer programming language) the model descriptions (especially those used for describing the behavior of investors) is easily modified and adapted to different needs. This flexibility allows researchers to easily change descriptions of investors (or of other market components) to better represent the realities of specific financial markets and their context.
- *Normative understanding*: Because investors can be modeled to adapt and react to the environment conditions (e.g. taxes, transaction costs) we can observe the emergent behavior of markets under a diverse set of norms and regulations. Therefore we can make 'simulated' experiments with different regulations and observe if the desired aggregate results can be achieved and in what time frames.
- *Transparency (Observability)*: Through-out a simulation all the components and their states are perfectly observable. This provides the possibility of observing and controlling very specific parts of a financial market.

According to (Campbell, 2000), "*Behavior models cannot be tested using aggregate consumption or the market portfolio because rational utility-maximizing investors neither consume aggregate consumption (some is accounted for by nonstandard investors) nor hold the market portfolio (instead they shift in and out of the*

stock market)". So, in order to test such models we require detailed information about the trading behavior of each market participant. Because such private information is confidential a good alternate is using agent based simulators. In a laboratory setting, an agent-based model can provide qualitative knowledge about financial markets dynamics. To obtain quantitative results we need to link such a model with real data that will be used for the initial micro-specifications of the model's components.

Agent-based artificial markets also have certain limitations which we must take into account. Because we describe an entire market, a good model will unavoidably have a high number of parameters. Since some of these parameters are hard (even impossible) to measure in reality, we are forced to make assumptions. Like the assumptions that are the base of a theory, our initial parameters are the assumptions we make within our agent based model. The validity of these assumptions is a crucial condition for the validity of our results.

([Veryzhenko et al., 2011](#)), ([Veryzhenko et al., 2010](#)) explain in detail and provide examples of the main issues that must be considered when creating realistic agent based models of financial markets.

Robustness checks and code validation are required to ensure that results of simulations are not the product of chance or designer intervention. Moreover, when describing micro-level specifications (invest behavior, price formation, information generation and dissemination) it is important to rest objective and not use models that directly incorporate the causes of the searched-for phenomena. There are different approaches to testing the models used. ([Brandouy and Mathieu, 2007](#)) provide a conceptual framework for the evaluation of technical analysis and agent-based trading.

4.3 Topology of financial market simulators

To better understand and distinguish between agent-based financial market simulators (ABFMS) we propose a taxonomy which describes the principal design elements of an ABFMS. We started from the short classification proposed by LeBaron in ([Tessfatsion, 2006a](#)). We expand this typology with new classifiers that shed more light onto the design of an ABFMS.

An agent-based financial market simulation represents an abstract computerized version of a real financial market. Its objective is to model and provide access to the causal links between elements of investor behavior (or market components behavior) and aggregate securities prices. Through the observation of this simulated system we obtain new hypotheses for the dynamics of financial markets - especially for the dynamics that are not explained by standard financial methods and models. We propose a new set of characteristics to describe and distinguish between ABFMS. Each of these characteristics represents a basic design question that one

must answer in order to build a simulator. Researchers try to build parsimonious models of the reality and in the case of an ABFMS they struggle to find a small number of parameters that controls the simulation of a big number of real features of a certain financial market. (Mathieu and Brandouy, 2010),(Veryzhenko et al., 2011),(Veryzhenko et al., 2010) provide hints for a generic model of financial markets that can generate *realistic* market simulations.

By decreasing the number of parameters, simulators tend to become more abstract. In this way, the simulator design ignores certain important aspects which may be crucial for creating the logical link between the simulation and the real financial market.

4.3.1 Notable simulators

We present the main characteristics of six financial market agent-based simulators that have been often quoted in the literature.

The 'Santa FE' artificial stock market, (LeBaron, 2002) built by Blake LeBaron, Arthur W.B and Palmer R, is one of the first multi-agent based simulators. As initially called by the authors, 'an experimental computer simulated stock market', their model was created to understand the dynamics and time series proprieties of market prices. The authors created a market where mini-computer programs, called agents, could invest their wealth in a portfolio consisting of a risk-free bond (with a constant return) and a risky stock that has a dividend which follows an autoregressive process. Agents are assumed to be myopic with one period constant absolute risk aversion (CARA) and sharing the same absolute risk aversion coefficient. The actions of the investing agents have as objective the maximization of an expected utility function of the form:

$$\hat{E}_t^i(-e^{-\gamma * W_{t+1}^i}) \quad (4.1)$$

where γ is the agents shared absolute risk aversion coefficient, and W_{t+1}^i is agent's i estimated wealth in the next period. Because an agent wealth at time $t+1$ includes his demand of the risky asset at time t , the authors compute this demand function for each agent. The differences in the agents' demands are set by their forecast of the future market price and the stock's dividend. To make predictions of future prices, agents maintain sets of information called 'condition-forecast' which are rules that map a market condition (based on fundamental or technical information) to a forecast rule for the future market state. Using genetic algorithms agents combine and create such forecasting rules in order to make accurate estimations of future prices. The market price is set endogenously by a market clearing mechanism that resolves agent demands in relation to a fixed supply of stocks. Observing the resulting time series, the authors conclude that such a market model provides similar statistical proprieties with those of real market price series (e.g. volatility, predictability).

More interestingly, the authors point out that the horizon of agent forecasts (the longer the better) is important in achieving or approaching a rationally expected equilibrium. This market model is one of the first complicated artificial models of a stock market. Even if the model manages to qualitatively replicate some of the financial puzzle we described in chapter , we have to be aware of the possibility that these features are embedded in the initial specifications of the models: the autoregressive dividend function and that the fact that agents' demands (and consequently future prices) are a function of the agents' forecasts of future volatility (which is simplified to equal past volatility). In the design of our market we have actively looked not to make any assumption that can directly generate the time series properties we are looking for. The authors have written a sequel article, (Lebaron, 2002) in which they draw attention to some of the issues that need to be improved in their market design.

ATOM² is a complex software tool, that offers a *general environment for agent-based simulations of stock markets*. ATOM is maintained by a joint research team at the Lille University of France³ and from the University Paris 1 Pantheon Sorbonne⁴. Unlike other agent-based simulators we present, ATOM is more generic and can be adapted to a different range of research applications. ATOM has the possibility to simulate stock markets with market makers or with limit-order books. Moreover ATOM was designed to accommodate order books for more than one asset, making it one of the few current simulators with this capacity. From their collaboration with the NYSE-Euronext, the ATOM creators have duplicated the Euronext pricing methods and are able to create the same price series, as the real stock price series, starting from the market orders of a real stock market. Such pricing possibilities are enhanced by ATOM's capacity of simulating extra-day and intra-day continuous trading periods. A number of financial studies, like (Veryzhenko, 2012), (Derveeuw et al., 2007b) or (Veryzhenko et al., 2010), have been conducted using ATOM. A good review of ATOM's capabilities can be found at (Mathieu and Brandouy, 2012).

In 2001, economist Marco Raberto and his colleagues published an article describing their GENOA artificial stock market, (Raberto et al., 2001). As the authors mention, the aim of this model was to '*offer a simple understanding of the known stylized facts of financial time series, i.e., volatility clustering and fat tails in the distribution of short term returns*'. The model includes agents that don't receive any fundamental information and make offers to sell/buy, in every round, with a probability of P/(1-P). Orders sent to the market are limit orders with prices based on a formula like

$$P_i = p(h) * N_i(\mu, \sigma_i) \quad \sigma_i = k * \sigma(T_i) \quad (4.2)$$

where $p(h)$ is the stock's price at time t and $\sigma(T_i)$ is the stock's historical volatil-

² Available for limited testing at <http://atom.univ-lille1.fr>

³ Team SMAC Multi-agent Systems and Behaviors

⁴ <http://www.univ-paris1.fr/>

ity. At each price level, agents have different demands to buy or sell stocks. The resulting price is discovered at the intersection of the demand and the supply curves. The model exhibits, under certain conditions, the stylized facts found in real market prices. The results are not robust to the model size: a great number of agents eliminate the volatility clustering feature. As it was the case with the previously presented Santa Fe simulator, the GENOA market makes assumptions (limit order prices computed using the market historical volatility) that include the expected results.

The builders of the GENOA market have made improvements to their simulator, in two new versions GASM-2 and GASM-3. In the last version, authors investigate the relative strength of different agent strategies (based or not on fundamental economic information). They draw an interesting conclusion, that a strategy's success depends, on reasonable time scales, on market conditions.

In (Brock and Hommes, 1998) the authors represented a model of an Adaptive Belief System. The model uses a risk free asset with a constant return and a risky asset that pays a dividend with IID values. Similar to the Santa Fe example, agents in this model are homogenous in their conditional variance and they all try to maximize a myopic mean-variance function. Different from other models, these agents have different conditional expectations of fundamental economic information that depend not only on the economic information they perceive but also on their own past market errors (differences between the perceived correct value and the market realization). In this framework of beliefs, agents can be viewed as belonging to different types of common strategies (e.g. fundamentalists, which believe that prices return to fundamentals, or trend-chasers that believe prices than to trend). The adaptive component resides in the possibility of an agent to choose to switch between strategies of trading (fundamental to non-fundamental) in order to maximize his wealth. The results provide some clues to how market efficiency can be attained. In a simple case of homogenous beliefs the simulated market is perfectly efficient. When beliefs are heterogeneous departures from correct pricing can be persistent and it is not obvious that such departures can disappear in a relevant period, or as the authors explain it: "the answer to this question seems to be not as obvious as one might have guessed". It can be argued that since the market is not complete (fundamental agents cannot borrow money and sell stocks short) that arbitrage are limited. This argument can be made in the favor of any trading strategy. If a group that holds a certain belief about future prices can borrow and trade to move prices towards their belief than their strategy will be powerful (at least while their supply of capital endures). Such a strategy will survive in the market and not necessarily the strategy that is "correct".

Another interesting market simulation model has been proposed by in (Matassini and Franci, 2001). The authors model a market with a single risky asset that has no fundamental value model. Trading is done continuously through a limit order book where the best bid is executed against the best ask (if possible). Agents each have

few randomly drawn behavioral parameters: expected gain and maximum accepted loss before selling their positions and a time threshold that prompts an agent to change his idea of the investment (wait to long for an expected gain to realize and it is not profitable). When sending an order, traders compute a price based on the average of three sources of information: peer price expectations, internal information (which assures that a balance is maintained between bids and asks) and historical trends extrapolation. This model is able to reproduce fat tails and volatility clustering. Whilst the model seems realistic, especially with the usage of stop and gain limits, the results are not robust and the variation of certain parameters can completely change the behavior of the simulations.

In (Levy et al., 2000) the authors propose an agent-based model which they call 'microscopic simulation' of a stock market. The authors model the behavior of two types of investors: 'rational investors' (RI) which believe price converges to the fundamental values and 'efficient market investors' (EMR) which believe price already reflects all the fundamental information. All investors are risk averse and want to maximize the expected utility of their wealth. To simplify things, the authors assume that all investors follow a power utility function of the form:

$$U(W) = \frac{W^{1-\alpha}}{1-\alpha} \quad (4.3)$$

where α is the risk aversion parameter, the same for all agents. All agents maximize the expected value of this utility by finding the proportion of wealth to keep in risk-free constant return asset and in the risky asset (which pays a random dividend with a known probability density function). All investors know this probability distribution function but they form different expectations about the future price. Since RI investors believe prices can diverge from fundamentals (but also that these prices will quickly converge back to fundamentals) they compute the future price using the Gordon growth model:

$$E_t(P_{t+1}) = F_{t+1} = \frac{E_{t+1}(D_{t+2})}{r_f - g} \quad (4.4)$$

From equation 4.4 and the known dividend distribution, RI investors compute their expected wealth in the next period. By deriving this equation, investors provide demand levels for each price level. In contrast, EMR investors think assets are correctly priced and there are no possibilities to find '*bargain assets*'. Thus these investors will look only to maintain the optimal portfolio balance between the risky and the non-risky asset. To do this, they need to know the ex-ante distribution of returns. Since this distribution is unknown they construct it using the observations of past returns (ex-post distribution), thus supposing the process that created the returns is stable. These EMR investors can differ in the number of past returns they use to estimate the distribution of future returns and compute the future expected return (but this 'return memory' parameter doesn't influence (qualitatively) the

results). EMR investors can, as RI investors do, also specify demand levels for each price level.

$$\sum_{agent} Demand_{agent}(P_{t+1}) = N \quad (4.5)$$

The market price is found such that the sum of all the agents' demands equals a constant supply of stocks (N), by solving the equation above. With this, rather simple, 'microscopic' specification of a stock market the authors manage to create price series that replicate all of the 'stylized facts' presented in chapter . In a market with only rational investors prices are stable around fundamental values and no statistical anomalies are present. In a more complete market, that includes only 5% EMR investors with a memory of $m=10$ past returns, the resulting price series shows a cyclic price behavior that is similar a price with booms and crashes.




Figure 4.4: Figure from (Levy et al., 2000), chapter 7 page 162, showing the market generated price with 95% rational investors and 5% "efficient market believers".

These price boom-crash cycles appear in a predictable manner:

1. First a series of high positive dividends occurs
2. EMR investors update their estimations of the ex-ante distribution of returns using this new high returns
3. EMR investors increase their portfolio size of the risky asset because of the new

high expected returns, which drives the prices high (like a positive feedback loop).

4. Rational investors observe this non-fundamental increase in price and reduce their share of the risky asset
5. After a few rounds of price increases, EMR investors usually hold all the risky assets. Yet the high returns realized at the starting points a) and b) are not being realized anymore and the EMR investors start to revise their ex-ante distribution expectation using the new observed returns.
6. EMR investors start reducing their portfolio shares in the risky asset because its returns are not high enough anymore. These rebalancing actions generate the important fall in prices.
7. As prices descend to the fundamental value, RI investors start reallocating the risky asset to their portfolio (since it is not overpriced anymore). Afterwards, the price settles and the cycle can repeat itself from point a). During the crash and a few periods afterwards, EMR investors don't touch the risky asset since their ex-ante estimation of the return distribution is highly pessimistic due to the recent price crash.

The same boom-crash cycles can occur in the opposite direction. When very low dividends are realized the EMR investors will expect lower returns and reduce their share of the risky asset thus driving prices down. The rise in prices will follow when new higher returns follow.

This model can be criticized for showing a predictable price series that can, in reality, be arbitrated away. The authors make tests with a population of EMR investors (5% of the total) with heterogeneous memory spans (with a mean of 40 past returns and a standard deviation of 20). The resulting investor mix creates series of prices with bubbles and crashes but with a frequency, intensity and duration that are not (easily) predictable. In this case it is difficult to conceive a strategy that could detect and arbitrage this behavior. The authors also investigate the ratio of wealth between rational investors and efficient market believers and they show that EMR investors survive (and sometimes prosper) in the markets.

In (Farmer and Joshi, 2000) the authors propose a model of a financial market where investors are considered units that process information shocks and diffuse these inputs into prices. Investors can distort information and move prices away from their fundamental levels. The model proposed by Farmer and Joshi uses a market maker that computes the market prices by trying to lower the excess of orders (too many buy/sell orders and the price goes up/down). The investing agents, which populate this artificial market, are divided into two types:

1. Trend followers (or chartists): they base their demand on the idea that the

market price moves in trends. The orders of these agents increase the volatility and autocorrelation of returns.

2. Value investors: They believe prices can stray away from fundamental values and they try to profit off these mispricing. They aim to buy undervalued and sell overvalued assets. Value investors use state-dependent triggers (maximum/minimum price deviation from their perceived fundamental values) for entering and exiting positions.

This model generates volatility clustering because of the trend followers' reaction towards perceived price or return trends. The authors point out that heterogeneous populations of trend followers, which differ in the length of trends they can perceive, are able to generate interesting dynamics that lead to the stylized facts. The authors focus on the mathematical resolution of the price dynamics and often try to provide analytical models that reduce investor behavior to a representative individual. They conclude saying that when non-linearity appears in the investors' behavior than a "representative individual" doesn't exist or can produce a distorted time series. Because market returns have autocorrelations such phenomena cannot be generated by fundamental value drive investors alone, hence markets are populated by other types of investors (which survive). One of the drawbacks of this model is, as indicated by the authors, the fact that the market maker has infinite inventory and is not risk averse.

After reviewing a few notable multi-agent simulators we go through the main methodological and design issues a researcher is faced with when working with simulators. We describe the details and importance of each necessary parts of a simulator. Each designed component is important because it imposes explicit and implicit simplifications and assumptions about the studied financial reality.

4.3.2 Simulated market

One of the first choices to be made in the design of a simulator is in the type of simulated financial market: a stock market, bonds market, stock or bond derivatives market or another type of market. In real life, these markets are not independent and they influence each other. If we simulate only one market we are, in a way, eliminating the influences of other markets. Most researchers defend themselves with the argument that the joint analysis of a bond and stock market calls for a too complicated and extremely difficult to understand model. Consequently, most studies focus on only one type of market (usually stocks) and assume that other markets have no noticeable effects.

We account for the lack of a bond market through the information process. Implicitly, we suppose that the information concerning the underlying asset's economic value and risk is perfect and complete - meaning the company's asset value is correctly reflected in this information. Since most listed companies own financial

instruments (among which we often find bonds and their derivatives) we are actually assuming that bond markets are *efficient* and that they correctly price these financial instruments. Similarly, we choose to ignore the influences of stock derivatives on our simulated stocks. One should make the choice of the simulated market with awareness of the simplification of the real world and should provide a clear statement of the limitations that his choice implies.

4.3.3 Simulated time

The perceived order of a system depends on how abstract its model is. Looking at financial markets on a yearly basis we generally observe price growth and small volatility. Looking at markets on a monthly basis we see a very different dynamic - at least higher volatility levels. Going further and looking at a market on a daily basis we will observe very volatile returns and fast changing trends.

The choice of simulated time is crucial for a simulator's design simply because financial markets have different dynamics and participants on different time scales. We rarely see technical traders on a time scale larger than that of a week. Likewise we don't often encounter "buy and hold traders" on an intra-day time scale. The choice of a simulated time will often impose constraints on the choice of simulated markets (some intra-day bond markets don't even exist), agent strategies (enter/exit timing), exogenous information (daily information has different characteristics than weekly or monthly information), price formation mechanisms and the types of dynamics we can reproduce. This choice of time is important and has to be well defended.

4.3.4 Information

Prices (should idealistically) reflect underlying fundamental values through the process of information integration into price. Therefore the information flow is a crucial component for financial markets. There are two types of information that are available to market participants: exogenous and endogenous information. Endogenous information involves data about the structure of the market itself (participants, securities structure, price formation methods, etc.) and it is usually well known by the majority of investors. Exogenous information represents the data concerning the securities' underlying value. Such fundamental information has different formats and usually reflects a distilled value showing the economic and financial positions of a company (translated in an expectation of future cash flows and return rates).

Most financial market simulators replicate stock markets. The exogenous information for these markets usually consists in proxies for the underlying value (dividends, price/dividends, etc). We believe that a second source of information as a proxy for the underlying risk is essential for a parsimonious representation of a real financial market. Assessing the underlying risk of a stock (the economic risks of a company) is at least as important as assessing its underlying value. Future

cash flows are predicted based on economic data of the company and its business environment. This means that a risk factor is always computed and necessary for a good valuation of the company's stock. The exogenous risk information can be a proxy for different risk factors like liquidity risk, market risk (related to financial assets the company holds) and other business risk. Using two flows of exogenous information (value and risk related) an agent can better assess the confidence interval for the price of a security and correctly discount it. Differences in agent capacities (knowledge and lack of bias) produces more accurate pricing and risk evaluations.

A debate can be made about whether a 'fundamental' value exists or not. In this study we view the 'fundamental value' of a stock as a theoretic construct that has the goal of providing a source of information for the investors.

4.3.5 Bias in information processing

Investors receive information in two forms: quantitative information (financial statements) and qualitative information (letters to shareholders, analyst reports, economic news). Quantitative information can be interpreted fairly objectively by the investor. We have to address the issue of qualitative information that can, and often will, be processed in a biased way.

Each investing agent has his own manner and capacity to interpret information. Some agents use information in its original form while others bias the information in all different manners. For example, agents that are optimists have a tendency to increase the magnitude good information (like fundamental value increase or risk decrease) and decrease bad information. Inversely, pessimist agents decrease good information and increase bad information. A more comprehensive review of investor biases was made in chapter .

When we talk about investor biases we refer to intrinsic characteristics of investors (meaning humans) and not to voluntary actions. The biases, described succinctly above, refer to ways through which investors alter, involuntarily, the information they perceive. These biases can be changed by the investors' learning process. One can argue that an agent cannot change its emotional biases yet an agent can rationally alter the information he uses as the input for his investment strategy. Therefore, we say that the investor's biases can be rationally "unbiased" by making corrections to the exogenous received information. If an investor observes that his eyes have some form of refraction dysfunction (e.g. myopia) he can use external means of corrections (like glasses or contact lenses). Similarly, an investor knows that the information he receives is biased (by him or others) and he can apply a rational process for correcting these biases before using the data for this portfolio decisions. An optimistic investor that expects the growth of a company for next year to be 20% can correct his expectations (provided he knows he is an optimist) to a more 'unbiased' 17%, 16% (or lower).

4.3.6 Risk and time preferences

When modeling an investor with a simulated agent we make assumptions about how investors create their portfolios: how much cash do they hold and how much do they invest in the risky asset. An agent can be modeled as an optimizer of mean-variance, of expected utility maximizing or none of the two. The risk attitude of investors is not simple to model and one should describe more than one characteristic risk attitudes to support the heterogeneity assumption of our methodology. Risk aversion is shown, in (Brandouy et al., 2012b), to be an important factor in the performance of investor's strategy.

In regards to portfolio horizons, investors can be heterogenous from myopic (looking to quickly exit positions) or are ready to hold positions for a longer-term. Modeling investor time preferences is important, as explained in (LeBaron, 2001), and all choices should be clearly justified. A good choice is to permit a mix of preferences, short and long horizon for investment, and observe their co-evolution.

Referring to risk measures most agent-based simulators don't focus on risk and use a classic form of volatility measure. A newly published model, in (Thurner et al., 2012) focuses on another aspect of risk: systemic risk. The authors model an entire 'stylized' economy consisting of informed investors (e.g. hedge funds), uninformed investors and banks that provide credits. The simulated agents send demand orders that are functions of the future price (the higher the price the lower the demand). While the model is interesting, in reference to its purpose, there are some features that make it unrealistic (e.g. the demand of the uninformed investors is a weak-mean reverting function of the fundamental value). As we discussed in chapter , a model should not have as input the proprieties that we desire in the output. In this case, the price is forced by design to revert to fundamental values. The authors draw interesting conclusions about how the relations between leverage, bank precaution regulations and systemic risk. What we should retain is the idea that "systemic risk" may be relevant in pricing and can be taken into account by individual investors. Thus, a more comprehensive market model should take the systemic risk factor into consideration.

4.3.7 Agent learning

The agent learning reflects the learning behavior of real investors. Several methods for learning have been researched and acknowledged as correct. Before explaining some of these methods we want to point out that all markets possess a latent form of learning: evolutionary learning (also called passive learning). By evolutionary learning we express the fact that a system becomes better adapted to its environment, provided that less performing actions eventually disappear from the system. On a micro level, an agent is influenced by the system's evolutionary learning if he will be excluded from the system because of other superior competing agents. In

the case of financial markets, the performance measure is the return rate. Thus, agents that have inferior strategies will, in the long run, lose all their wealth (provided the games in zero-sum or there are periods with economic downturns) and will eventually be eliminated from the market. We say that a financial market has a form of evolutionary learning because it eliminates underperforming investors. The necessary conditions for this evolutionary learning are:

- The system is relatively closed: not many new, and less 'adapted' investors, enter regularly the market. These new-comers become a prey for the more 'adapted' investors.
- Trading is frequent. Low volumes don't permit exchanges of wealth from less to better performing strategies, at least not on decent time scales.

In the long term, if evolution works as expected, a market will reach a point where all the agents will have the same well performing strategies and all trades will cease (since no agent will be willing to buy high / sell low). This consequence is actually implied by the efficient market theory. This passive method of learning doesn't require investors to actively adapt to the market because only the prevailing strategies will persist in time.

The other type of learning requires voluntary action from agents. This type of learning is called active learning because it requires agents to actively improve their strategies for investing. In the case of a simulated financial market there are several questions a researcher should answer in order to equip his agents with the capacity for active learning. The first question should be about the nature of the objective function for the agents. Should they measure utility of wealth, value of wealth according to prospect theory or their wealth relative to mean wealth of the population? Should investors measure period returns relative to market returns (or non-risky returns) or should they just consider their absolute returns? Each of these objective functions implies and produces different dynamics in the markets. This choice is therefore highly relevant and important for the construction of a market model.

Once the investor objective function is set we look at how agents actually learn. There are basically two forms of active learning: innovation and imitation. Firstly, imitation implies that an agent will observe and imitate the strategies of better performing agents. Underperforming strategies will be discarded much faster and evolution will take its course. Secondly, innovation implies that agents will change the parameters of their strategies (information bias parameter, risk bias, risk aversion, trigger level for selling/buying, frequency of trading, etc.). This process involves even faster learning and adaptation of competing strategies. When all (or most) agents have this learning ability the simulated financial market becomes an ecosystem of fast evolving and competing strategies. Real investors exhibit both forms of active learning. The combination of these two types produces interesting dynamics in a simulated financial market.

4.3.8 Agent strategy bias

Markets are driven by information and extreme information can produce extreme effects in financial markets. The human brain has a part, called brain stem, which is considered a primitive brain. The brain stem can inhibit the cortex (the big rational brain) and take lightning fast decisions and send commands to the locomotory system (like jumping out of the way of a car or avoiding a flying object). These reactions are triggered by extreme emotions which are in turn caused by the detection of an impending grave danger.

Extreme information that arrives in financial markets can, and often does, cause erratic behavior in the market. Extreme market movements, in the form of more than $3 * \sigma$ deviation from historic returns are seen more often than we expect from the predictions of classical financial theory (see leptokurtic returns, chapter 3.6). The reason for these market outcomes is rooted in the "not completely" rational and repeatable behavior of investors.

Therefore we propose that agents, in simulated financial markets, should possess some form of trigger for rash and fast reactions. These reactions are caused by exogenous factors: lack of information, extreme information, rumors or by endogenous factors: loss aversion level, high compulsory consumption, end of investment period. Most models assume that if investors are biased or not completely *rational* they will behave like this in a consistent and predictable manner. Such extreme behavior can better capture the unpredictable and indeterminist behavior of investors. Of course, if this behavior is underperforming it will be eliminated through evolution. Therefore it is difficult to argue that introducing such an extreme behavior feature can disturb market efficiency for a sufficiently long period of time.

4.3.9 Price formation

The price formation mechanism is one of the principal components of a simulated financial market. Real markets use order books that match corresponding transactions and offer information about price distribution during a time period. Simulated markets use a variety of mechanisms that can simplify the usage and tractability of their models. (Bottazzi et al., 2005) make a good comparison about market organization and how this can affect price formation. We adhere to the opinion that simulators should use true order books for price formation (as described in the (Epstein, 2006), page 1194). We make a brief introduction into the main models used in price formation.

One price formation mechanism consists in a market maker that announces prices at which agents send buy/sell orders. Consequently, the price is adjusted progressively to balance the excess demand or supply, as observed in (Farmer and Joshi, 2000). Another mechanism consists of agents meeting randomly and trading assets for an ad-hoc agreed price (see (Beltratti and Margarita, 1992) for an example).

Although these two mechanisms simplify the agent models they also eliminate some microstructure features that are presently studied in real financial markets. The most realistic method is the one described in (Chiarella and Iori, 2002) or (Raberto et al., 2001). The model uses a limit-order book to resolve transactions and find market prices (see chapter 3.2 for details about price formation in real financial markets).

The dynamics of different types of price formation mechanisms have been studied by many researchers such as (Raberto et al., 2005). In the three versions (GASM 1, 2 and 3) of their GENOA simulator, Raberto and Cincotti show how different types of price formation mechanisms can impact the efficiency of markets. Even if this GENOA simulator is often quoted when providing evidence of 'stylized facts' and other findings, some of these versions (notably GASM 1) make assumptions on the behavior of investors that implicitly generate some empirical proprieties (e.g. all investors generate orders stochastically as a function of historical volatility).

4.4 Critics of agent-based artificial stock markets as a methodology for research

Agent based simulators were first introduced in the 1980's and they provided the means to model systems using a "bottom-up" approach: model the building blocks of a systems and observe its emergent aggregated level proprieties. Like with any new approach, this method was meet with a lot of critic and disapproval, especially from the defenders of classical research methods (either experimental or pure theoretical). In this chapter we review and address the main critics for our methodology.

1. *"Simulations are very abstract models of a reality and they cannot capture the reality and complexity of financial markets"*.

By definition, a model is a simplification of a phenomenon. A multi-agent simulator doesn't try to capture the entire complexity of a market. Contrary, an agent-based model aims to create an abstract representation of a market that behaves, in aggregate, in a similar manner as real markets. Moreover, the language used in describing financial markets (computer language) permits the modeling of complex rules and behaviors (which is extremely difficult to do using mathematical equations). Different from other approaches, agent-based modeling makes assumptions on a micro-level basis (e.g. investor behavior) and permits the evolution of such an environment in order to observe the emergent macro-level patterns. Therefore, complex and unexpected behavior can emerge from a simple-defined simulator.

2. *"Agent-based simulations are computational black-boxes and their designers can tweak the model in order to obtain the results they need. Therefore, the results discovered are not more than what was already built inside the model"*.

This critic is perfectly true for all types of models. The important difference is that researcher "tweaking" can be better hidden in a computer program (at least when it has an important volume of computer code) than inside a mathematical function. To counteract such 'visibility' problems, agent-based simulators have a simple and powerful propriety: repeatability. A computer simulation program can be made public and tested easily by researchers from the entire world. By making the source code public, the specification of the model can be reviewed. Anybody can judge if the results are somehow embedded inside different parts of the model design. In opposition, classic models cannot always be tested (by other researchers) because the analyzed data is confidential or the methodology used is highly complicated.

3. *"Because simulations are numerical we cannot use them to extract general theories."*

The purpose of agent simulator is closer to testing theories rather than constructing new ones. In finance, an agent based simulator can be used to see if, in the long term, the survivors from a mix of investors are only the ones that respect the assumptions implied by a certain theory (e.g. CAPM theory). As argued in chapter , our methodology offers the means to test the attainability of the theoretical valid 'market efficiency' state. Therefore, rather than extracting new theories we provide better way to test existing theories.

4. *'The results are merely chance and represent ad-hoc effects'.*

It is possible that a certain result can be generated from a very specific (and maybe unrealistic) initial world setting. To counter this problem robustness tests are employed. By varying input parameters (e.g. agent behavior mix, number of investors, time frame) we can see if the observed result persists or not. Moreover, tests can be repeated and confidence intervals can be measured for every output variable.

5. *"Agent based models are not mathematical equations for which there are numerous rigorous methods for testing"*

A computer program is perfectly equivalent with a mathematical function. Therefore a computer program has the same proprieties, and advantages, as a mathematical function. More than a classical mathematical function, an agent based model can represent sets of behaviors that are not always tractable. Furthermore, complex models (like investor behavior) are easier to express using an agent-based model than with a set of mathematical functions.

6. *"Agents are not deductive."*

Acknowledging point 5, we say that an agent's action is similar to the application of a mathematical function on a set of states (the agent and the environment states). Thus an agent's actions and all the simulator's results are deductions. Moreover, the transformation of input data (economic information) into output (market prices) data is strictly deductible and reproducible. We can rerun the same simulation with the same input parameters and obtain the same results each time.

4.4.1 Equivalence between an agent-based artificial market and a mathematical analytical model.

Inside a computer market simulation, agent behaviors follow simple rules that deduce an action (order) from a set of environment states (information, states of other agents, market prices). Such behavior is encoded inside a computer program using computer language. As proved by Alan Turing (see (R.E, 1995)) for every computer program it exist an isomorphism with a mathematical analytical function. Building on this fact we show how, at an abstract level, an entire agent based simulator is equivalent to a mathematical model. Like many others systems, e.g. weather system, such a mathematical model cannot be always tractable and it is often estimated with computer approximations.

Let f be the function that defines the behaviors of an individual agent i as i :

$$f_i : A * M \mapsto Order \quad (4.6)$$

A represents the set of all the possible states of agents (portfolios, objectives, risk view and other behavioral considerations). M represents the states of the observed market as well as information about next states (historic prices, volumes, information, analysis, etc).

The market evolves from one state to another and this evolution is modeled by a function m defined as:

$$m : M * Order^n \mapsto M \quad (4.7)$$

where $Order$ represents the complete set of possible orders that are sent to a market and n represents the maximum number of orders a market can process between two states. Such a function exists since the process of price discovery is deterministic and it depends only on the received orders and the last market state. After a change in the market states the agents' states change (their wealth evolves). This change is modeled using a function g with parameters that depend on each individual agent i :

$$g_i : A * Order * M \mapsto A \quad (4.8)$$

Therefore, an agent-based simulation is equivalent to the calculation of a set of equations like

$$\begin{aligned}
 &A_{1..n}^0 - \text{agent initial states} \\
 &M^0 - \text{market initial state} \\
 &M^t = m(M^{t-1}, f_i(A_i^{t-1}, M^{t-1})) \\
 &A_i^t = g_i(A_i^{t-1}, M^t)
 \end{aligned}
 \tag{4.9}$$

The first two equations are the specifications of the market (which includes economic information) and the agent's endowments (wealth and behavior). This part equates to the specifications of the modeled market and the assumptions we make about agent's behaviors. The last two equations represent the co-evolution of the market and agent states. These equations are developed using mathematical language only if the behaviors of agents (and of the price formation) are easily translatable into equations (which is often not the case). Therefore, a computer simulation allows us to compute the series M^t (market proprieties: autocorrelation, volumes, prices) and A^t (investor proprieties: cash, stocks, risk aversion, investment horizon, etc) and observe their statistical proprieties and financial meaning.

Contribution - Our simulator (LUMA)

Contents

5.1	Simulator, environment and assumptions	79
5.2	Information and fundamental value	81
5.3	Orders and the price formation mechanism	84
5.3.1	Orders	84
5.3.2	Price formation mechanism	86
5.4	The agents, their biases and investment decisions	91
5.4.1	The agent	91
5.4.2	The bias	93
5.4.3	The investment decision	94
5.5	Review of configuration parameters	102

*LUMA is available online (along with detailed instructions) at
<http://cerag.org/lucianstanciu/LUMA/index.html>

5.1 Simulator, environment and assumptions

The computer-based financial market simulator we have created is called LUMA. Our simulator's objective is to simulate a series of realistic daily stock **market opening prices**. In other words, we simulate the **extra-day price series** in a stock market. This choice of trading time is justified by two reasons: 1) price series of intra-day continuous trading incorporate much more noise and less information about fundamental values than during extra-day trading (statistical properties of these intraday price series are not especially relevant for our research questions) and 2) most financial research studies, especially regarding price anomalies and behavior biases, have daily price series as reference.

LUMA replicates a financial market formed by two assets, which are accessible to all investors: a risk-free financial asset with a 0% return and a risky financial asset that, for the duration of a simulation, pays no dividends. The choice of a 0% bond is made in order to simplify the model and it has no impact on the price

dynamics. One can consider that an investor's cash does not accumulate interest and the only possible gains are capital gains from investing in the risky asset. This choice of assets does not limit our results and it has been used in other published studies like (Farmer and Joshi, 2000).

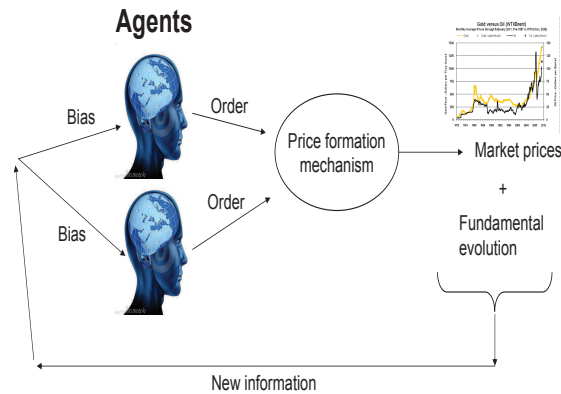


Figure 5.1: Flow of information inside the LUMA simulator

The investors (called agents) that are modelled in LUMA receive a flow of biased information. The information is biased in a way specific to every agent. The information received by agents is used to make investment decisions (according to each agent's subjective considerations). Following the investment decision, agents send orders to the market to buy or sell the risky assets (see Figure 5.1). After a short period of time, during which all agents have the chance to send their orders to the market, the trading period ends and the new market price is discovered. This process lasts for a predetermined number of rounds (unknown to agents) and the resulting series (price, returns, volatility) are analysed.

Agents start a trading simulation with a customizable amount of cash, number of risky assets, information biasing properties and investment decision methods. All agents in LUMA have proprieties that do not change in the course of a single simulation (except for their level of cash and risky assets). This choice implies that investors do not actively adapt (learn) to the market. Although it is true that no particular investor changes strategy (or other parameters) during a simulation, we observe that a form of "Darwinian evolution" takes place. Underperforming investors lose their cash to better performing investors (see chapter). In the next chapters we explain in detail every component of the LUMA simulator and we

conclude by showing the results of a few basics test of the concept.

5.2 Information and fundamental value

According to the standard modern finance definition, the price of an asset is defined as the discounted sum of all future payments that are received from the asset. To compute this price an investor holding the asset must have information about these future payments, their amounts and dates of arrivals. Moreover, the investor must also know the correct discount rates for each of the future cash flows. To simplify this informational problem, without eliminating important features or reducing the model's applicability, we have simulated an asset, a company's stock, which pays no dividends. Thus, investors can only make capital gains and have a single piece of information to discover and/or predict - the fundamental value. Since the cash owned by investors doesn't offer any return, it is 'rational' for them to invest all of their cash in the risky asset (of course, when they consider it is profitable). And in order to invest, agents have to discover how the fundamental value of the asset has evolved. They do this by interpreting, via their own personal biases, an economic information stream. We call this stream of economic information the fundamental value of the asset.

The fundamental value has a normal distribution of returns and it is generated, for each particular simulation, using the formula below and the two parameters that describe the fundamental's return and standard deviation:

$$V_t = V_{t-1} * e^{N(\mu_V, \sigma_V)}, \forall t > 0 \text{ and } V_0 = M = P_0 \quad (5.1)$$

where M is the starting IPO price of the stock, μ_V is the average return and σ_V the standard deviation of the fundamental's return values and P_0 is the first market price (fixed exogenously similarly to an IPO pricing).

This fundamental value is, like in real markets, not visible to investors. We say that this value is not observable at least because of the sheer complexity of a listed company as well as the great number of remote entities/events that can affect, in every moment, the value of a company. Therefore, in computing this economic value a number of assumptions and simplifications are made which render this informational stream not directly visible to investors. There are agents that prefer not even to seek this fundamental value and choose only to play 'the market'. Yet for other agents, that believe prices should reflect the asset's fundamental value, a stream of information is needed.

We have chosen an information stream, I_t , that offers a signal about the 'distance' (in price terms) between the fundamental value of the asset at the end of a trading period and the market price at the beginning of the period:

$$I_t = V_{t+1} - P_t \quad (5.2)$$

Since investors try, at least in theory, to estimate the 'real' value of a stock we can say that P_t represents the market's estimate of the asset's fundamental value at time t . We can rewrite formula 5.2 as:

$$I_t = V_{t+1} - E_{m,t}(V_t) \quad (5.3)$$

where $E_{m,t}$ stands for the market's (m) expected value, at time t , of the fundamental value of the asset at time t . Looking at 5.2 we say that the information available refers to the amount of change of the fundamental value from time t to $t+1$. In other words, this information refers to how much the real fundamental value has evolved in relation to the market's latest estimation of the fundamental value. In real markets, 5.4 looks more like

$$I_t = E_{a,t}(V_{t+1}) - E_{m,t}(V_t) \quad (5.4)$$

where $E_{a,t}$ represents the best combined estimates of analysts, economists, etc. of the company's value. In real financial markets we have only estimates of fundamental values that are either market-based (through price) or non-market-based (analyst, research and individual opinions). Because such fundamental information is not very useful in itself we have chosen to model the information stream received by investors as a series of negative and positives signals related to the market price.

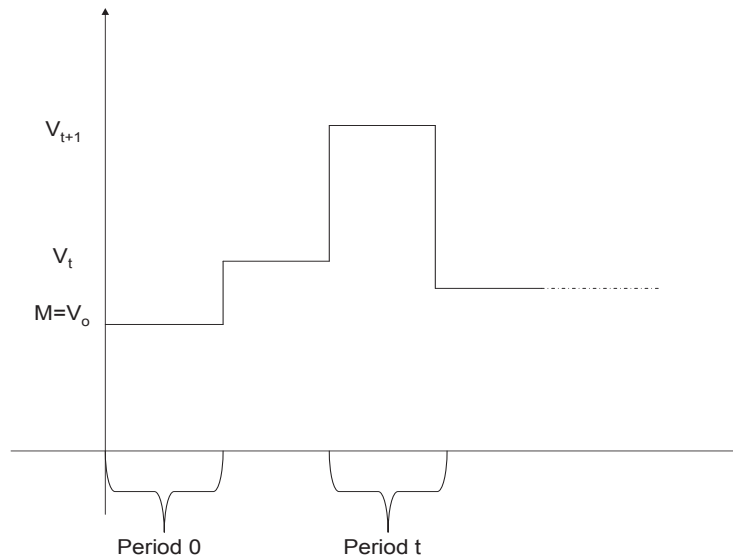


Figure 5.2: Trading periods and information flows

At the beginning of a trading period t the fundamental value of the asset moves instantly from the value of V_t to V_{t+1} . Agents can observe the latest market price P_t

and also the new information I_t that tells them how close/far is this latest market estimation (P_t) from the new fundamental value of the asset V_{t+1} . During this trading period t , an agent x will interpret (through their own personal beliefs) the economic information ($I_{t,x}$) and use his individual investment strategies to send orders to the market. At the end of the period, the new market price, P_{t+1} , should in theory be equal (or very close) to the new asset's fundamental value V_{t+1} . It is important to underline that at the beginning of period t , the fundamental value of the asset instantly changes from V_t to V_{t+1} . Only after this change occurs the investors receive the new information stream I_t . Therefore, during trading period t the market will try to estimate the new asset's value, V_{t+1} .

Looking at the structure of the information signal, it is easy to see that an investor can simply add the signal I_t to the last known market price P_t and obtain the new fundamental value V_{t+1} . While this is true, research in psychology has showed that humans can have very different perceptions of the same objects. Therefore, the same information stream I_t can be *perceived* (or interpreted) differently by different agents. We refer to an information $I_t < 0$, representing a decrease in fundamental value, as 'bad news' and similarly an information $I_t > 0$, showing an increase in fundamental value relative to the last market price, will be referred as 'good news'.

To understand better the information stream that agents receive we show in Figure 5.3 the co-evolution of a fundamental value and the market price in a market simulation. We see that the market starts at the same level with the fundamental value and afterwards follows the economic value more or less closely.



Figure 5.3: Fundamental value and market price for an asset and the perfect information stream

Our choice of asset information enables investors to focus only on the fundamental value of the asset. Of course, like in any model, this is a simplification which is rarely found in real markets where assets and asset classes are interlinked and their pricing depends on much more than their fundamental value. Therefore we suppose that the main information stream available, the traded asset's fundamental value, encapsulates not only asset-specific information but also macro-environment information.

The underlying objective of our study is focused on the ways in which information, which describes the value of an asset, is converted and transformed, through the investment decisions of biased agents, into a market price.

5.3 Orders and the price formation mechanism

5.3.1 Orders

During a trading period, agents use economic information and send limit or market orders to the stock market. These orders are stored in a limit-order book until the end of the period. At the end of a trading period, the limit-order book is closed and the price discovery process begins. As pointed out by Blake LeBaron in (Tsfatsion, 2006a) the limit-order book provides the most realistic method for price formation.

We first present the structure of an order and then show how the new price is discovered from a limit-order book.

After receiving and analysing new information, agents send orders to the financial market. An investment order has four parameters like in the table below:

The first parameter, **Agent ID**, represents an identifier that uniquely describes each agent. Using this information, the market can resolve matching orders and makes the correct transfers of cash and assets between the counterparts of a transaction. The second parameter, **Direction**, signals whether agent wants to buy or sell assets. This parameter is also relevant for the interpretation of the third (and fourth parameter) that has different meanings depending on the direction of the order. The third parameter, **Number of assets**, indicates the maximum desired number of assets the agent (identified by Agent ID) is willing to buy or sell. This implies that the agent involved in the transaction can have his order executed with the exchange of at most this number of assets or of a smaller number of assets. For example, the agent with the ID 3, see , wants to buy 10 assets. This order can be executed with the agent receiving 1, 2, 3 up to 10 assets.

The fourth and last parameter, **Max/Min amount**, indicates the maximum/minimum price at which the agent is willing to buy/sell assets. Looking at the order listed in we see that the agent with ID 5 is willing to sell 15 assets from his portfolio at a minimum price of 8 for each asset. If the order is executed, the agent is guaranteed a price of at least 8 for each asset sold.

In real markets, investors have now numerous possibilities to customize their orders like:

1. **Trigger orders** - at a certain price a new order is created to buy/sell assets at the best price possible. These orders are used to limit losses and to cash in on gains, quickly and safely.
2. **Quantity orders** - they are similar to our limit orders and they have a quantity restriction that forces the order to be executed for the entire specified quantity (buy/sell all or cancel the order).
3. **Strategic orders** - combination of different orders like: buy or sell when price goes up x%, alternative orders that specify buying at a certain price or selling at another price (and only one of these orders is executed).

These complex orders can be achieved through the combination of different simple orders. They are conceived to address specific needs of investors and are focused

Agent ID	Direction	Number of assets	Max/Min Amount
3	Buy	10	7.5
5	Sell	15	8

Table 5.1: The informational components of a limit order

on intra-day trading operations. Due to the nature of our simulator, extra-day trading, the orders we use (limit or market orders) are sufficient to create realistic market dynamics.

5.3.2 Price formation mechanism

During each trading period agents send orders to the market. All these orders are recorded in a limit-order book. At the end of the trading period, the limit-order book stops accepting new orders and a procedure for determining the new asset price is started. After the new price is determined and disseminated to agents, all matching orders are resolved and the remaining orders are discarded, leaving the limit-order book prepared for a new trading period.

The new market price is determined using a four step approach. Every step has a specific objective to accomplish. At each step we choose among the prices selected at the previous step those prices that satisfy the new objective. Provided a step offers more than one price that satisfies its objective we use this list of prices and go to the next step. If we arrive at a single price quote, the price formation stops and this single price is used as the new market price. If we reach the final step of price formation having more than one possible price, we use a special fixing procedure using as reference the market price from the previous period.

At the end of each period we discover the new market price using the following 4 steps:

1. **Step 1** will select prices that maximize the volume of executable transactions. At a price level, the volume of transactions is given by the minimum between the cumulative buy and sell quantities demanded. We compute the maximum volume of transactions and we search prices that accomplish this volume. If there is only one price that achieves the maximum volume, we select this price as our new market price and the process of price discovery stops. If there is more than one price that allows for the same maximum volume of transactions we hold this list of prices and move to the next step.
2. **Step 2** will select the prices that minimize the unfilled or unmatched quantities demanded for selling and buying. For all prices available at this step, that allow for a maximum executable volume, we compute the minimum level of surplus. The level of surplus is the absolute difference between cumulative buy quantities and cumulative sell quantities at the same price level. Using the minimum surplus level we eliminate prices that have a surplus that is higher than this minimum. We are left with a list of prices that maximize executable volume and minimize the surplus of unresolved orders. If we have only one remaining price, we choose this price as our new market price and the process stops. If we have more than one price we move to the next step.

3. **Step 3** has the goal of finding the price that agrees with the market buy/sell pressure. The existence of market buy pressure, at a price level, is indicated by the positive sign of the surplus of orders. If the Cumulative Buy Quantity is higher than the Cumulative Sell Quantity we can say that the market exerts buy pressure at this price level. For every remaining price from Step 2 we observe their respective market pressures. If only buy pressure exists, meaning that we have a positive surplus at every price level, we will select the highest price level as our new market price. Since all this residual demand for buying will increase the price in the next period, we will select the highest price available in order to minimize the effect after price formation. Alternatively, if the market pressure is for selling at every price level, we choose the lowest possible price in order to minimize future drops in price after the new price formation. If one price is chosen, the price discovery stops here. In the case of the existence of both buy and sell pressure, or if all surplus is zero, we cannot choose a price because the remaining surpluses (if any) of buy/sell orders are identical in magnitude (they all satisfy Step 2) and different in sign and so can bias the price in any direction. In this case we go to the final step.
4. **Step 4** will choose, among the prices from Step 3, a price that is closest to the reference price. In markets where intraday trading occurs the reference price is the price of the last intraday transaction. We use as reference price the last market price, meaning the price for the previous period P_t . Next, we narrow the list of prices available to just only two prices, using these principles :
 1. If all the prices from Step 3 show a surplus of zero, we will hold the minimum and the maximum price from the list of available prices.
 2. If we observe a mix of positive and negative minimum surpluses, we will hold the two prices that mark the changing in sign. At this point in Step 4 we are left with only two prices P_1 and P_2 (we assume $P_1 > P_2$) and the reference price P_t . We select the new market price P_{t+1} using the equation

$$P_{t+1} = \begin{cases} P_1 & : P_t \geq P_1 \\ P_2 & : P_t \leq P_2 \\ P_t & : P_t \in [P_1, P_2] \end{cases} \quad (5.5)$$

This method of price formation ensures the best price for every participant by maximizing exchanged volume and avoiding to incorporate new biases into the new price P_{t+1} .

After computing the new market price, the simulator resolves all the orders in the limit-order book at this new price. After all possible transactions are resolved a new trading period can begin. To better understand how a price is discovered we will go through a complete example. In 5.2 we can see an example of a limit-order book.

Agent ID	Direction	Number of assets	Max/Min amount
3	Buy	10	12
4	Buy	7	11.5
1	Sell	15	11
2	Sell	7	13
5	Buy	5	12.5
6	Sell	15	11.75
9	Buy	4	14
10	Sell	4	10

Table 5.2: Example limit-order book

Agents have sent 8 orders to the market. The best buying price is 14, for 4 assets, and the best selling price is 10 also for 4 assets. Following the first step in the price formation mechanism we compute the buy/sell quantities at each distinct price and then we discover the cumulative buy and sell quantities.

BUY		SELL		
Cumulative Buy Quantity	Buy quantity at price	Price	Sell quantity at price	Cumulative Sell Quantity
5	5	12.5	0	41
12	7	11.5	0	41
12	0	15	0	41
16	4	14	0	41
16	0	13	7	41
26	10	12	0	34
26	0	11,75	15	34
26	0	11	15	19
26	0	10	4	4
26	0	7	0	0
26	0	5	0	0
26	0	4	0	0

Table 5.3: Market price determination - computing cumulative quantities at all price levels

As we can see in Table 5.3, most assets can be sold at a price of at least 13 and the most assets can be bought for a maximum price of 12. We compute the maximum executable volume for each price level, as indicated in Step 1 of the price formation procedure. The maximum executable volume indicates the number of assets that can be exchanged after discovering the new asset price.

BUY		SELL	
Cumulative Buy Quantity	Price	Cumulative Sell Quantity	Maximum Volume
5	12.5	41	5
12	11.5	41	12
12	15	41	12
16	14	41	16
16	13	41	16
26	12	34	26
26	11,75	34	26
26	11	19	19
26	10	4	4
26	7	0	0
26	5	0	0
26	4	0	0

Table 5.4: Step one of price determination - maximum volume criteria

The resulting prices, in Table 5.4, according to the maximum volume criteria are 11.75 and 12. Because we cannot make an unbiased choice between these two prices we continue with Step 2. In this step we compute the minimum surplus, meaning the number of assets that will not be exchanged, at every distinct price level.

BUY		SELL		
Cumulative Buy Quantity	Price	Cumulative Sell Quantity	Maximum Volume	Minimum surplus
5	12.5	41	5	-36
12	11.5	41	12	-29
12	15	41	12	-29
16	14	41	16	-25
16	13	41	16	-25
26	12	34	26	-8
26	11,75	34	26	-8
26	11	19	19	7
26	10	4	4	22
26	7	0	0	26
26	5	0	0	26
26	4	0	0	26

Table 5.5: Step two and three of price determination - minimum surplus and market pressure

After computing the exchange surplus, as can be seen in Table 5.5, we still cannot discriminate between the two remaining price choices (11.75 and 12) because these prices have both the same minimum surplus level (8 assets remain untraded). Therefore we move on to Step 3 where we look at the market pressure. We observe that the market pressure for both price choices is for selling (at the price levels 11.75 and 12 we have 8 more assets that are available for selling than for buying). According to the price formation mechanism we choose the lowest of the available prices (11.75) thus ensuring the best price and the least bias for the next trading period. The new price, 11.75, now serves as a basis for the execution of all compatible orders. The orders are filled and the resulting exchanges are listed in Table 5.6.

Agent ID	Operation	Quantity exchanged	Cash received	Counterparties Agent IDs
3	Buy	10	-117,5	1
4	Buy	7	-80,5	1,6
1	Sell	15	176,25	3,4
2	Sell	0	0	not filled
5	Buy	5	-57,5	6
9	Buy	4	-47	6
6	Sell	11	825	5,9 (partly filled)
10	Sell	0	0	not filled
	Assets Exchanged	26	Assets not exchanged	8

Table 5.6: Transaction resolution after price discovery

We see that most of the agents have their orders resolved. For example, agent 3 had initially an order to buy 10 assets at a maximum price of 12 (see Table 5.6). The agent has his order fulfilled at a price of 11.75 and he obtains 10 assets for a total sum of 117.5. His counterparty, agent 1, also completely fills his order. Agent 1 sells 10 assets to agent 3 as well as 5 assets to agent 4 thus receiving 176.25 in cash.

We point out that with this price formation mechanism, market orders do not have any effect on the future price. From one point of view this is normal because market orders don't provide any economic information for setting the asset's price. From a market pressure point of view it is arguable that large market orders can influence prices by rapidly going through limit orders in one or another direction. Because of the nature of our economic and financial environment, this effect is ignorable. In our simulator, market orders have the purpose of providing liquidity to the market without having any direct influence on prices.

5.4 The agents, their biases and investment decisions

5.4.1 The agent

There is no market without buyers and sellers. Even when information is missing markets can function provided that investors are willing to buy and sell. So, we assert that investors are the basic and most important building blocks of a financial market. Consequently, the investor's representation inside a computer based simulation, in the form of an agent, is maybe the most important characteristic of part of our model.

An agent is a computer model of a real life investor. Like any model, our agent

is a simplified version of an investor. An agent is a computer program that receives and interprets a stream of information. Using this information, and his personal characteristics, the agent makes an investment decision and sends an order. In our model, an agent receives, at the start of each trading period, two informational streams about the new economic developments (I_t) and the last market price (P_t). In accordance to the theory behind behavioural finance, (Kahneman, 1973) and (Kahneman et al., 1981), our agents have cognitive biases which distort the economic information they receive. Thus, agents will not observe the perfect information stream I_t (see equation 5.2) but a biased interpretation of this informational stream, $Bias_x(I_t)$. For example, "optimist" agents will overestimate information about the fundamental value's growth.

For an extra twist of reality, we created a mechanism through which agents can receive delayed information. Instead of the latest economic information I_t , some agents can receive out-dated information I_{t-k} . This feature models the fact that not all investors have access to the most up-to-date economical information about an asset. Therefore an agent can receive the latest news ($k = 0$), a day old news ($k = 1$) or even older ($k > 1$).

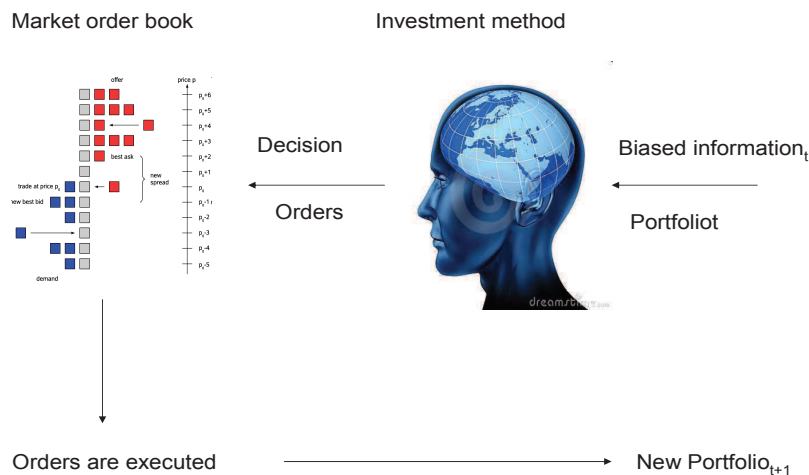


Figure 5.4: Information interpretation for investment decision making

After interpreting (biasing) the information streams, an agent will use his investment strategy to make an investment decision. Using this decision the agent will make an investment action and choose between buying, selling or holding his assets. Because cash does not provide a return, the agent is not worried about an optimal repartition of his wealth between a risky and a non-risky asset. A buy or

sell decision will prompt the agent to send an order to buy stocks to the extent of his cash holdings or to sell his entire stock portfolio.

At the beginning of a test, the simulator will create a mix of agents that are characterized by a set of parameters. Agents start trading the asset. Agents which use all their cash or sell all their stocks cannot send any more orders to buy respectively sell stocks. Consequently, agents that exhaust both their cash and portfolio of stocks will no longer be able to send orders to the market. We do not allow for short-selling or borrowing since the price effects of actions supported by such cash sources will cancel each other out.

5.4.2 The bias

The personal characteristics of an investor, called biases, are simulated in our environment using an affine transformation of the economic information. Thus, an agent P , with personal characteristics a_P and b_P will interpret the economic information I_t as

$$Bias_p(I_t) = a_P * I_t + b_P \quad (5.6)$$

Different combinations of an agent's bias parameters represent well known behavioural characteristics of investors (as can be seen in in 5.7).

Nr	A	B	Characteristic
1	1	>0	Pure optimistic (decreases bad news and increases good news)
2	1	<0	Pure pessimistic (increases bad news and decreases good news)
3	A=1	0	Perfectly Informed, not biased
4	A>1	*	Optimistic with good news and pessimistic about bad news
5	1>A>0	*	Moderated, >Underestimates good as well as bad news
6	-1<A<0	*	Contrarian, Moderated
7	A=-1	0	Contrarian, Perfectly Informed
8	A<-1	*	Optimistic with bad news and pessimistic about good news

Table 5.7: Bias parameters and associated behavior trait

We observe that the two bias parameters of each agent, A and B, play different roles in characterizing the agent's behavioural traits. Parameter A represents a characteristic of inflation/deflation of economic value. An agent with $A > 1$ (or Nr. 4 in the table above) will overestimate the information, regardless of its sign.

In a market dominated by this type of agent, market prices fluctuate around fundamental values because agents will make exaggerated price corrections, based on their overestimated information. Thus we say that, apriori, the 'A' parameter has an influence on price volatility, or more precisely on price excess volatility.

The second parameter B, that models the way investors interpret information, implies a vertical shift in the $A * I_t$ line. For a positive B value, see 5.7, an agent will be optimistic about all news, meaning that it will enhance positive information with a B amount and it will decrease negative information with the same amount. Inversely, a negative B value will have the effect of decreasing positive information and increasing bad information thus rendering an agent to be 'pessimistic'.

An interesting question that arises is if these biases change in time. An investor can realise that his way of interpreting information is 'biased' and can try to "unbias" himself. We can devise an extension to our model where an agent can try to 'unbias' itself. By transforming equation 5.6 into the formula below,

$$Bias(I_t) = (a_P * I_t + b_P - c_p) / d_p \quad (5.7)$$

an agent tries different sets of values for the c_p , d_p parameters in order to arrive at the perfect information stream. In our results chapter we argue that doing this active 'unbiasing' method is equivalent to a normal evolution where underperforming ('biased') agents will theoretically exhaust their wealth and stop trading (thus disappearing from the market). This self-elimination will leave only performing (*not biased*) agents to continue trading. During our simulations we see exactly which agents survive since not having a bias is not necessarily a guarantee for survival in a financial market.

5.4.3 The investment decision

After receiving new information, an agent executes his investing strategy and generates new orders. Depending on the agent's motivation (liquidity, profit, etc.) and market beliefs we have created a typology of investment strategies (see Figure 5.6).

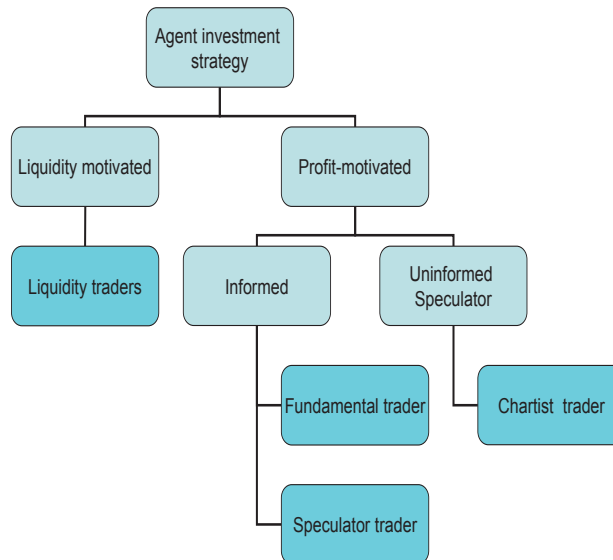


Figure 5.5: Typology of investment strategies

The first choice in distinguishing between different strategies is that of the objective of investment. We make a difference between active and passive investors. Passive investors are those that use markets as a pure means of conserving and transferring wealth through time. They do not get actively involved in making 'good' trades. Instead, they buy assets when they have extra cash and sell the assets when they want to consume. They expect to receive the market return for their investment period. We call these market participants 'liquidity investors'.

In opposition to the passive way of the 'liquidity investors', active investors engage in efforts to make superior gains on the market. These investors expect to receive more than the market returns. They believe they can outperform the market by using their skill and information. Although both types of investors expect some returns for their invested savings, only the active investors are those trying to increase their earnings above the market level. Therefore we refer to active investors as being 'profit motivated'. This typology is not intended to be exhaustive. Instead we describe as many available investor types as possible using a few concentrated and distinctive patterns (which are modelled in our simulator).

5.4.3.1 Liquidity investors

Liquidity traders are investors that have extra cash. These traders invest their cash for short periods. They do not look at market fundamentals and are motivated by a number of reasons which we generically call "liquidity" reasons. Liquidity traders

send only market buy or sell orders at random time periods.

Because price is determined from the limit order book, these traders' orders do not have any direct effect on the formation of the market price. In a continuous trading market, the liquidity traders have an effect on prices since more market buy (sell) orders will resolve more bid (ask) prices and eventually increase (decrease) the current price. In our case, of extra-day trading, liquidity traders offer liquidity to the market (to the extent of their cash holding). In real financial markets these traders are often companies with extra cash to invest for short periods (overnight, week-end, etc.). This liquidity trading strategy is modelled by the method:

Strategy of liquidity trader x

Inputs: Agent's $Cash_x$ and $Stocks_x$ parameters and the selling probability $P(Sell) = 50\%$

Actions

1. $p =$ random real from $[0,1]$
2. If $p \geq P(sell)$ and $Cash_x > 0$ than send Market Buy order for a maximum amount of $Cash_x$
3. If $p < 1 - P(sell)$ and $Stocks_x > 0$ than send Market Sell order for a number of $Stocks_x$ stocks

In this case, with a probability of 50% a liquidity trader will buy stocks with all his cash or sell all of his stocks. This state-independent probability will ensure a constant (on average) volume of market orders. These traders can eventually adapt their probability function to depend on recent returns or volatility levels - like buy less when volatility increases. This type of state-dependent behavior introduces, by design, correlations between volume and returns (and/or volatility). From the above strategy we infer that 'liquidity' traders can buy high and sell low so at a loss. It is interesting to see if and when the wealth of these traders' shifts in the favour of more informed investors or if these liquidity investors manage to become richer. By adjusting simulation parameters we can distribute liquidity traders into different strategy groups. If we lower the probability of selling, these traders can become what the literature calls 'buy and hold' investors. With $P(Sell)=0.01$ a liquidity trader will buy 99% of the time (so in practice he will buy once with all his cash) and he sell his portfolio after a certain holding period (maximum 100 trading days, on average). Lowering $P(Sell)$ even more would increase the holding period of this new trader (see Figure 5.6).

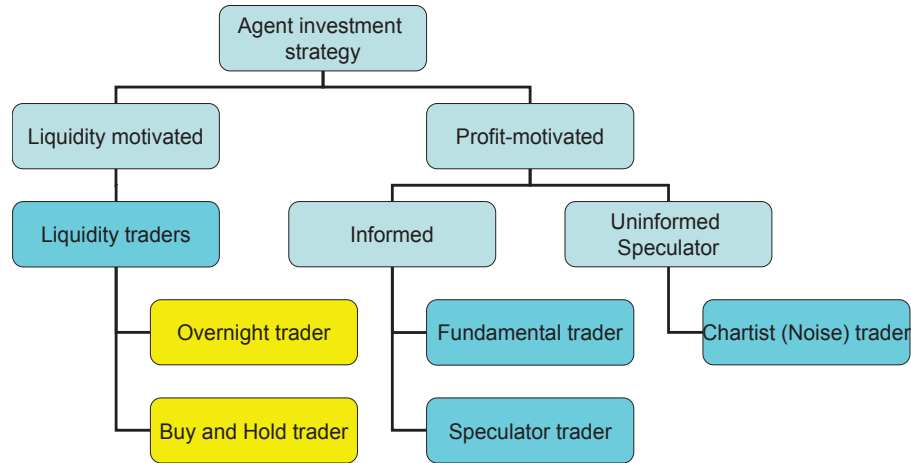


Figure 5.6: Typology of investment strategies (extended)

On the opposite side, if we increase $P(\text{Sell})$ to 99% we model the behavior of an over-night investor. This investor has a holding period of 1 days (1% of his time he buys) and only stores his cash in stocks for a very short period. As with the main liquidity investor type, these two extreme variants of liquidity traders have an impact only on trading volume and not on prices.

5.4.3.2 Profit-motivated investors

Investors that enter markets to actively increase the value of their portfolio are called profit-motivated investors. They try to predict future prices or the tendencies of the market. These investors can be divided into two groups, depending on their beliefs about the financial market. *Informed investors* are those which believe markets are not entirely efficient and they will try to profit from these inefficiencies. *Uninformed speculators* (or pure speculators) are those that believe markets are efficient, meaning that the fundamental information is integrated in prices, and they consider that trading based on the available economic information is not profitable. Therefore, these investors, called "chartist", "noise" or "technical" traders will use non-economic information to forecast and profit from future market moves.

The fundamental trader is the investor who thinks the market price should reflect the economic value of the underlying company. If this relation doesn't hold, this

investor will act in the direction of the 'correct' price. The strategy of this type of trader can be described as:

Strategy of informed fundamental investors agent x

Inputs: $Bias_x(I_t), P_t$, Minimum expected return for trading of agent x, $r_{x,min}$

Actions:

1. Forecast, at time t, the agent's expected value of the asset's market price at time t+1 $E_{x,t}(P_{t+1}) = P_t + Bias_x(I_t)$
2. Compute the confidence interval of this price

$$IC_{P_{t+1}} = [E_{x,t+1}(P_{t+1}) * (1 \pm r_{x,min})]$$

3. If $P_t < min(IC_{P_{t+1}})$ then send buy Limit Order at buy price $min(IC_{P_{t+1}})$
4. If $P_t > max(IC_{P_{t+1}})$ then send sell Limit Order at sell price $max(IC_{P_{t+1}})$

In the first step, the trader uses his available information (biased as it is) to forecast the future market price. Looking at equation 15 on page 81, that describes the formation of the market information, we see that an informed trader without any bias could perfectly forecast the next fundamental value of the asset. Even if a fundamental trader can forecast this value (biased or not) he knows that the market will not necessarily move the price towards this fundamental point immediately. Therefore the trader will take a precaution measure by acknowledging that the next market price will be close the correct fundamental value. This is way the trader computes, in step 2, a confidence interval using his minimal return requirement. This type of trader believes the next future market price will be situated somewhere inside this interval. If the market price is outside of this interval it means that, if all goes normally, there is an opportunity for a minimally accepted profit since the market will make an important adjustment to prices in the near future. This parameter, of minimally accepted return, can be considered as a function of the agent's risk aversion and the perceived market risks. The future price confidence interval is computed using a subjective parameter, minimal return, that contains the investors risk aversion levels as well as his interpretation of the market and asset risks. In a very competitive and precise market, that can correctly and quickly integrate economic information into prices, the market prices will always be very close to the fundamental realities. Therefore the risk of market misestimating is small and likewise the minimal required returns of traders are smaller. A trader with high aversion for risk will want to enter a trade only if there is a significant profit possibility. Therefore higher risk aversion implies a higher minimal return expectation. Depending on the trader characteristics, his risk aversion coefficient

can fluctuate depending on his perception of the market (more risky, less risky, etc). In our conception of the market, this parameter should rest constant for all investors. Indirectly, the daily average risk aversion and market risk are not constant because not all investors' trade daily. We can understand that the smaller the minimal required return of investors the better (and faster) prices readapt to fundamental values. For example, an average of 5% minimal required returns, will create a market where prices will fluctuate randomly in a +/- 5% band around the fundamental value. When the price is too close (>5%) to the fundamental value than the informed investors are not motivated to approach the market price even closer. If we would make the minimal required return a function of a constant risk aversion and a proxy of market risk (like past volatility) we would be repeating known design mistakes by introducing explicit autocorrelation in return volatility (see chapter for more details).

We can rewrite the equations, from steps 3 and 4 of the informed fundamental investor strategy, in a return form. This will show that the an informed trader takes a decision based on his expected returns:

1. ...
2. ...
3. If $E_t(r_{t+1}) < 0$ and $|E_t(r_{t+1})| > r_{x,min}$ then make limit sell order
4. If $E_t(r_{t+1}) > 0$ and $|E_t(r_{t+1})| > r_{x,min}$ then make limit buy order

In other words, such a trader considers that if the market price is too far from the fundamental value than it will intervene in order to profit off of this gap. When betting that the market will adjust towards the fundamental value, the informed trader believes that markets will be once again efficient and they will soon correctly price the asset. For this reason, these traders are the ones that confront the crowd-formed trends and introduce stabilizing effect in the market. This behavior is opposite to that of an informed speculator who chooses to go with the crowd even though he has information that shows markets are mispricing the asset.

In opposition to *informed fundamental investors* we introduce the typology of *informed speculators*. These speculators have access to the same level of, biased, information as informed investors. They have the same resources both in cash and assets and they also share the same capacity for forecasting a confidence interval for the future fundamental price. The only difference between these two informed trader types lies in their conception of the market. *Informed fundamental investors* believe that markets are efficient but can sometimes make estimation mistakes. When these mistakes occur, informed traders step in and act in the direction of correcting the error: when the market overestimates an asset informed investors sell and when the market underestimates an asset they buy. On the other hand, *informed speculators* believe that markets are not that efficient and assets are mispriced for long uninterrupted periods. Therefore, informed speculators will not always go against the

market trend, even if their information tells them that the market is fundamentally wrong. Theoretically, it is possible that these speculative investors can create a positive feedback loop to reinforce and sustain a trend in market (provided they act at the same time). Different from other speculator investors (like technical traders) the informed speculators trade at incorrect prices only to a certain point. After a mispricing limit, informed speculators decide that the market is too far away from the fundamental value. In 5.7 we show the evolution of a fundamental and market time series in a hypothetical market in order to explain how informed speculators get involved in the markets.




Figure 5.7: Speculators and the evolution of fundamental value (green) and market price (red)

The green line represents the fundamental value and the red line the market value of an asset. According to their strategies, informed fundamental investors will enter the market and will start selling the asset (in order to lower its value towards the fundamental value). These investors will take action when the market price P_t will be around the r_{min} threshold (an average minimal expected return). The selling orders of these investors will press on the market in an effort to lower the price, which is above the fundamental value. In the same time, the informed speculators can also start acting in the market but in the direction proposed by the market: away from the fundamental value. These speculative actions, of buying even if the asset is overpriced, can increase the actual trend. Such actions are executed only if the speculators believe they can profit off this operation. But informed speculators can profit off such a trend if they buy with the market and sell against it after

some time. Indeed, these speculators will go with the market for some time and after the prices are considered too far from the fundamentals, like passing over the R_{max} threshold, they will try to close their positions. At this point the speculators will have a similar behavior as the informed investors. Speculator investors can be viewed as a hybrid between informed investors (with a high minimal expected return) and uninformed speculators (chartists that amplify market trends when the minimal expected return condition is not fulfilled).

The second major type of profit-motivated investors is represented by the **uninformed investors**. These investors believe economic information is not a good indicator for forecasting future market moves. They justify this belief in two ways: 1. economic information is already integrated into prices and the price movements between moments when economic information is released are generated by market "psychology" or 2. markets rarely use economic information and instead they conceive strategies in order to make fast profits. Because they hold such beliefs, uninformed investors do not try to obtain and use economic information. Instead they construct mathematical, statistical or other forms of indicators that help them discover profitable market 'opportunities'.

Inside our simulator, we modelled these investors, also called "chartists", using simple memory functions that determine trends inside the market price series. Each uninformed investor has two strategy maps with a length of S . Each strategy has S units bellowing to one of these states $[-1, 0, 1]$. Each investor has two strategies: one strategy is used to determine a buying signal and another to determine a selling signal. The strategy configuration is determined randomly for each investor at the beginning of a simulation. Because we are using a model without active learning, see chapter , the investors will be using the same strategies during the entire simulation. A strategy of length 3 can look like:

$$Strategy_1 = [1, 1, -1]$$

$$Strategy_2 = [-1, -1, 0]$$

The numbers $[-1, 0, 1]$ represent properties [**negative, indifferent, positive**] that returns must have to active a strategy. Each strategy set is confronted with the history of returns and if all the properties match than the strategy is activated.

Therefore at time t , $Strategy_1$ is activated if $r_{t-2} > 0$ and $r_{t-1} > 0$ and $r_t > 0$. $Strategy_2$ is activated if $r_{t-2} < 0$ and $r_{t-1} < 0$.

Depending on the strategy's purpose, the chartist will either sell or buy stocks when his strategies are activated. His investment action will be executed using limit orders. Because the chartists don't know the future value of the market prices, but only the direction of future movements, they will try to sell/buy at "**good**" prices. So the chartists will send limit buy/sell orders at the last market price plus/minus a random variable.

Buying orders are sent with a price of $Ask = P_t - \tilde{x}$ and the selling price

$Bid = P_t + \tilde{x}$ where $\tilde{x} \in N(\mu_c, \sigma_c)$. This method ensures that chartists always try to get out as much as possible from their trades (buying low and selling high). Moreover, it also allows orders that are not always higher than the market price. With this simple trading procedure chartists can detect and profit off of market patterns. Moreover, if enough chartists have similar strategies, their actions will create self-fulfilling prophecies.

Strategy of uninformed chartist investor agent x

Inputs: Buying Strategy S_b , Selling Strategy S_s , μ_c, σ_c

Actions:

1. If buying strategy S_b is activated then send limit buy order at price

$$P = P_t - N(\mu_c, \sigma_c)$$

2. If selling strategy S_s is activated then send limit sell order at price

$$P = P_t + N(\mu_c, \sigma_c)$$

5.5 Review of configuration parameters

We described the different components of our LUMA market simulator: the economic environment, the flow of information, the formation of prices and the strategies of the simulated investors. Our research tool is fairly complex because it is intended to model a highly complex real object: a financial market. Since our goal is to have a parsimonious description of a market we made efforts to limit the number of parameters to be used in customizing our simulator. These parameters describe individual parts of the simulator and can be divided into four categories:

- **Asset parameters**

- fundamental value daily average return and volatility: μ_V, σ_V

- **Agent parameters**

- proportion of agent types: % uninformed, % technical, % informed, % special types
- uniform distribution of bias parameters $a \in [a_{min}, a_{max}]$ and $b \in [b_{min}, b_{max}]$
- uniform distribution of minimal expected returns for informed traders $r_{min} \in [R_{min}, R_{max}]$
- strategy length for chartist traders, S
- mean and dispersion of price quotes for chartists, μ_C and σ_C

- **Market parameters**

- Total number of agents, N
- Starting wealth and portfolio of agents

This generic parameter set describes the inputs to our LUMA market simulator. In the next chapter we use the LUMA simulator to answer our research questions. The general approach taken will be to find different settings configurations that will generate desired aggregate behaviours. After finding the 'seed' settings, or range of settings, we explain what the emergent proprieties are and how different parts of the markets interact and generate particular features in the resulting market series.

Answering the research questions

Contents

6.1	Proof of concept validations - Benchmark simulations . . .	105
6.2	Biases and market prices	113
6.2.1	Biases involving lags in information interpretation	124
6.3	Volatility clustering and other "stylized" facts	129
6.4	Bubbles: how prices can grow above the asset's fundamen- tal value	138
6.5	Survival and evolution of agents	147

This chapter is dedicated to using the LUMA simulator and answering the research questions. For each research question we investigate different possible causes of the researched effects. Through simulation we show the causal links between investor behavior biases and aggregated market proprieties.

6.1 Proof of concept validations - Benchmark simulations

To test that our financial market simulator functions properly, we make a simple tests that have the following properties:

1. the modelled economic situations are easily described
2. the link between the investor input parameters and the simulated financial market outputs are intuitive

We run such tests to verify that the simulator has not got any endogenous faulty or unpredictable behavior. Also, we test the correctness of each of the important components of the simulator (price formation, information bias, agent orders). These tests will help the reader become more familiar with the settings and the usage of the simulator. Later on it will be simpler to perform and verify more complex tests.

In our first test we see the evolution of prices and of the fundamental information in the absence of any trading agents. Like in all the following tests, our fundamental

information will have a normal distribution with a daily return governed by the parameters μ_V and σ_V . Running a simulation without agents implies that there are not trades and the market price is constant and equal to the first pre-set price P_0 (which is the IPO price). For every simulation test we will provide a table with the main parameters of the simulation. When some parameters are not needed (like agent-specific variables in this case) they will not be listed in this settings table. By using these simulation descriptions we are able to easily understand the link between the inputs and outputs of the simulation tool. In the Table 6.1 we list the simulation's parameter specification for this first benchmark test.

We run a simulation using the above parameters and we fix the starting price (and fundamental value) at 40. As expected, no investor implies no trades and therefore no changes in the market price after the initial public offering price.

The fundamental value fluctuates randomly and follows a normal distribution. Since there is no point in using different distributions this fundamental economic value will always have a normal distribution. We run a second benchmark test using the same parameters as before with the exception of the agents. Noise traders (that send market buy/sell orders at random times) are introduced and the results are similar as before. As confirmed by Veryzhenko, Mathieu, Brandouy (2010), noise traders (or Zero Intelligence Traders) are not sufficient to generate realistic market behavior. Because no trader sends price quotes to the market (meaning limit orders) the formation of a new market price is not possible. To observe new market prices we need to have a market with at least a few investors that can send limit orders to the market.

We make a benchmark test with perfectly informed profit-motivated investors. These investors don't have any biases and they correctly compute the future fundamental value. Moreover, these investors engage in actions to move the market price towards the fundamental value. Because these agents make perfect estimations they are willing to trade if the minimum expected return is positive.

We run a simulation with the above parameters and observe that trading takes place. The market price follows perfectly the fundamental values. Since investors know exactly the new fundamental values, and they expect at least 0% returns, they will send orders to buy or sell the asset at exactly the fundamental value.

Parameter	Value	Explanation
μ_V	0%	No long-term growth in fundamental value
σ_V	1%	1% standard deviation of the fundamental value return
N	1000	Number of trading days
Nr agents	0	No agents

Table 6.1: Parameters for Test 1, first benchmark test (no agents)




Figure 6.1: Representative result - Market price (red) and fundamental value (black) for Test 1 - benchmark without investors

Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	100	Number of trading days
Nr agents	200	100% Informed investors
Minimum expected return r_{min}	0%	0 for all investors
Bias a_p	1	Investors have perfect information
Bias b_p	0	

Table 6.2: Simulation parameters for Test 2 with perfectly informed investors that do not speculate

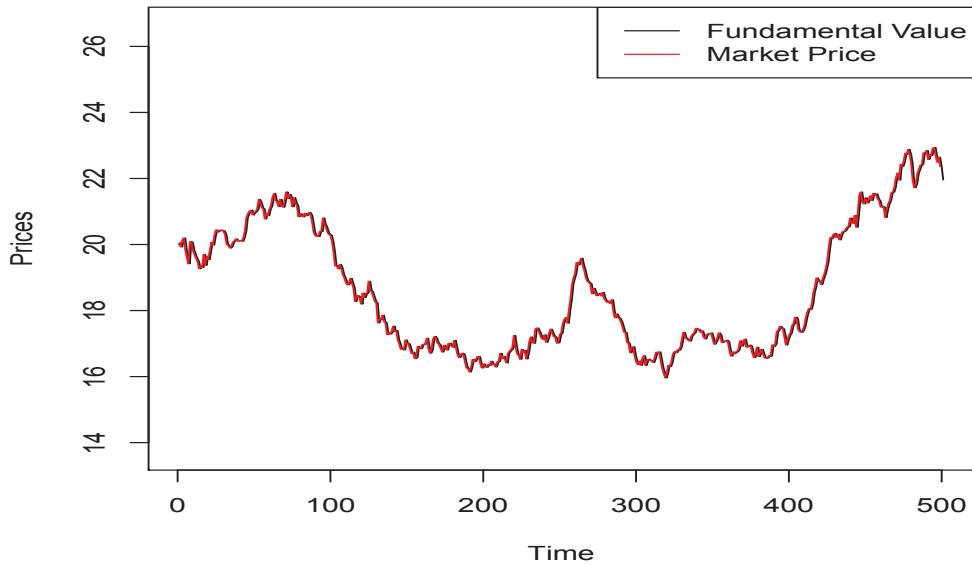


Figure 6.2: Representative result for Test 2: Fundamental (black) and market price (red) in a market with perfectly informed investors

Trading is plentifully since all investors are happy to buy or sell at the correct fundamental value. If we introduce risk aversion factors (represented by more than 0% minimal expected return) then the investors will try to sell/buy for a profit. This is similar to introducing a transaction cost C payable by both parts in a transaction. All investors would be willing to buy the asset at a discounted price, $P_t - C$, or sell it at $P_t + C$ (where P_t is the correct fundamental value of the asset). In this situation all trading would stop since no investor would be willing to buy/sell at a loss.

Proposition 1. We say that in an efficient market, populated only by rational and perfectly informed investors, trading can occur only in the absence of transaction costs.

If no transaction costs (or risk aversion or perceived market risks) exist, the market price will be kept equal to the fundamental value. Therefore all the 'stylized facts', discussed in chapter , can be attributed to the distribution of the fundamental value (since the market does not make any modification on this value). If this is not true or not verifiable then we should look at other factors or ingredients to a market that can distort market prices and create 'stylized' facts. It is clear now that trading away from the fundamental value or in the presence of transaction costs can only take place in markets populated by non-rational or not perfectly informed investors. We will go into more details about such market composition in the following tests.

The previous test assumed that informed investors were willing to trade if they expected at least a zero return, $r_{min} > 0$. But what would happen if investors

would have higher expectations of returns? A financial reason would be that a stock should be strictly more than a risk-free asset (which pays 0%). We modify the last test and introduce a level of heterogeneity regarding the minimum expected returns of informed investors.

An investor i , with a strictly positive minimum expected return (also implying risk aversion), will always try to buy at a price of $F_t * (1 - r_{min})$ and sell at a minimum price of $F_t * (1 + r_{min})$. Because all investors are perfectly informed (they are not biased) nobody will be willing to be a counterparty (since it would imply selling or buying at a loss) and no trading will happen. To avoid this situation we decide that, in the majority of tests, informed investors will all have a minimum expected return of 0%. This assumption implies trading will be possible between informed investors (even with perfect information) and most importantly an efficient market will be one where the market price will move fast towards the fundamental value of the asset (since investors with information do not ask for returns on risk). Because no trading can be possible when perfectly informed investors are risk averse (demand an positive r_{min}) we extend the first proposition and say that:

Proposition 2. In an efficient market, populated only by rational and perfectly informed investors, trading can occur only and in the absence of transaction costs and only between investors that are not risk-averse (demand a 0% r_{min}).

This observation is important because it describes an idealistic but unattainable situation. In real financial markets trading is plentiful in spite of the assumptions of proposition 2. In real markets we always have costs for transactions and investors are willing to trade only if they are well paid for the risk they take (they are risk-averse, $r_{min} > 0$).

Therefore, the trading volume we observe in real financial markets, where informed investors are risk-averse, is possible only if one of these conditions is true:

1. Investors are not always rational
2. Investors are not all perfectly informed
3. Other types of investment strategies exist (aside that of an informed investor that tries to profit out of the difference between the market price and the asset's fundamental value).

If we rerun Test 2 with investors with a variety of biases (a_p and b_p are variable) trading will plentiful. Investors have different versions of the "correct" fundamental value and are willing to trade - in a transaction both parties believe ex-post that they are making a profitable trade. We also test assumption c) and introduce alternative types of investors.

First we simulate a market with perfectly informed risk-averse investors and noise traders. We remind that noise traders only send market orders so they will

Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	100	Number of trading days
Nr agents	200	100% Informed investors
Minimum expected return r_{min}	(0%,10%]	Informed investors have minimum returns expectations more than 0% and a maximum of 10% (drawn from an uniform distribution at the start of each simulation)
Bias a_p	1	Investors have perfect information
Bias b_p	0	

Table 6.3: Simulation parameters for a test with perfectly informed investors that do not speculate

be counterparty to the best limit orders. As reported by (Black, 1986) such noise traders are essential for the liquidity of markets.

If informed investors would have been accepted 0% minimum expected returns then the market price would equal the fundamental value. In this situation trading occurs and prices fluctuate above the asset's fundamental value (because investors expect a return on their investment and therefore sell higher or buy lower than the fundamental value). Because informed investors always want a positive expected return the wealth of noise traders will be transferred to the informed investors. The speed of transfer depends on the fundamental value's fluctuations and on the distribution of the informed investor's expected returns.

We can observe that prices fluctuate along the fundamental value. Informed investors always try to sell above or below the fundamental value depending on their information and return expectations. Trading occurs because noise traders act as counterparties. Eventually, wealth is transferred from noise to informed traders. When all wealth is transferred to well-informed traders the trading will stop (due to the absence of counterparties).

In the last benchmark test we simulate a market only with technical traders. These traders look for certain specific patterns of past returns (like up, up, down or down, up, down) and take specific actions. They aim to predict and take advantage of expected future price movements. In this kind of market the economic information is useless because no investor uses it.

Because of their beliefs, chartists do not use the fundamental information. They look for "patterns" in return series and believe these patterns repeat themselves. As we can see in Figure 6.5 the chartists are coordinating on trend patterns. Once the price drops two or three times in a row, some traders will have their selling rule activated and will start selling. These selling actions will move the price down

Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	100	Number of trading days
Nr agents	500	70% Informed investors 30% Noise traders
Minimum expected return r_{min}	0.01%, 20%]	Informed investors have minimum returns expectations more than 0.01% and a maximum of 20% (they are all risk-averse)
Bias a_p	1	Informed investors have perfect information
Bias b_p	0	

Table 6.4: Parameters for a test with risk-averse perfectly informed investors and noise traders

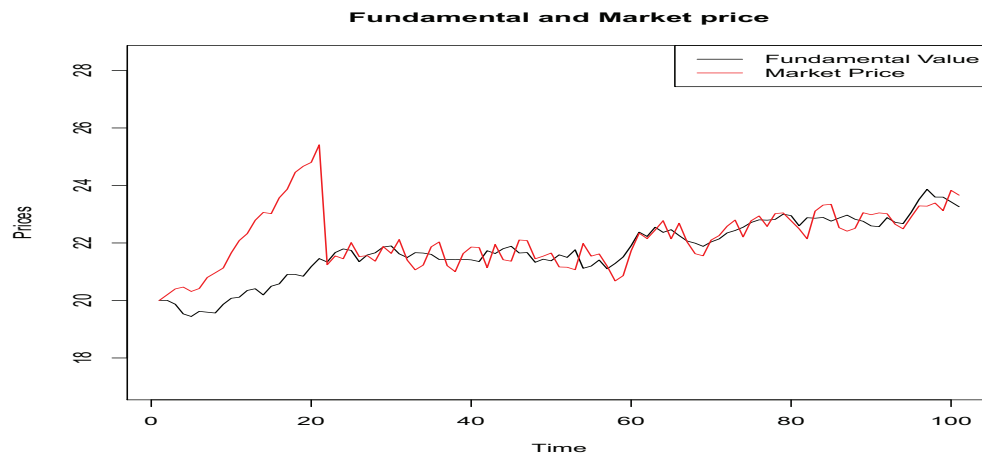


Figure 6.3: Price time series (red) for a market with informed ($r_{min} > 0$) and noise traders (parameters in Table 6.4)

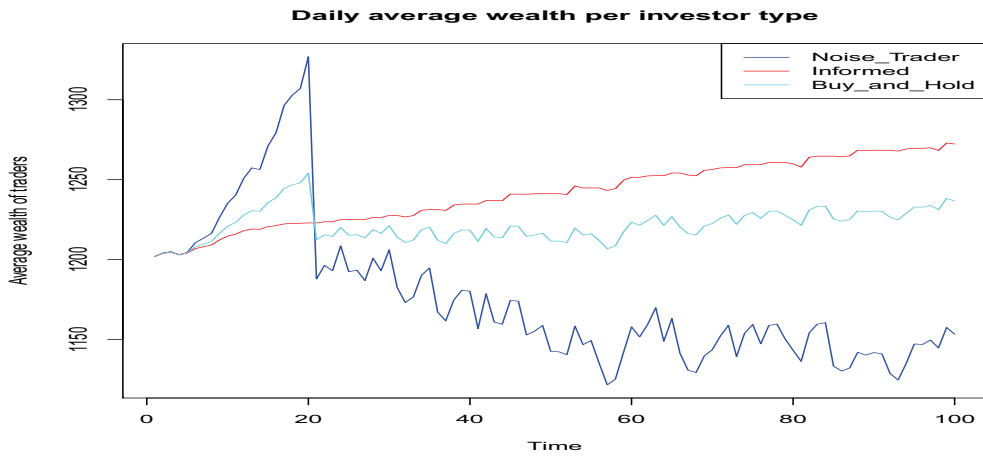


Figure 6.4: Wealth evolution for a market with informed ($r_{min} > 0$) and noise traders (parameters in Table 6.4)

Parameter	Value	Explanation
N	300	Number of trading days
Nr agents	300	100% Technical traders
Memory length	3	Traders use rules regarding the last 3 returns (r_t , r_{t-1} and r_{t-2}).
μ_C	5%	Traders try to sell/buy with a random mark-up of $R \in N(\mu = 5\%, \sigma = 8\%)$
σ_C	8%	

Table 6.5: Simulation parameters for a market having only technical traders

even more. The drop in price will continue until the chartists (which reinforce the price fall) have sold all their assets. At that moment other strategies will begin to dominate. In Figure 6.5 we see that the second price trend is generated by traders who believe prices should rise more and more.

The amplitude and duration of these price trends depend on the traders' wealth and their desire for profits (μ_C and σ_C) and are independent of actual asset's fundamental value. We show below another example where traders are much more "undecided" in their expected profits: their mean desired return is low $\mu_C = 1$ yet the volatility of their expectations is much higher than before, $\sigma_C = 10\%$. This high volatility implies that the bids and asks of the trades will be more dispersed and will generate high market returns (see Figure 6.6).

Observing the price evolution from Figure 6.6, we see price bubbles that are created and destroyed much faster than in the previous test. We are able to observe bubbles that loose intensity after the initial huge rise in price (up to 400 or 600 starting from the IPO level of 20). This market price behavior is generated by traders which wager on rising prices and buy at very high price levels. After the bubble bursts (the rise being limited only by the chartists' available cash), these traders will eventually sell their assets at low price level thus absorbing big losses. These losses hinder their ability to make wild speculations in future situations.

We proposed a few simple tests to test the behavior of traders. We discovered that noise traders cannot trade alone in a market since they don't send any price quotes. Perfectly informed rational traders are able to create a market if they agree to trade, either buy or sell, at the 'correct' price (they are risk neutral). Yet, such an efficient market cannot exist in the presence of transaction costs. Introducing transaction costs (or positive risk-aversion) stops all trading because no counter-parties exist. The only ways in which trading can exist in a market with informed risk-averse investors are: investors have heterogeneous informational biases or other non-fundamental based investment strategies exist (e.g. noise or chartist traders).

In the end we observed the effects of chartist traders inside a market. These traders can move prices by themselves and have the tendency to create price bubbles. All of these behaviors are rather simple and intuitive to predict. What happens when we introduce features that approach our agents to their real counterparts? What happens when we simulate financial markets with a mix of different investor types? We answer these questions in the next chapter.

6.2 Biases and market prices

In this chapter we study the effects of different investor information biases on the formation of market prices. We start with a simple market setting (like the one described in Table 6.4) where we mix perfectly informed investors with biased investors.

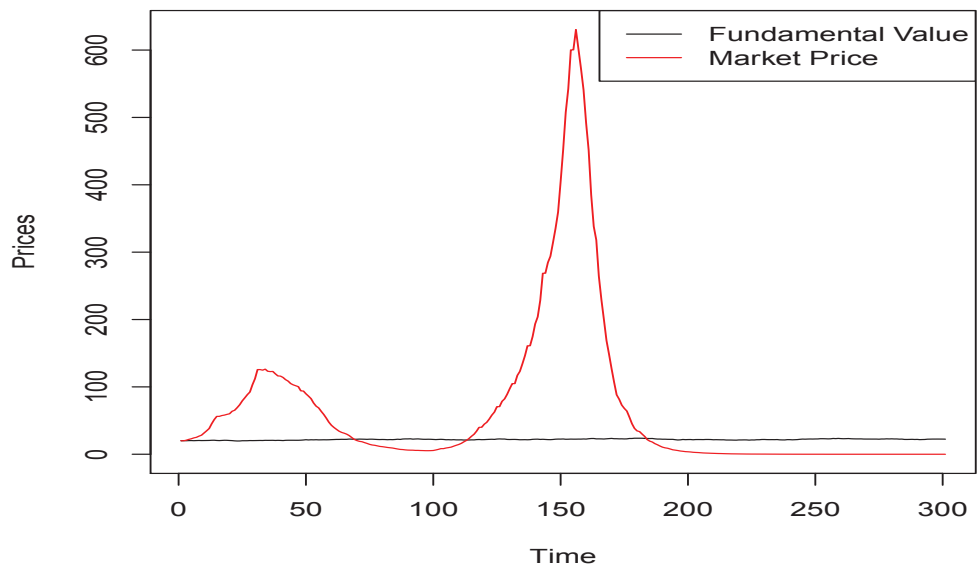


Figure 6.5: Price time series (red) for a market only with chartist traders (parameters in Table 6.5)




Figure 6.6: Price time series (red) for a market only with chartist traders with highly volatile return expectations. Market price in red and fundamental value in black

Test 4		
Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	2000	Number of trading days
Nr agents	500	50% Perfectly Informed investors ("GOD" investors) 50% Biased Informed investors
Bias a_p	[0.7, 1.3]	Biased investors will interpret information with a maximum of +/- 30% error.
Bias b_p	0	
Wealth		500 in cash and 10 stocks for each investor
Minimal expected return	0%	All informed investors are not risk-averse (they trade for at least the "perceived" correct price)

Table 6.6: Scenario with informed investors, with balanced biases and without biases

In this scenario half the investors in the market know perfectly the future fundamental value and trade accordingly (remember that they have 0% minimal expected return so they are willing to buy or sell at the correct market price). The biased investors will send orders with prices inside an interval of -30% to 30% around the correct fundamental value. We observe, in Figure 6.7, that the simulated market price is very close to the fundamental value (the mass actions of the "GOD" investors keep prices close to fundamentals). With a few expectations, most price deviations are very small, at a level of 0.02 points out of the mean market price of 20. An interesting perspective is offered by the evolution of the average wealth per investor type. The biased informed investors, represented in red in Figure 6.8, have on average the smallest wealth levels. The biased investors have returns that are surpassed (most of the time) by a benchmark strategy of Buy & Hold (a virtual investor that would not have made a single trade from the beginning of the simulation). We observe that the perfectly informed investors (green line in Figure 6.8 and marked 'GOD' in the legend) have a more stable average wealth and they maintain a better performance than the biased investors and even better than the buy & hold strategy.

Because informed investors are uniformly biased (the information a_p bias parameter has a uniform distribution) they will have complementary effects on the market price: some investors will overestimate while others will underestimate the price. This effect is observed by the small differences between the fundamental value and the market price.

In the next scenario, Test 5, we use the same parameters as for the last test (see Table 6.6) and we modify the information bias parameter. This time, informed investors have a bias parameter a_p drawn randomly from the interval [0.9, 1.4]. This bias will impose on informed investors a tendency to overestimate the fundamental

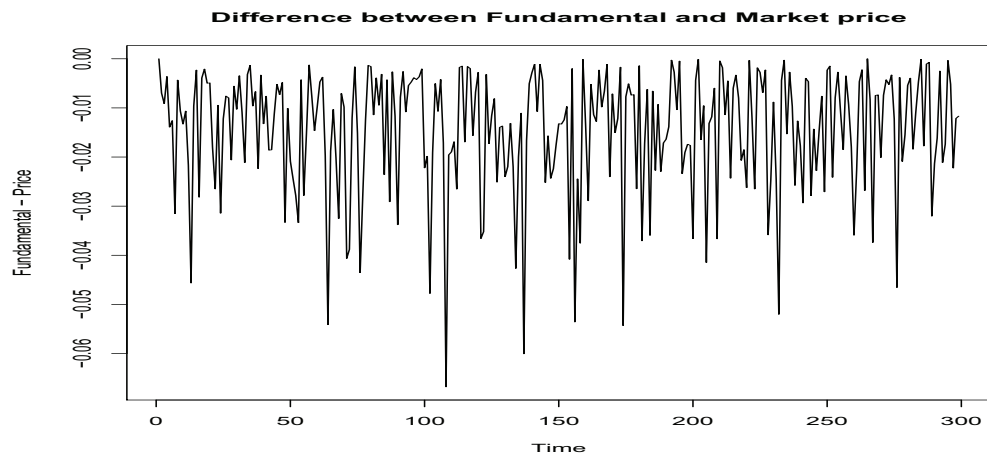


Figure 6.7: Graph of the difference between fundamental value and market price (market parameters from Table 6.6)

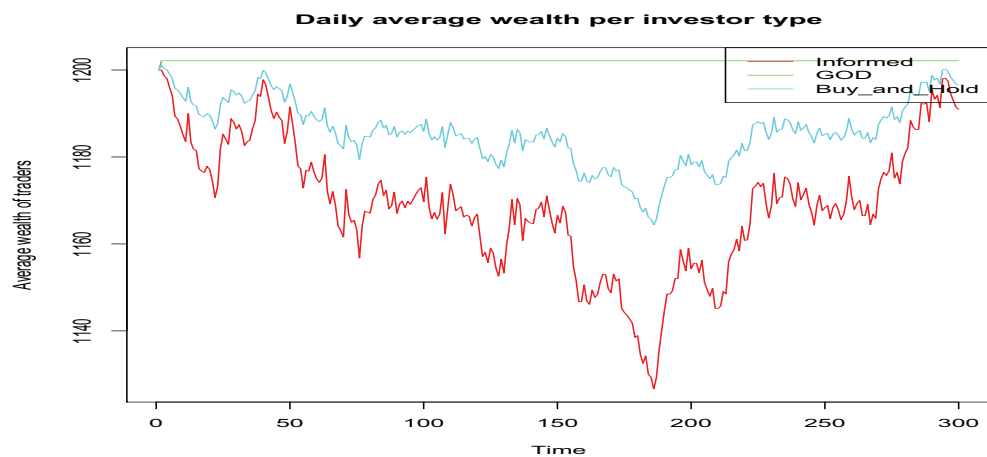


Figure 6.8: Graph of the average portfolio Sharp value for each investor type (market parameters from Table 6.6)

value. Such an overestimation is visible in Figure 6.9 where we see that the market price is often above the fundamental value (albeit with small differences but not negligible). When the market price goes above the fundamental value the 'GOD' investors will sell their assets and the biased informed investors will buy them at prices about the fundamental value.

Because biased investors consistently misprice the asset they will most often buy high and sell low. These investors will trade with each other (since some are less biased than others) or with the perfectly informed investors (GOD). We see in Figure 6.10 that the biased investors perform worse than the buy and hold strategy.

To better understand how the fundamental value fluctuates away from the market value we will explain how investor biases distort resulting prices. In Figure 6.11 we can see the graph of a theoretical fundamental value in three points: V_t , V_{t+1} and V_{t+2} . At the first time moment the market price is exactly the fundamental value. The new information flow, $I_t = V_{t+1} - P_t$, will inform investors about the evolution of fundamental value in the current period. If the majority of the investors are biased then they will interpret this information as $a_p * I_t + b_p$. Therefore the new market price P_{t+1} will be a biased function of V_{t+1} . Going on in such a way, the market price will fluctuate around the fundamental value. These differences between market price and fundamental value will not be cancelled out through averaging since the distance between price and fundamental value is constrained by the availability of funds of the biased investors, price levels and existence of sufficient counterparties. We underline that as long as there are sufficient informed investors (with $a_p > 0$) than the market price will have a tendency to revert to the asset's fundamental value. As a consequence of this investor bias ($a_p > 0$) we also observe the existence of excess volatility.

In Figure 6.12 we notice that the volatility of market returns is most often above that of the fundamental value. As we will see this excess volatility persists in situations where other stylized facts are present.

We observe a consistent average overestimation of prices (in each test we take the

Summary statistics for 100 simulations of Test 4	Average of excess volatility	Average fundamental to market price distance
Mean	0,00038403	-0,03260471
Standard deviation	8,62346E-05	0,00479993
T value	44,5331688***	-67,9274698***

Table 6.7: Summary statistics for 100 trial simulation with input parameters from Test 4 with biased a_p parameters




Figure 6.9: Graph of the difference between fundamental value and market price with optimist biased investors)

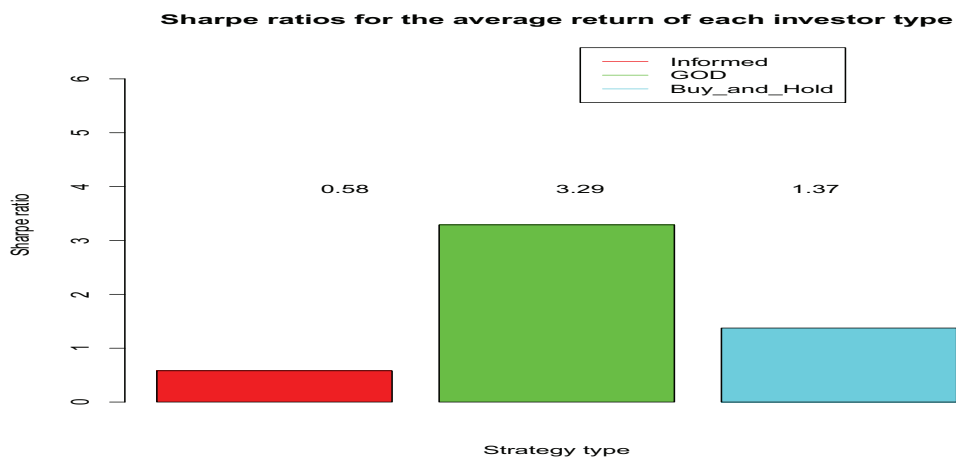


Figure 6.10: Graph of the average portfolio value for each investor type with optimist biased investors




Figure 6.11: Dynamics of market price and fundamental value with biased investors




Figure 6.12: Excess volatility in a market where informed investors are optimists, a_p in $[0.9, 1.4]$

average excess volatility of the average distance between fundamental and market prices and we average these measures across all the simulations).

After establishing that a population of biased ($a_p > 0$) informed investors can produce persistently biased prices we make tests to see the influence of the factor b_p on the creation of market prices. We remind the biasing equation, for an investor p ,

$$Bias_p(I_t) = a_p * I_t + b_p \quad (6.1)$$

from chapter 5, in which a positive/negative b_p value represents an optimistic/pessimistic look on assets.

Running a simulation with the parameters from Table 6.8, yields a market price above the fundamental value with a distance equal to an average of investor randomly drawn b_p biases.

In Figure 6.13 we see that the market price jumps at the beginning of the simulation to the biased levels and this bias persists throughout the simulation (see statistics in Table 6.9).

The bias b_p is not additive and market prices do not grow consistently. This bias just keeps prices at a distance from the fundamental value. It is important to point out that the market price bias is the result of the joint actions of a number of investors having heterogeneous levels of optimism b_p . If we rerun Test 6 with a negative bias (mostly pessimist informed investors) b_p in $[-2, 0]$ we will not observe the intuitive result of market prices being consistently below the fundamental value. Instead market prices will rest very close to the fundamental series. This effect is due to the action of the GOD perfectly informed investors. When the pessimists try to lower prices with low quotes, the GOD investors will continue to provide high bid prices. If the market price drops, even a little under the fundamental value, GOD investors will seize the opportunity and start buying the undervalued asset. Because of the price drop GOD investors will be able to buy more assets with the same amount of cash. These factors will result in a relative well valued asset. To obtain a market that consistently undervalues assets we have to increase the proportion of pessimist investors to perfectly informed investors. We run a new test with the parameters presented in Test 6. To observe an effect of persistent underestimation of the asset we have to have a market where the pessimist investors have a dominant financial position (here they start with 90% of total wealth) and with an important array of pessimism degrees (from -4 to 0).

In Figure 6.14 we observe that the market price (red) falls and rests below the asset's fundamental value.

Even if the pessimism level of the investor is wider than in Test 6 (the lower limit is now -4) we see, in the statistics summary Table 6.11 that the underestimation level is lower than when investors were optimists.

Test 6		
Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	2000	Number of trading days
Nr agents	500	50% Perfectly Informed investors ("GOD" investors) 50% Biased Informed investors
Bias a_p	1	Biased investors will be optimistic about the fundamental value.
Bias b_p	[0,2]	
Wealth		500 in cash and 10 stocks for each investor

Table 6.8: Scenario with informed investors, with and without biases




Figure 6.13: Example time series variation with optimistic biased investors (market price in red and fundamental value in black)

Summary statistics for 100 simulations of Test 6	Average fundamental to market price distance (persisting overpricing bias)
Mean	-0,81382763
Standard deviation	0,09008741
Student t value	-90,3375544

Table 6.9: Summary statistics for test 6 - shows that market price is consistently overestimating the fundamental value prices (in each test we take the average excess volatility of the average distance between fundamental and market prices and we average these measures across all the simulations)

Test 7		
Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	2000	Number of trading days
Nr agents	500	10% Perfectly Informed investors ("GOD" investors)90% Biased Informed investors
Bias a_p	1	Biased investors will be pessimistic about the fundamental value.
Bias b_p	$[-4, 0]$	
Wealth		500 in cash and 10 stocks for each investor

Table 6.10: Test 7 - Market with dominant pessimistic informed investors

Summary statistics for 100 simulations of Test 6	Average fundamental to market price distance (persisting underpricing bias) b_p in $[-4, 0]$	Average fundamental to market price distance (persisting overpricing bias) b_p in $[0, 2]$
Mean	0,74505984	-0,81382763
Standard deviation	0,1136774	0,09008741
Student t value	65,5415975	-90,3375544

Table 6.11: Summary statistics for test 7 - shows that market price is consistently underestimating the fundamental value




Figure 6.14: Test 7 - Example price series in a market where pessimistic investors represent 90% of the market

This smaller gap between the market and fundamental value is due to the actions of the perfectly informed investors. When the market price drops below the fundamental value the GOD investors will buy as many stocks as their cash permits them. Afterwards, if this underestimation persists the GOD investors will just hold their portfolio of stocks and will wait for the market to correctly evaluate the stocks. Holding only undervalued stocks (no cash) will render the portfolio value of perfectly informed investors more volatile, as we can see in Figure 6.15. Of course, informed strategy that take into account time horizons would change this results (informed investors could close their positions even if the asset's price did not revert to its fundamental value).

In the absence of new investors or internal changes in investor behavior, the market will persistently under-price the asset (the perfectly informed will hold on to stocks and wait for them to be well valued). In the previous case, of asset overestimation, the GOD investors would have sold all of their assets (because they were overpriced) and kept only cash.

***Research Question #1:** What is the mix of biased investor behaviors that produces consistently biased market prices?*

Answer #1:

Informed biased investors, that over or underestimate the magnitude of the information regarding the distance between fundamental and market prices, can create

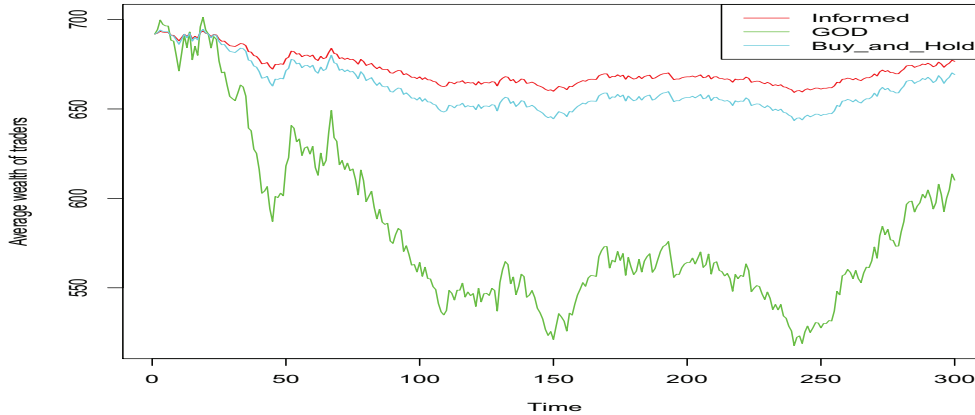


Figure 6.15: Test 7 - Evolution of wealth for different types of investors in a market where pessimist investors dominate (stocks are undervalued)

consistently biased prices. The less evenly distributed their biases are (the majority of investors tend to over or underestimate information) the more important is the average deviation from fundamental values. The more relative wealth these biased investors hold the easier the market will misprice assets. A mispriced state can persist as long as new and better informed investors don't enter the market in a sufficient enough volume. Investors that have "correct" price estimations will either sell all their assets (if overpriced) or buy with all their cash (if underpriced). When all the "smart" money is used and prices are still 'biased', these non-biased investors loose control over prices. Short-selling abilities are irrelevant since the biased investors can also take credit to buy more mispriced assets.

6.2.1 Biases involving lags in information interpretation

In this chapter we study the effects of a very specific type of bias: lag in information. We remind the equation through which information is constructed and biased for an individual agent:

$$I_t = V_{t+1-lag} - P_t \quad (6.2)$$

From the original equation we replaced the index $t + 1$ of the fundamental value with the index $t + 1 - lag$. In all the previous tests the lag variable of informed investors was 0. This implied that all investors considered at trading period $[t, t + 1]$

the information that explained the new fundamental value of that period. In reality this is not the case because investors sometimes look at outdated information and infer from them their conclusions about the present time. The existence of a lag in using economic/financial information can also be explained by the existence of some psychological biases like 'hindsight' or 'availability' bias (see chapter for complete descriptions).

We study the effects of this parameter, *lag*, on the formation of prices using a simple market structure, described in the table below.

In Table 6.12 we have described a market with no growth and where all the investors have unbiased information which can be lagged for a maximum of 8 days. If at time t an investor receives information with a lag of 1 he will be estimating the next market price, P_{t+1} , based on the information about the fundamental value of yesterday, V_t , instead of the current period's fundamental value V_{t+1} . We run a simulation with these parameters and we observe the appearance of stylized facts (leptokurtosis and volatility clustering).

We can observe in Figures 6.16 and 6.17 that the market price does not exhibit the same distribution as the fundamental value it was supposed to follow. The phenomena of 'fat tails' is clearly visible, extreme market returns are more probable than they should be (according to the efficient market theory). The summary statistics Table¹ 6.13 shows a clear indication about the persistence of stylized facts through 100 simulations runs with the same parameters.

If we consider the effect of lagged information on market price we can say that the market price is formed using orders based on current and the last seven fundamental values, e.g. :

$$P_{t+1} = f(P_t, V_{t+1}, V_t, V_{t-1}, \dots, V_{t-7}) \quad (6.3)$$

$$R_{t+1} = f(R_t, R_{t-1}, \dots, R_{t-7}) \quad (6.4)$$

From these equations we can infer that the market returns has a direct form of auto-dependence due to the fact investors that use lagged information. If we look at the autocorrelation diagram for the market returns, Figure 6.18, we see that returns exhibit non-linear autocorrelation (confirmed in repeated simulation runs - see statistics summary in Table 6.13). Moreover, if we increase the span of the lag parameter (in this case to 20) we can observe in Figure 6.19 that the non-linear autocorrelation of market returns becomes stronger.

One can observe that in Figure 6.18 and 6.19 the autocorrelation levels of absolute return (at the power 1.5) are almost double than the levels of autocorrelation when the maximum investors lag is 8. In either of the two cases, average and long

¹The Hurst statistic test parameter was used. A value H strictly between 0.5 and 1 indicates long-memory behavior of returns. Student t test considering the null hypothesis as Hurst parameter=0.5

Test 8		
Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	5000	Number of trading days
Nr agents	200	100% Informed investors (not speculators)
Information lag	[0; -8]	Each investor will have a (fixed) lag drawn randomly from the interval [0..8]
Bias a_p	1	Investors have perfect information
Bias b_p	0	

Table 6.12: Market with perfectly informed investors, with lag in information




Figure 6.16: Distribution of returns for the fundamental value resulting from a simulation using the specifications in Table 6.12




Figure 6.17: Distribution of returns for the market price resulting from a simulation using the specifications in Table 6.12

Summary statistics for 100 simulations	Excess kurtosis	Non-linear autocorrelation of absolute returns (HURST test)
Mean	2,99520934	0,69654308
Standard deviation	1,07042628	0,08612114
Student t value	27,981463	22,8216998

Table 6.13: Statistical parameters for excess kurtosis and long memory of returns for 100 tests with the parameters from Table 6.12




Figure 6.18: Autocorrelation of absolute market returns resulting from a simulation using the specifications in Table 6.12 and a maximum lag of 8




Figure 6.19: Autocorrelation of absolute market returns resulting from a simulation using the specifications in Table 6.12 and a maximum lag of 20

lag intervals, the non-linear autocorrelation is very persistent (its last for even 1000 time periods). In real markets, see the Introduction, we observe levels of autocorrelation that are even much higher (above 20%) and that decay rapidly and are usually not significant after a lag of 200 rounds. We have pointed out that volatility clustering (as well as extreme returns) can be generated realistically for very simple fact: the use of old information for the formation of new prices. Lagged information, even if used by unbiased investors, will generate non-linear autocorrelation in market returns. If we combine the lagged property of information usage with other biases, like a_p or b_p parameters, we will observe richer market dynamics whilst still maintaining nonlinear autocorrelation.

With this observation, we conclude that 'almost' efficient markets, where the differences between market prices and the fundamental value are very small, can exhibit stylized facts. This statistical effect is caused, in this case, by a simple lag in the investor information.

6.3 Volatility clustering and other "stylized" facts

Volatility clustering is one of the 'stylized facts' presented in chapter 2. Market return series that display this property are characterized by periods with different volatility levels: low/high volatility tends to be followed by low/high volatility. In classical studies, see chapter , such effects were explained by investors that had one of the following properties:

1. Simulated agents compute their future expected price using a function of past volatility like:

$$E_t(P_{t+1}) = f(P_t, \sigma_{R_t, \dots, R_{t-k}}, \dots) \quad (6.5)$$

2. Investors' change, for different reasons, their risk estimations and they switch between regimes of high/low expectations of returns which induces them to make bids/asks that generate high/low returns.

These approaches generate times series with volatility clustering but the result is actually embedded inside the model by its design. The two market models 1) and 2) can be approximated by an autoregressive function with generates market returns. Although this is how we can numerically estimate market return it does not consist in a proper behavioral, causal explanation.

In our simulator we have tested different mixes of investors and observed that chartists mixed with perfectly informed investors can generate such stylized phenomena. We run a simulation with mixed technical/informed investors and following parameters:

Test 9		
Parameter	Value	Explanation
N	500	Number of trading days
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
Nr agents	500	40% Technical traders 60% Perfectly informed investors
Memory length	3	Traders use rules regarding the last 3 returns (r_t , r_{t-1} and r_{t-2}).
μ_C	1%	Traders try to sell/buy with a random mark-up of
σ_C	5%	
Agent Wealth	Cash and Portfolio	2000 in cash and a portfolio of 10 stocks for each agent at the start of the simulation

Table 6.14: Simulation parameters for Test 9 - volatility clustering with technical traders

We observe, in Figure 6.20, that the market price sometimes departs from its fundamental value. This departure happens when a considerable high number of chartist trader have strategies that are activated. These chartists send orders to buy (or sell) and their collective action moves the resulting market price away from the fundamental. When the market price is below/above the fundamental value the perfectly informed investors will try to buy/sell the asset. To be able to restore the prices to their correct levels the informed investors have to put more pressure on the prices than the chartists investors do.

It is possible that chartist traders can absorb all (or most) of the stocks of the informed traders and they afterwards dominate the market price. At these moments, the chartists can self-coordinate as a group and are able to control the market price through their orders. The speed of growth and amplitude of the resulting price bubbles are determined by the expectations of chartists (the bigger the return expectations the faster the price will raise) and on their cash reserves (the more cash chartists hold the bigger can the bubbles get). We see bubbles appearing in Figure 6.20(at the beginning at the simulation and around simulation time 290) or in Figure 6.21 (starting at time 50 and ending at around 450). During such periods market prices go above/under fundamental values and fluctuate without any apparent logic. As one can observe bubbles can be negative when prices drop far below fundamental values. This is a natural phenomena, albeit not very frequent, when an asset can be highly underpriced because investors sell such stock in bundles when important opportunities arise. We recommend (Allen and Gale, 2000) for examples of positive and negative asset bubbles.




Figure 6.20: Example market price behavior with 40% chartists, 60% perfectly informed (market price in red and fundamental value in black) - simulation parameters from Table 6.14

Until the market price reaches again its fundamental values the market prices are virtually the result of orders coming from chartist traders. During these mispricing periods, the price quotes attached to orders are generated using different volatility levels (σ_C) than that of the fundamental value (σ_V). Therefore this is a possible explanation for why market prices have different volatility regime: different regimes implies different types of investors (with different return distribution expectations) that drive market prices. Because of this change in regime we can observe the volatility clustering phenomena. Looking at Figure 6.22 we can see the return time series for one example simulation test.

We can observe two compact and distinct periods: high volatility levels (time 0-120 and 300-400) and low volatility levels through-out the rest of the simulation. As argued in (Shiller, 1980) the volatility levels of the fundamental returns do not explain the volatility observed in the market returns series. In our case our fundamental return have, by design, a normal distribution.

The only determinant of these price anomalies is the self-coordinated² actions of chartist traders (all the other traders are perfectly informed investors). The appearance of these different volatility regimes, one dominated by informed traders

²This coordination is involuntary and it happens when a sufficient number of chartist use similar strategies for buying/selling. These strategies are activated in similar conditions and the chartists' orders have a significant impact on market prices.

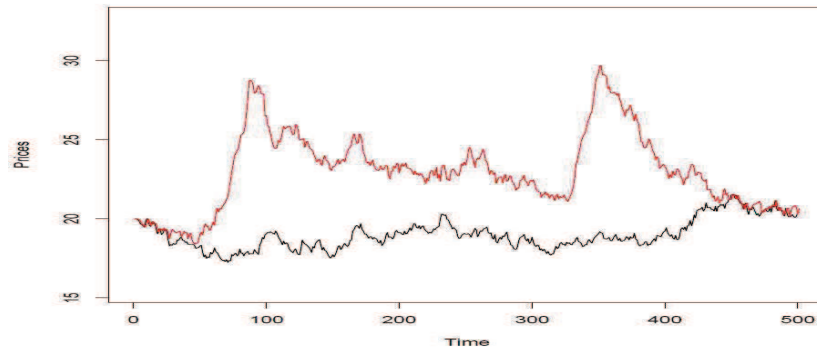


Figure 6.21: Another example market price behavior with 30% chartists, 70% perfectly informed (market price in red and fundamental value in black) - simulation parameters from Table 6.14




Figure 6.22: Example market price return series with market simulation parameters from Table 6.14

and another by technical traders, also influences the distribution of the market returns. We observe, in Figures 6.23 and 6.24 that market returns are leptokurtic, as opposed to the fundamental returns which are normal.

We know that the distribution of the fundamental value is normal and has the parameters that are maintained fixed for all simulations ($\mu_V = 0$ and $\sigma_V = 1\%$). The market returns have much more instances of extreme returns (2 to 4 times σ_V deviations). This is caused by the volatility of the expected returns of the chartist traders ($\mu_C = 1$ and $\sigma_C = 5\%$). Because the chartist investors have high return demands, they send orders that generate more volatility in market prices.

Other than increasing the probability of extreme market returns, the different volatility regimes create the effect of non-linear autocorrelation in returns. We can see in Figure 6.26 that normal returns are not correlated thus showing that the market cannot be predicted using only past returns. On the other hand, absolute returns are, in a certain measured, explained by past absolute returns (see Figure 6.25). This time, the levels of return non-linear autocorrelation have a level similar to the ones observed in financial markets (see Introduction chapter). We test our results for robustness and place the summary statistics in the Table 6.15.

In the Table³⁴ 6.15 we first show that the stylized facts (using statistical tests for excess leptokurtosis and non-linear autocorrelation of returns) are present in multiple runs of the market with parameters from Table 6.14). Moreover we need to test for the emergence of price bubbles. Therefore we compute, for every simulation, the average/median/maximum distance between the market price and the fundamental value using the equations 6.6.

³The Hurst parameter was used. A value H strictly between 0.5 and 1 indicates long-memory behavior of returns. Student t test considering the null hypothesis as Hurst parameter=0.5

⁴F is the fundamental value and P is the market price. We measure statistics about the distance between the market price and the fundamental value

Summary statistics for 100 simulations	Excess kurtosis	Autocorrelation of absolute return (HURST parameter)	Median $F_t - P_t$ (m)	Average $F_t - P_t$ (u)	Maximum $F_t - P_t$ (M)
Mean	1,39481349	0,81418867	5,708036	12,2303034	97,90651
Standard deviation	1,0222444	0,1016898	3,683139	8,98876249	86,06556
Student t-value	13,5981217***	31.4***	15,49775***	13,6062149***	11,37581***

Table 6.15: Robustness test for simulation with parameters from Table 6.14




Figure 6.23: Probability distribution of market returns - simulation parameters from Table 6.14




Figure 6.24: Probability distribution of fundamental value returns - simulation parameters from Table 6.14




Figure 6.25: Test 9 Autocorrelation diagram for absolute market returns - simulation parameters from Table 6.14




Figure 6.26: Test 9 Autocorrelation diagram for normal market returns - simulation parameters from Table 6.14

$$\begin{aligned}
 M &= \max(F_t - P_t) \\
 m &= \text{median}(F_t - P_t) \\
 \mu &= \text{mean}(F_t - P_t)
 \end{aligned}
 \tag{6.6}$$

We see that our results, as described in Table 6.15, verify the following hypothesis:

1. Maximum value is at least 10 times bigger than the median value ($M - 10 * m > 0$). If this is true than we have, in every repetition of our Test 9 simulation parameters, at least one price bubble (a market price much higher than the fundamental value). This test shows only that there are moments were market prices spike way above the fundamental value. Yet, such sudden market moves can last for only one period (if there is a sudden spike in market prices) and do not constitute proper bubbles.
2. To prove that our market generates, robustly, price bubbles, we measure the relationship between the average and the median distance between price and its fundamental value ($\mu - 1.5 * m > 0$). If this is true than the average value is biased because of extreme observations (high prices around the peak of the bubbles). This condition shows that bubbles appear. We have to also test that such bubbles break.
3. To prove that price bubbles appear and also crash we check if the average price distance is not higher than twice the median value ($\mu - 2.2 * m \leq 0$). If this is true than it shows that the periods with very high price peaks do not last very long - hence bubbles crash. If price distance would stay high then the average measure would be affected we increased considerably above 2 times the median.

We make hypothesis tests, for our three bubble measures, and observe in Table⁵ 6.16 that our predicted results, a population of investors with sufficient chartists create price bubbles, are statistically significant and appear consistently in our simulation outputs.

With these robust observations we have shown that price bubbles (and also volatility clustering, along with excess kurtosis) are caused by a chain of events that we summarize here:

1. Prices fluctuate along with fundamental values. Market price is controlled by informed investors at follows the asset's fundamental value.
2. A specific pattern of market returns appears (by chance) and a significant number of chartists send limit orders, with involuntary coordination, which shift the new market price away from its fundamental value.

⁵We could not reject the H_0 hypothesis ($u=2*m$) and this implies that the average is not significantly bigger than twice the median value.

Summary statistics for 100 simulations	Test $M > 10 * m$	Test $u > 1.5 * m$	Test $u > 2 * m$
Mean	51,8515	3,59499	0,71654984
Standard deviation	71,0521	5,71153	5,50356029
Student t-value	7,29767***	6,29427***	-0,77749239

Table 6.16: Hypothesis tests for the presence of price bubbles

3. When the market price is below/above its fundamental value the informed investors will start buying/selling assets motivated by the potential profit.
4. If more chartists are willing to buy/sell assets than the fundamental investors can sell/buy then control of prices will move towards chartists. This period, while informed investor try to profit of the sudden mispricing, lasts until fundamental investors consume all of their resources (either all their cash from buying undervalued stocks or they sell all their overpriced assets).
5. From this moment the informed investors have exhausted their resources for new orders and the market price is completely created by the orders of the chartists. Credit lines are irrelevant since both types of investors can get credit for further investment.
6. Market prices now fluctuate as a function of the expected returns of the chartists and of the volatility of these expectations (that volatility is usually very different from the distribution parameters of the fundamental value). We are now in a different volatility regime where chartist trades generate market prices. Because chartists have trend following strategies market prices will evolve in trends and thus form bubbles.
7. Market price bubbles persist as long as chartists have enough resources and there are enough counterparties (usually also chartist investors) that enable the exchange of assets at prices away from fundamentals. When these resources diminish or counterparties disappear, the orders of the informed investors (with their remaining assets) will drive back the price towards the fundamental levels. This step can also happen because of an accidental reversal in market price/fundamental value trend. When such a reversal arrives (can happen by chance) a part of the chartists will have their selling strategies activated and will start trading in the new price direction.
8. After bubbles crash and when market price reach the level of the fundamental value the limit order book will be filled with orders from the informed investors. The actions of these informed investors can have a price stabilizing effect.

We can now complete the answer to the research question #1 by adding that a biased market, that sometimes exhibit the 'stylized' facts, can be created in a number of ways: combining different types of investors (informed or not informed) or having informed investors that use lagged information (see chapter 6.2.1).

One can make the argument that since market prices return to fundamental values, chartist traders will loose money and soon disappear from the market. We will see in chapter 6.5 that this is not always true and the survival of different strategies depends on the horizon of investment and mix of different types of investors.

6.4 Bubbles: how prices can grow above the asset's fundamental value

In this chapter we look for answers to the research questions #4 and #5. More specifically we search for the basic behavioral causes of the creation and destruction of price bubbles.

Market bubbles can be defined as phenomena during which market prices rise abruptly, at levels above what the fundamental information can explain, and this rise is followed by an abrupt fall in prices. The essential element in the definition of a bubble is that it consists in a price rise and fall and that the rise is far **above** the fundamental value. Unlike classical studies where the asset's fundamental value is unknown, using our research methodology we can quickly observe such phenomena by comparing the market price with its fundamental value (a market with risk-neutral perfectly informed investors is efficient when market prices are at the asset's fundamental value level). In the previous chapter we showed how prices are influenced by informed investors or chartist investors (non-informed profit motivated). For a bubble to exist, investors have to send orders to sell and buy and price quotes that are higher and higher than the fundamental values.

We first look at a market with informed investors. For such investors to send orders to buy or sell assets above the fundamental value they should have biased information. We have seen that investors with an a_p bias generate a market with excess volatility, but market prices deviate only shortly from fundamentals (the market price returns fast towards the fundamental value). When informed investors have a b_p bias they will generate market prices which will rest at a relatively constant distance above or below the fundamental value. With such a model specification we will not be able to obtain price bubbles since the price expectations of the informed investors is "correctly" linked only to the fundamental value. One "creative" way would be to transform the parameters a_p and b_p into time functions such as:

$$a_p(t) = a_p(0) * (1 + \alpha * \sin(t)) \quad (6.7)$$

6.4. Bubbles: how prices can grow above the asset's fundamental value

$$\beta_p(t) = \beta_p(0) * (1 + \beta * (1 + *cost(t))) \quad (6.8)$$

The market prices created by informed investors with dynamic biases such as $a_p(t), b_p(t)$ would go above the fundamental values and then down in the shape of a geometrical sinus function. We can even control very well the amplitude and duration of price bubbles by acting on the periodicity of the *sin* and *cos* functions. While this approach would create the desired effect, price bubbles, it would not provide a sound financial explanation.

We can affirm that for market price to go above the fundamental values we should have at least one of the conditions:

1. When the market price goes above the fundamental value (for whatever reason) the informed investors will sell the overpriced asset. Their counterparties will be investors that do not base their price expectations on the fundamental value, e.g. chartist traders. Having bought all, or most, of the traded assets these non-informed investors will continue to bid up the price and create a speculative bubble. This bubble will grow with certain a speed and up to a height depending on the aggressiveness and cash reserves of these speculative traders.
2. Informed investors send orders to sell at higher and higher prices and "other" investors should be willing to buy at such prices. When the market price passes above its fundamental value the informed investors will NOT try to sell the asset (to cash in this capital gain). Instead they will speculate this growth and try to buy the asset to sell it at an even higher price. Investors with such behavior (a mix between an informed and a chartist trader) will be called "SITH" investors and will be studied later on.

In the first case we consider a market where perfectly informed investors are mixed with chartist investors. Such a mix of population, 40% chartist and 60% informed investors have already proved to consistently generate bubbles in Test 9 (with summary statistics in table 6.15). One can argue that the population mix we have chosen is arbitrary and might not be realistic. Therefore, we will perform Test 10 where we will vary the proportions of informed and chartist investors and try to determine a minimal concentration of chartist traders that are able to generate market bubbles.

We vary X, the percentage of informed investors, from 100% downwards and for each level we run 50 simulations. We want to find out what is the minimal percentage of chartist traders that can produce market bubbles. It is important to consider that all investors have the same level of wealth at the beginning of the simulations. For all levels of X% we run 30 simulations and record, for each, the

Test 10		
Parameter	Value	Explanation
μ_V	0%	No growth in fundamental value
σ_V	1%	1% standard deviation
N	350	Number of trading days
Nr agents	500	x% Perfectly Informed investors (100-x)% Chartist investors (not informed)
Memory length	3	Chartist investors will invest using speculative rules of concerning the last 3 market returns
μ_C, σ_C	(1%,7%)	Traders try to sell/buy with a random mark-up of $R \in N(\mu = 1\%, \sigma = 7\%)$
Wealth	1000,10	Each investor starts with 1000 in cash and 10 stocks. IPO price of the stock is 20

Table 6.17: Simulation parameters for Test 10: we vary the proportion X from 100% downwards

Percentage of Informed investors (x%)	Percentage of Chartist investors (1-x)%	Median bubble size	Average bubble size	Maximum bubble size
70	30	0	0,31	3,86
69	31	0	0,91	9,45
68	32	0	0,69	5,58
67	33	0,015	1,7558	14,069
66	34	0,04	4,42	40,13
65	35	1,19	5,49	38,115
64	36	1,91	9,306	56,182
63	37	3,59	8,6267	50,584
62	38	5,27	9,65	65,43
61	39	6,64	13,145	70,53
60	40	6,47	18,791	128,4

Table 6.18: Summary statistics for results of multiple simulations on different levels of relative proportions of chartists to informed investors.

6.4. Bubbles: how prices can grow above the asset's fundamental value

average/median/maximum value of the distance between the market price and the fundamental value.

We list the test results at levels starting from 70% informed investors (a lower percentage chartists cannot persistently move away prices from their fundamental values). Our results indicate that a growing number of chartist investors manage to slowly move prices away from its fundamental value. In Figure 6.27 we show an example time series with 31% chartist traders where we observe very small price bubbles. At a level of 35% chartist investors we notice that the average bubble size is 4.42 units (meaning an average price increase during a bubble of 22% over the fundamental value). Bubbles are now bigger and more persistent, as observed in Figure 6.28, and they can move prices at an average maximum of 40 units (at these bubbles peaks market prices are double the fundamental value). It is clear that increasing the fraction of chartists will generate market series with bigger and more frequent bubbles.

It is important to understand that the minimal conditions of price bubbles depend on a number of factors. First of all, it is important that the mass of chartist investors have enough cash resources to buy at least all the stocks of the informed investors. This is a necessary but not sufficient condition for the emergence of a price bubble. Even if chartists manage to hold all (or most) stocks they will still need extra cash to be able to fuel the price trend and buy stocks at increasing prices. Moreover, after informed investors have sold their overvalued trading can continue only if enough heterogeneity exists between chartist investors.

We can observe that, with 31% chartists, small price bubbles emergence (time periods 95-130 and 135-160). When increasing the fraction of chartists, like in Figure 6.28, bubbles appear more frequent. To understand in more detail why and how these price deviations are formed we examine the transfer of assets during a simulation.

In Figure 6.28 we observe more features of an example simulation of a market with the same parameters as before, where 35% of the investors are chartists. The market price forms more and bigger bubbles. When the bubbles rise informed investors sell their (overvalued) assets and trading continues between chartist.

We see, in Figure 6.29, that most assets are quickly transferred from the informed to the chartists investors. After this moment market prices move away from fundamentals. As indicated by the excess volatility graph, in Figure 6.30, the simulated periods with more volatility are those when trading is mostly done between chartists. The traded volume graph, from Figure 6.31, indicates that trading has a lower volume when prices are disconnected from reality (chartist driven). This happens because informed investors are not in the market anymore. With these robust observations we give a few qualitative answers for the existence of price bubbles.

Research question #4: *What are the market micro conditions for the emergence of price bubbles?*




Figure 6.27: Example simulation results in a market with 31% chartist (see table 6.17 for the market parameters)




Figure 6.28: Example simulation results in a market with 35% chartist (see table 6.17 for the market parameters)

6.4. Bubbles: how prices can grow above the asset's fundamental value

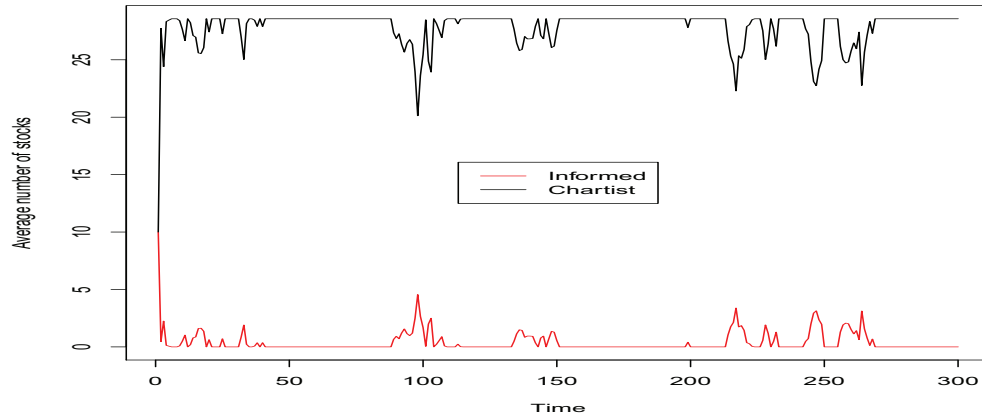


Figure 6.29: Portfolio holdings of investors in simulation with a market with 35% chartist (see table 6.17 for the market parameters)




Figure 6.30: Excess volatility of market returns in simulation with a market with 35% chartist (see table 6.17 for the market parameters)




Figure 6.31: Trading volume in simulation with a market with 35% chartist (see table 6.17 for the market parameters)

Answer: The necessary conditions for the appearance of price bubbles are:

- a. The existence of investment strategies where the prices of bids are not linked with the asset's fundamental value
- b. Investors with strategies from a) must have enough cash to
 - i. buy most stocks held by investors with fundamental based strategies.
 - ii. be able to trade at higher than fundamental level prices

Price bubbles can be initiated by random movements of the fundamental value that result in an overestimation of the asset. In the presence of a group of uninformed (or badly informed) investors this overestimation can generate a coordinated buying action. Informed investors will sell the overpriced asset but chartist investors absorb the entire offer. If chartist still have enough cash they can continue trading, among them, the asset at higher and higher prices. The price bubble will burst when there is no more cash to buy stocks at high prices.

From the simulation runs we have tested we can offer a few quantitative hints for the market conditions that can bring a price bubble:

1. Ceteris paribus, the lower the asset price the more chances we have for bubbles (this is due to the relative buying power of chartist - lower prices can help chartists acquire more stocks)

6.4. Bubbles: how prices can grow above the asset's fundamental value

2. Ceteris paribus, the less the number of tradable stocks the more the chances of price bubbles (because chartist have fewer stocks to buy)
3. Ceteris paribus, the cheaper the money available for stock speculation the more chances we have of a price bubble.
4. Ceteris paribus, the more chartist investors (or liquidity providers) are the more chances we have for price bubbles (few chartist imply low liquidity and few counterparties for trading at high prices).

Because a financial asset can also be underpriced we have sometimes observed a "reverse" price bubble. In this case the asset can be underestimated and the price can fall way below its fundamental value. We observe such behavior in Figure 6.32 where at the beginning, time between 20 and 50, the price falls below its economic value and this mispricing persists.

The traded asset rests undervalued because the informed investors simply do not have sufficient financial power to move the price to its correct level. We can, in Figure 6.33 between $t=20$ and $t=50$, observe that informed investors buy as many undervalued assets as they can (since it can be bought at a discount).

Unfortunately these informed investors exhaust their cash reserves and the price is afterwards controlled by the uninformed chartist investors. Although the behavior of speculative investors is symmetric in both directions of market returns (chartists will sell if they believe prices will go down as much as they will buy if they believe prices will go up). Such "**reverse bubbles**" do not appear or are very limited in their amplitude (the maximum distance between the fundamental value and the market price). This is due to an asymmetry that arises from the ownership of the majority of assets, during a price bubble. In a normal (rising) price bubble assets move from informed investors (that sell the overpriced asset) towards the portfolios of uninformed investors. After a sufficient rise in prices, the bulk of the tradable assets are held by investors that have nothing to do with the fundamental value (such asset transfer also happens in real-estate bubbles where speculators end up owning most of the land and buildings). Therefore prices can continue to move and rise (the only limit being the available cash of trading counterparties). During the opposite phenomena, when market prices are forced down (and the asset is underpriced) informed investors will buy more and more assets. Therefore, the informed investors will eventually buy assets (at a discount) with all of their cash. When the bulk of the assets is in the hands of the informed investors the only new prices will be set by these informed investors - of course, at levels indicated by fundamental values. In such moments the group of investors that holds most assets is informed (in opposition to the situation where the bulk of the asset was held in portfolios of uninformed investors). This asymmetry in stock ownership, when undervalued informed investors step in and when overvalued informed investors step-out, is why a reverse price bubble will be smaller in amplitude (than a normal price bubble).




Figure 6.32: Example simulation results in a market with 31% chartist (see table 6.17 for the market parameters)

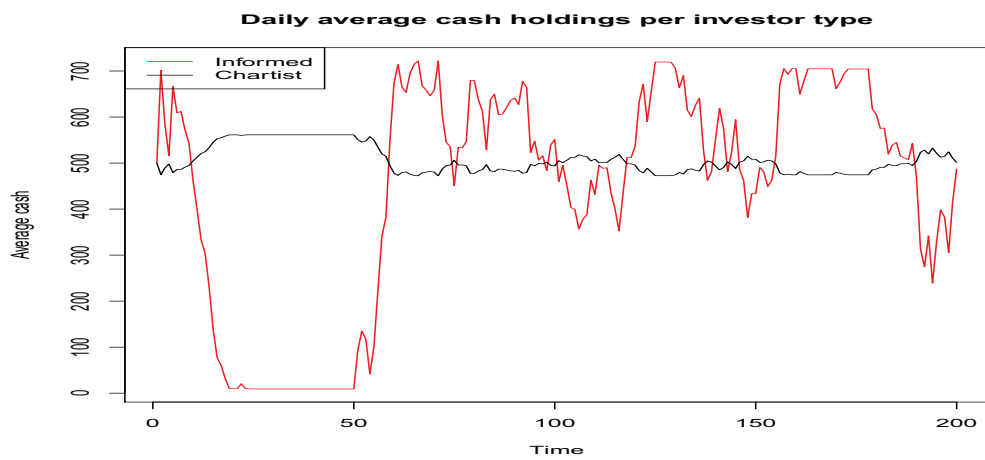


Figure 6.33: Average cash holding portfolio for each trader type inside the simulation example from 6.32 - with parameters found in Table 6.17

The extent and the duration of a price bubble depend on the total wealth of the chartist traders (or the type(s) of traders that control the price during the bubble) and on their aggressiveness (how much do they want to get as returns). If the chartist traders expect higher returns, or more volatility returns then they will bid up or down prices in a faster way, like in Figure 6.34.

We observe that the bubbles from Figure 6.34 exhibit prices that rise up to 17 times from their original level. The level of such a growth is limited by cash of the chartist as well as by the total number of tradable assets. With sufficient simulations a quantitative relationship can be found between the amplitude and duration of bubbles and the micro-level market conditions.

In this chapter we have showed a simple market composition that can cause price bubbles, namely using chartist traders and perfectly informed investors. It is important to underline that such bubbles can be generated by any type of uninformed traders. Price bubbles can occur as long as such investors have positive future price expectations and have enough financial resources to dominate and sustain trading at high prices.

A possible critic to these results consists of the lack of short-selling possibilities. Short-selling implies that informed investors could continue trading with "borrowed" assets in an attempt to drive prices down to fundamental levels. This option implies that investors are able to borrow money. Therefore it would be fair to say that even chartist investors could borrow. Thus short-selling could decrease the possibility of price bubbles only if chartist investors would not be able to borrow and continue speculating. For this reason, of short-selling effect being eliminated by investor credits, we decided not to allow such borrowing capacities. We believe that with or without complete borrowing opportunities price bubbles can appear with the same frequency. From a market quality point of view it is, of course, better to allow only short-selling and block possibilities of buying stocks with credit.

6.5 Survival and evolution of agents

In this chapter we will answer the research questions #2, #3 and #6. The main issue we investigate concerns the survival of different investor behaviors in the market. According to financial theory, as discussed in detail in chapter 1.3, investors which are not rational or not perfectly informed cannot survive in markets because they will make mistakes and lose their money to better informed investors.

To establish which investor types will survive and which will not it is important to clearly define how we will judge survivorship. In a market that exhibits price bubbles, asset holders will show significant returns before the bubble is peaking but would have almost zero returns (or negative) once the bubble bursts. If we compare wealth returns between different investor types we will observe different results before, during and after a price bubble. Moreover it is important to consider




Figure 6.34: Example simulation with bubbles and crash from a simulation with parameters like in Table 6.17, but changing $\mu_C = 8\%$ and memory length of 2 (instead of $\mu_C = 7\%$ and memory length of 3)

the openness of the observed market. In the case of our LUMA simulator the market is considered to be closed since no new investors enter the market during a simulation run. This is not the case for a real financial market where new investors arrive constantly in the markets. In our tests we assume that if new investors enter the market they do not change the initial proportions of investment strategies.

The research questions we want to answer (2, 3 and 6) all revolve around a central point: do rational fundamental investors (that want to keep market prices close to their fundamental values) eventually drive out of the markets all other types of investors (irrational, speculators, etc)? Apparently, the proposal of economist Milton Friedman, in (Friedman, 1953), is still true for most researchers: "*People who argue that speculation is generally destabilizing seldom realize that this is largely equivalent to saying that speculators lose money, since speculation can be destabilizing in general only if speculators on the average sell when the currency is low in price and buy when it is high.*" The common understanding of this statement is that speculators gain wealth by doing the exact opposite action: buying low and selling high (where the fundamental value is the reference). Moreover, the general assumption (as formalized by the efficient markets theory) is that investors that do not speculate in such a way, meaning in trying to redirect market prices towards fundamental values, will eventually lose money.

We contribute to this dialogue of ideas by saying that a non-fundamental specu-

lative agent can, in opposition to having a losing strategy (buying high and selling low), manifest two other profitable strategies:

1. (classical) Buy low and sell high - a "stabilizing" strategy that provides returns for the risk of market prices not returning fast to their fundamental levels.
2. Buy high and sell higher - price bubble profit
3. Sell low and buy lower - reverse price bubble profit
4. Combining 1), 2) and 3)

We will now show that strategies 2) and 3) can be profitable and investors using such strategies can persist and flourish in financial markets, when the right conditions are achieved. To clearly differentiate between these three abstract trading strategies we have made a graphic visible in Figure 6.35.

Strategy 2) and 3) are in the trading arsenal of our chartist traders. Strategy 2) can be characterized as a positive trend amplifier and strategy 3) as a negative trend amplifier. We will now run a number of simulations inside our LUMA research tool and we observe the evolution of agent wealth.

We first look at a market where all investors use strategy 1). Such investors differ only by the level of their biases. We are interested to observe if these investors make more money when they are less biased or not. Because we want to observe the transfer of wealth between agents we will simulate assets without growth. In the first test we compete perfectly informed investors against biased informed investors.

We run 50 simulations of test 10 a) (where biased investors dominate) and we look at the evolution of the Sharpe ratio of the investment strategy of each investment type. Market prices are often moved above the fundamental level because of the actions of biased investors (see Figure 6.36). Intuitively they should sometimes be moved under the market price (if the initial distribution of biases creates a pessimistic population). This underpricing does not happen because, even if in minority, the perfectly informed investors can absorb more stocks at decreasing prices - thus stopping the undervaluation. As we can observe in Figure 6.36, biased investors dominate and the asset overpricing persists (trading is done between the biased investors). Therefore perfectly informed investors sell their overpriced assets and are driven out of the market and hold only cash (see Figure 6.37). Even if the Sharpe ratios in Table 6.20 are similar (0%) this is because the growth of the asset was 0%. The conclusion is that biased investors, when they persistently dominate the market, will capture the growth of the asset (and perfectly informed investors will exit the market).

We make the same simulations and introduce positive minimal expected return (r_{min}), for all investors, between 1% and 10%. All investors try to sell/buy the asset

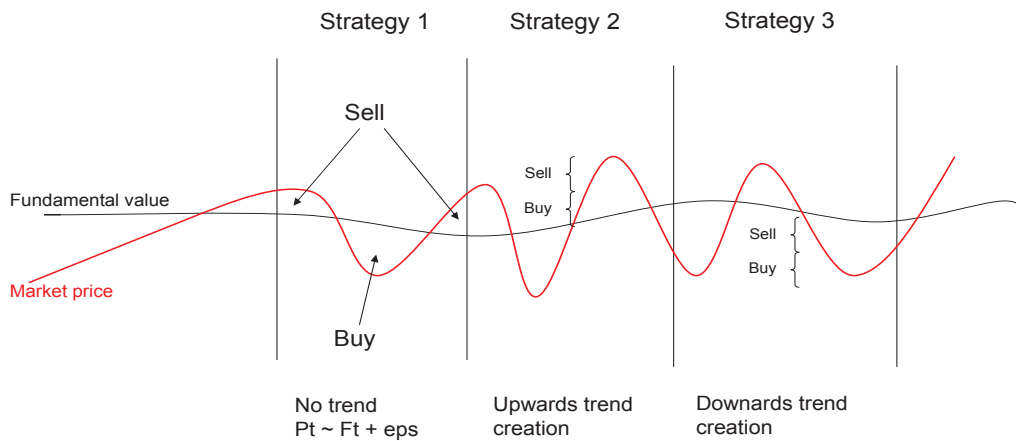


Figure 6.35: Presentation of abstract trading strategies

Parameter	Value	Explanation
Test 10 a (b)		
μ_V	0.0%	No daily growth in the fundamental value
σ_V	1%	1% standard deviation
N	500	Number of trading days
Nr agents	500	30% (70%) GOD (perfectly informed) investors 70% (30%) Biased informed investors
a_p	[0.5, 1.5]	Biased informed investors have a parameter of information amplification randomly chosen for this interval
b_p	[-1, 1]	Biased investors will, on average, not skew their information
Wealth		Each investor start with 1000 in cash and 10 stocks

Table 6.19: Simulation with rational investors, biased and not biased (which are dominant)




Figure 6.36: Market price movement - simulation parameters 10 a) in Table 6.19




Figure 6.37: Average wealth evolution per investor type - simulation parameters 10 a) in Table 6.19

with a premium/discount over their perceived correct value. This generates a discrete price movement⁶, see Figure 6.38, since all investors wait for the fundamental value to move sufficiently away from the market price. Relative to the previous test a difference arises in the ownership of the stocks. Even if the asset is on average overpriced the perfectly informed investors won't be able to sell all of their stocks because of the premium they demand is incompatible with the discount the biased investors expect, see Figure 6.39. Therefore the perfectly informed investors will hold some assets and also capture, in part, the returns of the underlying asset. If the market overpricing would be more than the premium plus the discount expected all stocks would be transferred to the biased investors.

In Table 6.20 we observe that investor haves similar Sharpe ratios. It is important to underling that there are three distinct situations:

- without positive minimal expected return (r_{min}): Biased investors own most stocks and capture the underlying asset's growth
- with positive minimal expected return (r_{min}):
 - If the overpricing is high enough biased investors will own most stocks
 - If the market overpricing is not high enough perfectly informed investors will hold some stocks (because they won't be able to sell them)

Regardless of the situation ($r_{min} > 0$ or not) the biased investors, because they dominate the market, are able to capture most of the stock's returns. Even if they are better informed the GOD investors are not able to take advantage of the biased investors.

We look now at a market where perfectly informed (GOD) investors dominate. We rerun 50 simulations for Test 10 b (where 70% of investors are perfectly informed) and display the average Sharpe ratios in Table 6.21 below.

This market can be called efficient because the market price follows the fundamental value (with an error with 0 mean). In the Table 6.21 we can observe the average distance to the fundamental value and in this case, of null minimal expected returns, is very close to 0. This efficiency is assured by the massive presence of perfectly informed investors which have null minimal expected returns ($r_{min}=0$). Because transactions between all investors take place at the newly formed market price than all the investors take advantage of the fundamentally correct market price, even if some investors have biased beliefs. We see in the table above that the investor's Sharpe ratio's are not statistically significant different than 0. This implies that transfer of wealth between the two types of investors is not happening (or it is at an extremely slow pace). If we run the same simulation and introduce a

⁶By discrete price movement we mean a price series with periods of zero change in prices (due to lack of transactions).




Figure 6.38: Market price movement in a market where all investors have positive minimal required returns - simulation parameters 10 a) in Table 6.19 and $r_{min} > 0$




Figure 6.39: Evolution of ownership of stocks - simulation parameters 10 a) in Table 6.19 and all investors have positive minimal required returns, $r_{min} > 0$

Summary statistics for 50 simulations for Test 10 a)	$r_{min} = 0\%$	T-Stat	r_{min} in (0,10]%	T-Stat
Mean Sharpe ratio S_i for perfectly informed investors	0		-0,11574451	-0,76585911
Mean Sharpe ratio S_b for biased investors	-0,16078453	-1,37661581	-0,08714432	-0,99584931
Mean Sharpe ratio S_b for buy & hold investors	-0,1611071	-1,38096891	-0,09406803	-0,8224964
Average distance between fundamental and market price	-0,52117657	73,3530412	-0,27527378	15,3321041

Table 6.20: Sharpe ratios for test 10 a) (biased investors dominated) when investors are and investors have $r_{min} = 0\%$

Summary statistics for 100 simulations for Test 10 b)	$r_{min} = 0\%$	T-Stat	r_{min} in (0,10]%	T-Stat
Mean Sharpe ratio S_i for perfectly informed investors	-0,00328227	-0,02680143	-0,12056442	-0,96329246
Mean Sharpe ratio S_b for biased investors	0,04915127	0,44385786	-0,12136926	-0,96914497
Mean Sharpe ratio S_b for buy & hold investors	0,04862858	0,60990783	-0,11918663	-0,95289965
Average distance between fundamental and market price ¹⁰⁵	-0,00214254	-2,24004638	-0,20050505	-10,1108777

Table 6.21: Summary statistics for test 10 b) - perfectly informed investors dominate

positive minimal expected return, the perfectly informed investors act in the market only when prices stray away from fundamentals above the risk premium level they expect. Because the biased investors also expect a minimal return very few trading takes places between the two investor classes and the market prices will rest biased. Like in the test 6a (with $r_{min} > 0$) biased investors, even if they are in a minority, control market prices and continue trading because of their heterogeneity. This fact is proved by the significantly negative mean average distance between the fundamental and the market price (value of -0.2 in the Table 6.21).

Research question #2: *Can a biased investor(s) survive in a financial market?*

Answer: *Yes. If the market prices are persistently upward biased then better informed investors will sell the overpriced asset. If their selling action is useless, and prices remain biased, then control of market prices will be maintained by biased investors. These investors will continue trading, at biased prices, and consequently capture the returns of the fundamental value. If well informed investors have sufficient resources they can maintain prices at fundamental levels. This good market pricing benefits everybody, even biased informed investors. Therefore, the transfer of wealth from biased to non-biased investors does not happen or it happens very slowly. Biased investors can therefore survive in markets. Moreover, if we consider that the cost of biased information is less than that of perfect information, we say that better informed investors will loose more than the biased informed investors.*

A market dominated by informed investors can be called 'efficient', by the definition of (Fama, 1970). Looking again at this financial market we draw the attention to a particular aspect: a market dominated by perfectly informed rational investors is efficient. The efficiency of this market benefits not only the aforementioned investors but it also serves the investors which have biases (because of personal fallacies or of their unwillingness to pay the price for the correct information). We can then ask a simple question: Why would perfectly informed rational investors (GOD) be willing to pay a price for their perfect information when they do not obtain superior gains than investors with less information? It is interesting to point out that such an intuition was exposed 32 years ago by (Grossman and Stiglitz, 1980) and remained known in the financial literature as the *paradox of impossible informationally efficient markets*.

Therefore we can consider that, in an efficient market, perfectly informed investors can make a rational decision and decide to stop paying the cost of perfect information and still have the same returns. When the majority of investors takes this decision, the market is transformed and it starts being 'less' efficient.

What would happen in a market where non-fundamental based strategies dominate prices? We simulate in Test 11 such a market where chartist traders dominate. The parameters of this test are found in Table 6.22.

In this type of market, prices depart frequently and consistently from their fun-

Parameter	Value	Explanation
Test 11		
μ_V	0%	No growth in the fundamental value
σ_V	1%	1% standard deviation
N	500	Number of trading days
Nr agents	500	15% GOD (perfectly informed) investors 15% Biased informed investors 70% Chartist investors
μ_C	2%	Profit expectations of chartist investors
σ_C	8%	
a_p	[0, 5; 1.5]	Biases of informed investors (exaggeration of information but no skewed bias)
b_p	[-1; 1]	

Table 6.22: Test 11 - Simulation of a market dominated by chartist traders and with perfectly informed traders

damental values. This is due to the trend speculation action of the chartist traders. We observe an example simulation in Figure 6.40 with the evolution of prices and of the average wealth of each investor type. The wealth of chartist investors fluctuates heavily because of the trends they create. In Figure 6.43 we see that during the periods when the market price bubbles, the majority of the stocks are held by chartist investors. Most importantly, when prices fall back down most stocks are bought again by biased informed investors. This happens because biased investors accept to buy the stock at a premium from the fundamental value. In conclusion, trading happens in a few steps: a) when prices are very close to fundamentals everybody trades, b) when prices move away from fundamentals only biased and chartist investors trade, c) when prices start to bubble chartists hold all stocks and trade d) when prices fall the first to buy stocks from the chartists are the biased informed investors (and next are the perfectly informed investors).

Because biased investors are able to trade most often, and especially with the chartist investors, they manage to speculate and capture the wealth of chartist traders. In Table 6.23 we show the summary of Sharpe ratios for all the simulation runs of Test 11.

Both the perfectly informed and the biased informed agents manage to make more than the buy & hold returns, at the expense of the chartist trades. But the biased investors manage to earn on average the highest risk-adjusted returns.

Research Question #3 *Can biased investors earn more than non-biased investors? If yes, in what conditions?*

Answer: *In a financial market the price can be moved up and away from their fundamental values when investors with biased information (or without any information) have more financial power than the well informed investors. In such*




Figure 6.40: Market and fundamental value - test 11 simulation parameters in Table 6.22




Figure 6.41: Average wealth per investor type - test 11 simulation parameters in Table 6.22




Figure 6.42: Average cash holdings per investor type - test 11 simulation parameters in Table 6.22




Figure 6.43: Average size of stock portfolio per investor group - test 11 simulation parameters in Table 6.22

Summary statistics for 100 simulations for Test 11	Sharpe Ratio	T-Stat
Mean Sharpe ratio S_i for perfectly informed investors	0,470085	6,219227
Mean Sharpe ratio S_b for biased investors	1,77264953	30,7699281
Mean Sharpe ratio S_c for chartist investors	0,13588381	3,1067333
Mean Sharpe ratio S_b for buy & hold investors	0,14916993	4,89394981

Table 6.23: Summary statistics for multiple runs of Test 11

situations well informed investors can try to speculate the mispricing but it can happen that they do not have enough resources to re-establish correct prices. If informed investors decide to go against a biased market they will take too much or too less risk and can make suboptimal risk-adjusted performances - as long as market prices do not revert to their fundamental values in a realistic time frame (or at least not before the investment horizon of well-informed investors). Under such circumstances the biased investors can survive and earn more than the perfectly informed investors. Moreover, from a pure competition point of view, in a market with biased prices (where well informed traders stop trading) the biased informed investors are left alone to compete and capture the wealth of non-informed investors like noise or chartist traders. Therefore, when a market strays persistently from efficiency, rational well-informed investors can be motivated NOT to engage against the market thus providing more conditions for market mispricing. In a different language: it does not always pay (enough) to speculate in favour of the objective fundamental value.

Such observations contradict the old adage saying that, in the long-term, "speculators" (meaning biased or not-informed) investors will be driven out of the markets by more 'rational' and better informed investors. This idea was commonly defended with the assumption that more "smart" money will be attracted to such a non-efficient market in order to seize the profit opportunities. Unfortunately the reverse of such an argument is possible too: more biased or not-informed investors can also be attracted to a market that is not dominated by 'correct' information and where profit opportunities can exist without the need to buy information.

It is our belief that the investors which make the biggest returns are those able to take over, as a group, the formation of price and maintain this dominance for sufficiently long periods (regardless of the strategy of common value reference they share). In other words the winning (and hence surviving) investors are those that dominate price formation for sufficient enough periods and not those investors that

try to establish correct 'prices' (based on economic fundamentals). Therefore, an investor that knows when to switch strategies (and always be in the dominant group) can make superior gains than investors that always stick to one fundamentals-based or pure speculative strategy. The implications of such mixed pure-profit strategies can sometimes be destructive (huge price crashes, bankruptcy, etc.) but such effects are not factored in the expected returns calculations of investors.

From these results we arrived, in a rigorous manner, at a conclusion that can also be extrapolated from some remarks of well known market investors like:

1. "It's not whether you're right or wrong that's important, but how much money you make when you're right and how much you lose when you're wrong."
George Soros
2. "Markets can remain irrational a lot longer than you and I can remain solvent."
John Maynard Keynes

We can now ask why well informed investors do the following actions:

1. Buy (sell) undervalued (overvalued) stocks when the price bias can persist long enough for their transaction not to become profitable (during their horizon of investment)
2. Continue to pay a price for better information when they cannot make a better risk-adjusted performance than the investors with less information (here the biased investors).

From these observations we naturally infer that it may be more profitable for well-informed investors to adopt a **bivalent or mixed strategy**:

1. Go with the market when the prices are biased in order to profit directly out of such mispricing (buy when prices increase and sell when they decrease, relative to the fundamental value). This implies a chartist strategy
2. When the mispricing passes over a certain threshold (when the risk of a significant price reversal is too great or the expected return is high enough) the investor changes strategy towards a fundamental well-informed strategy (and tries to profit of the mispricing by reverting the market price towards its fundamental value).

To see if this type of mixed chartist-fundamental strategy can survive (and make extra profits) we construct a new type of investor. Because this investor is able to revert the market price back to its fundamental value yet he still goes with the trend (adopts a chartist strategy up to a certain threshold) we will call him **SITH**. A **SITH** investor combines the strategy of a chartist and a perfectly informed investor. The parameters that govern the behavior of a **SITH** investor are:

a) Chartist strategy: μ_C, σ_C and memory length parameter L that govern the price expectations and the formation of speculative buying and selling rules (see on page 97 for complete description of a chartist strategy)

b) Threshold $x\%$, $x = \text{abs}(\frac{P_t - E_t(F_t|I_t)}{E_t(F_t|I_t)})$ representing maximum pricing distance, between the agent's expected fundamental value (conditional on the investor's information) and market value, above which the SITH changes strategy from a chartist to an informed strategy.

If an asset is fairly well priced, meaning the price distance to the fundamental value is less than $x\%$, the SITH will try to ride trends - adopting a chartist strategy. When high enough trends develop, the distance to the fundamental value is more than $x\%$, the SITH investor will switch to a fundamental informed strategy. We simulate a market with biased informed, perfectly informed and SITH investors.

We run a few simulations with the parameters of Test 12 and observe the same evolution of prices as in Test 10 a) (which was made with the same parameters except for the presence of SITH investors).

As we can see in Figure 6.44 and 6.45 trading is plentiful and the asset is still overpriced (because of the skewed optimistic bias of informed investors).

Because biased investors manage to persistently overprice the asset the perfectly informed investors are not able to adjust prices and end up holding only cash (see Figure 6.47). Trading continues between SITH and biased investors. The SITH investors speculate any trends that exist in market prices (which are a biased version of the fundamental value). They are able to profit off of these trends because the biased investors are always counterparties for the trend-following trades. Because SITH investor also have the perfect information they avoid creating big price bubbles. Moreover, if the biased investors push the prices to high the SITH investors will speculate and sell the overpriced assets (which they buy again at lower prices).

The average distance between the fundamental and market price is negative. As we can see in Table 6.25, this implies that the market consistently overprices the asset. The summary of Sharpe ratios show us that perfectly informed investor hold only cash (because they sell when the asset is overpriced). The SITH investor are able to make returns above the buy and hold level at the expense of biased investors.

It is even more interesting to regard the behavior of SITH investors in a market dominated by chartist investors, where considerable price bubbles appear. We run simulations with the parameters, in the Table 6.26, that form a market with 70% chartist traders and 10% of each other trader type.

A market as described by the parameters from Table 6.26 is dominated by chartist investors and thus has a propensity for the formation of bubbles. We run simulations with these parameters and discover a recurring type of result: bubbles are formed often and occasionally the market price returns to its fundamental level (see Figure 6.48).

Test 12		
Parameter	Value	Explanation
μ_V	0.01%	Very small daily growth in the fundamental value
σ_V	1%	1% standard deviation
N	500	Number of trading days
Nr agents	500	10% GOD (perfectly informed) investors 10% SITH investors 80% Biased informed investors
a_p	[0.8, 2]	Biased informed investors have a parameter of information amplification randomly chosen for this interval
b_p	[-1, 1]	Biased investors will, on average, not skew their information
μ_C	2%	Profit expectations of SITH investors when they emulate a chartist strategy
σ_C	8%	
X	[30%, 60%]	SITH investors will change from a chartist to informed investor strategy when the market price is above/below the fundamental value with x% (the exact threshold for each investor is drawn uniformly from this interval)

Table 6.24: Test 12 - Simulation with rational investors, biased (dominant), not biased perfectly informed investors (GOD) and also with SITH investors.




Figure 6.44: Trading volume - test 12 simulation parameters in Table 6.22




Figure 6.45: Difference between fundamental and market value - test 12 simulation parameters in Table 6.22




Figure 6.46: Average wealth per investor type - test 12 simulation parameters in Table 6.22




Figure 6.47: Average portfolio holdings per investor type - test 12 simulation parameters in Table 6.22

Summary statistics for 50 simulations with biased investors	Sharpe Ratio	T-stat
Mean Sharpe ratio S_i for perfectly informed investor	0	
Mean Sharpe ratio S_b for biased investor	0,13051136	2,1095634
Mean Sharpe ratio S_s for SITH investors	0,39361916	6,46804634
Mean Sharpe ratio for a buy & hold strategy	0,204812	3,3465827
Average distance between fundamental and market price	-0,50562706	-216,207024

Table 6.25: Summary statistics for Test 12 with dominating biased investors

Test 13		
Parameter	Value	Explanation
μ_V	0.01%	Very small daily growth in the fundamental value
σ_V	1%	1% standard deviation
N	500	Number of trading days
Nr agents	500	10% GOD (perfectly informed) investors 10% SITH investors 10% Biased informed investors 70% Chartist investors
a_p	[0.8, 2]	Biased informed investors have a parameter of information amplification randomly chosen for this interval
b_p	[-1, 1]	Biased investors will, on average, not skew their information
μ_C (SITH)	2%	Profit expectations of SITH investors when they emulate a chartist strategy
σ_C (SITH)	8%	
X	[30%, 60%]	SITH investors will change from a chartist to informed investor strategy when the market price is above/below the fundamental value with x% (the exact threshold for each investor is drawn uniformly from this interval)
μ_C (chartist)	2%	Profit expectations of chartist investors
σ_C (chartist)	8%	

Table 6.26: Test 13 - simulation with rational investors, biased, not biased perfectly informed investors (GOD or Yoda investors), SITH investors and a dominant proportion of chartist investors

We observe the order of the best performing investor type: SITHs are better than biased informed investors which are better than perfectly informed investors. The investors that get to loose are of course those that do not coordinate on a value, the chartists. By decreasing the percentage of chartist traders or by skewing the bias of the informed traders we can have markets with more activity and less frequent price bubbles. The results we describe next, for the most parts, rest unchanged.

As expected from the composition of the market we observe in Figures 6.48 and 6.50 above that the market price departs frequently from the fundamental value. The mass of chartist traders manage to control prices, and consequently volatility, and shift them from their fundamental references.

From the evolution of average stock portfolio, in Figure 6.52, we observe that during periods of bubbles most stocks are held by chartist investors. Once a price bubble starts informed investors start letting go of the overpriced stocks in the following order: first perfectly informed investors, next the biased investors (especially those that overestimate the fundamental value) and last the SITH investors (depending on their $x\%$ strategy threshold limit). Actually, in the build-up phase of a bubble the SITH investors imitate chartists and act as catalyst and even amplifiers of price. Once the SITH investors reach their maximum risk limit (their r_{min} is between 30% to 60%) they sell their stocks to chartists and cash in the profit (see Figure 6.51). Because the selling action of biased and perfectly informed investors are not able to stop the bubbles the SITH's strategy is profitable. A SITH manages to ride, up to a risk limit, the bubbles generated (or maintained) by the chartist investors. Therefore SITH investors they manage to capture wealth from the chartist investors at a higher rate than the biased investors. These results are robust, as we can see in Table 6.27.

Research questions #5: *Research Question #5: Can Yoda beat Sith? How can live longer, Yoda or Sith?*

Answer: *Yes, SITH investors (informed rational players that can act also as uninformed) can beat YODA investors (informed players that act only to speculate on mispricings relative to fundamental values - like perfectly informed or biased informed investors) when market prices are persistently biased (or rest biased for sufficiently long periods). When mispricing happens, informed investors either take too much or to few assets and loose the opportunities available for "going" with the trend. Standard quantitative tests that reject the hypothesis that non-informed investors (like chartists or SITHs) can survive in markets ignore the possibility that some investors, like our SITH type, can know the model of the fundamental value (or at least know if market prices don't reflect fundamental values) and choose to act as a non-informed (speculate trends) only in special moments (when market mispricing exists but is not high enough to imply high risks).*

In the previous chapter, we have seen that persistently biased market prices are possible for long periods, even with simple market conditions. The answer to this




Figure 6.48: Evolution of fundamental value and market price - test 13 simulation parameters in Table 6.26




Figure 6.49: Evolution of average wealth per investor type - test 13 simulation parameters in Table 6.26




Figure 6.50: Evolution of excess volatility - test 13 simulation parameters in Table 6.26




Figure 6.51: Evolution of average cash for each investor type - test 13 simulation parameters in Table 6.26




Figure 6.52: Evolution of average stock portfolio for each investor type - test 13 simulation parameters in Table 6.26

Summary statistics for 50 simulations with biased investors	Sharpe Ratio	Tstat
Mean Sharpe ratio S_i for perfectly informed investor	0,97673179	10,231286
Mean Sharpe ratio S_b for biased investor	1,97200353	27,2560436
Mean Sharpe ratio S_s for SITH investor	3,12170565	53,1889911
Mean Sharpe ratio S_c for a chartist investor	0,10076797	5,48089118
Mean Sharpe ratio for a buy & hold investor	0,12966009	7,5581478

Table 6.27: Summary statistics for Test 13 simulations

last research question gives proof that informed investors can be better off if they choose to, sometimes, ignore their economic information and 'ride' market trends. We conclude by saying that informed investors have many incentives to:

1. Use their economic information to speculate **local mispricings** - only when the mass of non-informed (or badly informed) investors cannot dominate prices for sustained periods of time. This strategy is profitable in a (close-to) efficient market.
2. Ignore economic information and make profits with the trends, up to a certain accepted risk level (translated in a distance from "actual" fundamental levels) when non-informed or badly informed investors manage to dominate prices for long periods. This strategy is profitable in a far-from-efficiency market.

Such mixed investment strategy, purely motivated by profits, can transform financial markets in beauty contests: investors try to predict the predictions of others (irrelevant of basic economic facts). This emerging "non-efficient" behavior of markets can be stopped (or be amplified) because of multiple factors:

1. sufficient new well informed (or non-informed) investors are attracted by the market
2. non-informed (informed) investors see their profits diminish and decide to quit the market

A concept from modern finance research called 'majority game' can better model this situation of inefficient-efficient market. The concept of 'majority game' describes a situation where players have to make a choice between two options and the winners are those players that make the majority choice. In the case of a financial market, viewed as a game, there are two optimal choices for investors:

- a) All choose to buy information and act on it (we have an efficient market where all investors make the market return)
- b) Everybody ignores the information and they play a beauty contest (we have an inefficient market). Risk is higher but so are promises of return.

We believe that real financial market moves from an efficient to a non-efficient state. In an efficient market, information is abundant, and investors have incentives to not buy it anymore. When more and more investors start trading, without information, prices can move persistently away from their fundamental levels.

In that case informed investors are better off speculating these trends using non-fundamental based methods (like SITH investors). When enough investors use non-informed based strategies prices are more and more volatile and the market can develop bubbles. Information start again to be valuable and when a sufficient mass of investors start using it they can capture wealth from pure speculative traders.

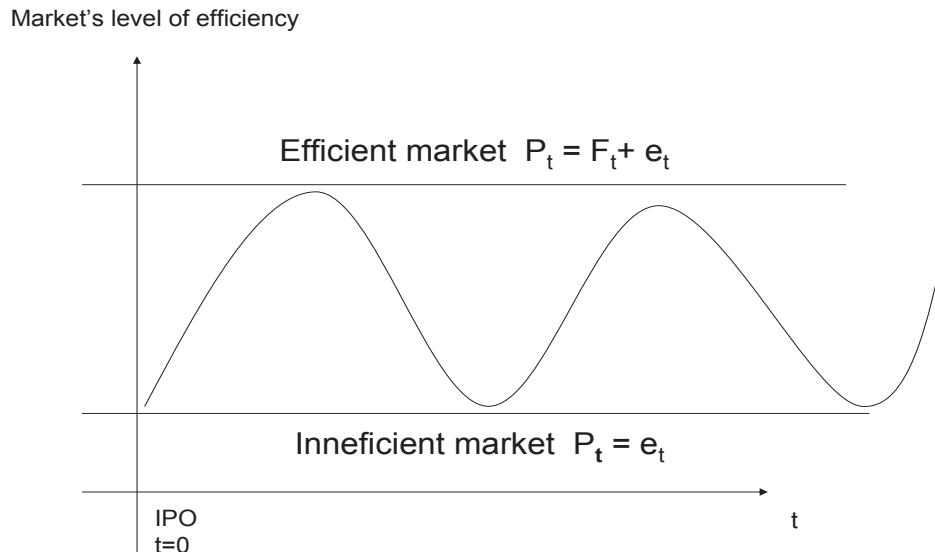


Figure 6.53: Theoretical transition of market from efficiency to inefficiency

The market can reach an efficient state and the cycle starts all over again. Therefore we propose a new definition relating the concept of efficiency with financial markets:

Sometimes efficient markets theorem: *Financial markets are "sometimes efficient". Markets exhibit a cyclical behavior: from an efficient state to an inefficient state. Most of the time market participants try to discover and adopt the dominant investment strategy: information or non-information based. Depending on the relative proportion of investment strategies the market will be move towards an efficient or an inefficient state.*

Corrolary: *Financial markets can be considered to be efficient at a strategy level. We define a **strategy-strong efficient market** as a market where above-market risk-adjusted returns cannot be achieved provided all investment strategy typologies⁷ are known. The strategy-level efficiency of a market implies that no new investement strategy can win superior returns. This does not implies that the market is also informationally efficient. A market can be informationally inefficient but no strategy can be used to profit off of the market inefficiencies.*

The dynamics of the transit between these two instable market states is complex and this complexity can be captured, in part, using simulation technology like our LUMA tool. The existence of these two market states (with different volatility regimes) can also account the market stylized facts.

⁷we refer to knowing the distribution of strategies and not necessarily which strategy (with specific parameters) is used by which investor.

In the next chapter we present a stylized game that shows, in a simple manner, how agents, faced with decisions under risk, can and should switch from objective fundamental based strategy towards mixed strategies that take into account the behavior of other participants.

The shark game: equilibrium with bounded rationality

Contents

7.1	Introduction	173
7.2	The model	175
7.3	Results	179
7.4	Conclusions	183

Abstract: *We present a toy model of a financial market. Hungry sharks represent investors and rewards for eating or not meat represent return on investments. Sharks adapt and learn what are the best actions using a human learning method, reinforcement-learning. Different payoff distributions can lead shark to find optimal strategies. We investigate the cases where equilibrium exists but it cannot be reached. We postulate that sharks, even without knowing, can engage in second-order dynamics. Their behavior mimics that of investors who, in their search for more profit, can ignore economic information and engage in "beauty contests" thus rendering markets inefficient. We manage to replicate this behavior with model-free learning and without any endogenous incentives.*

7.1 Introduction

Recent financial events, and their economic implications, have prompted a great number of actions starting from loss sharing to new policies for better governance of financial markets. From an abstract perspective, we say that, in the last decade, financial markets have either grossly mispriced asset risk (CDO risk mispricing as described in (Hamerle et al., 2009)) or they have badly priced asset future cash-flows and prices (as during the IT bubble (Morris and Alam, 2012)). Reasons for these mispricings are numerous and academics tend to revolve around the overwhelming evidence from behavioral sciences, see (Kahneman et al., 1981) or (Charness and Levin, 2003), namely that humans are boundedly rational.

In this context we propose a simple game¹, that mimics boundedly rational

¹This is an extended version of an article, with the same title, presented at the International

investors faced with decisions under risk. We will test if and when a rational expected equilibrium emerges. Besides its theoretical implications, our model is useful as a didactic and research tool for decisions under risk.

We investigate the emergence of equilibrium in a idealized world where investors are replaced by sharks and assets are replaced by meat. Sharks are hungry and they find a chunk of meat. All of the sharks are eager to eat but, using their excellent sense of smell, they detect that the meat has a toxic piece that could hurt or even kill them. The sharks know that a toxic piece is present but they do not know where that piece is. Therefore, sharks are faced with a risky decision: do not eat (and starve) or eat. If a shark decides to eat he has also to consider how many bites to take and for how long, so he can eat a maximum amount of meat and avoid the toxic part. In this setting, we investigate if reinforcement-learning sharks can manage to discover an equilibrium solution where a maximum number of pieces of meat are eaten and the toxic piece is avoided.

Toy models of financial markets have been created and used to observe the survival of boundedly rational agents in an competitive environment. One of the earliest models was the "El Farol bar problem" (Arthur, 1991). The citizens of El Farol like to visit the local bar. But if the bar is too crowded (more than 60% attendance) everyone has a bad experience. Each week, agents (who model El Farol's inhabitants) try to forecast what the majority will do. If an agent forecasts that the majority will go/not go the bar then he will adopt the opposite choice. Using mixed-strategies, the agents manage to converge the mean attendance of the bar to its maximum capacity. Yet the attendance level is always fluctuating around this mean and the agent strategies are always changing.

In (Challet et al., 2001) the authors coin this type of toy-model as a 'minority game'. As the name implies the best strategy is to choose the option that the minority chooses. All the gains of the minority agents are split equally. Under rational expectations such a game has no equilibrium and is a negative-sum game. Because of these features the goal of capturing financial market dynamics was not achieved.

This model of financial markets has evolved in a more realistic representation called a grand canonical minority game ((Galla et al., 2006) makes an excellent review). In this context, agents can choose, at any moment in time, not to make a choice (thus approaching more to the behavior of an investor that does not enter a market). This option generates fluctuations in the volume of the agents' actions (like trading volume) which in turn leads to statistical proprieties similar to stylized facts (?). Further research proposed models where agents had wealth that evolved in time. This ability created a weight effect where a agent's decision influenced the aggregate outcome as a function of his capital.

In (D.Sornette and Andersen, 2003) the authors propose 'the dollar game', a

Conference "Artificial Economics" 2012.

toy-model with features close to those of financial markets. The proposed payoff function forces players to learn and adapt to two-period-ahead strategies. Agents can choose to be in a minority (when they sell in anticipation of a price reversal) or in a majority (when they buy because they expect everybody else will buy). This framework allows for richer learning opportunities and generates dynamics that better reproduce those of real financial markets.

Minority-type games have also their disadvantages. The assumption of pure speculative markets makes superfluous the existence of exogenous information. This features imply that there is no equilibrium state because agents try, by design, to outforecast each other. The lack of an exogenous objective (like a fundamental value where the market should converge) does not allow for a measure of correctness or rationality. As the authors of (Challet et al., 2001) mention 'A rational approach is helpless'. Even if financial markets sometimes exhibit pure speculative behavior (meaning that market prices are disconnected from their fundamental values) it is not productive to create and work with models of financial markets where pure speculation is embedded by design.

The 'Shark Game' we propose adds to the existing academic literature by proposing a model that avoids the problems mentioned before. In our model a rational expectations equilibrium exists. Such an equilibrium implies two conditions: the aggregated outcome (like price) converges to the rationally expected solution (when price reflects only the fundamental value) and this state is also a Nash equilibrium (Nash, Jr., 1950) (no agent can benefit more by changing his strategy). Yet, real financial markets do not provide investors with the two necessary ingredients for achieving equilibrium: perfect information and a complete description of the underlying asset model. In our model, sharks encounter similar conditions as real investors: lack of complete information and lack of knowledge about the game's model. We observe how and when do the modeled sharks, using reinforcement-learning, arrive at an equilibrium. The simplicity of the game and its intuitive nature allows our model to be used in different contexts and problems (like different markets types or cooperation models).

7.2 The model

Consider a small sea golf where N sharks search for food. They discover a good piece of meat that unfortunately has a toxic part. The meat has M portions and a shark can eat a portion at a time. Sharks cannot distinguish the pieces of meat and therefore are not able to detect, in advance, which portion is toxic. For simplicity, consider the piece of meat as linear and that sharks can eat portions of meat starting from the first one.

A shark's goal is to accumulate the most points possible. A shark can receive A points for eating a good piece of meat. If a shark decides not to eat a piece of meat




Figure 7.1: Sharks and a piece of meat with 6 portions (where portion 4 is toxic)

he will get B points. However, if a shark eats the poisonous portion he will receive Z points. If more than one shark decide to eat the same portion of meat then it is not simple to decide how to split the reward. Therefore, for each new portion of meat available, sharks will be asked randomly for their decisions. If a shark, when asked, decides to eat then he receives the reward for eating (A or Z) and a new round begins with the remaining portions of meat. However, if a questioned shark decides not to eat, he will receive B points and another shark will be asked for his decision until either all sharks decide not to eat or at least one eats.

To understand better how the game develops we explain a scenario using the example from Figure 7.1:

1. Round 1: Shark 3 is sampled. He decides not to eat. Shark 2 is sampled and he decides to eat. The round is finished.
2. Round 2: Shark 1 is sampled. He decides eat. The round is finished
3. Round 3: Shark 4 is sampled. He decides not eat. Shark 1 is sampled and he eats. The round is finished
4. Round 3: Shark 2 is sampled. He decides to eat. The round is finished. Because the toxic piece has been eaten the game ends.

Following the payoff history in Table 7.1, we observe that in Round 1 only the 3rd shark is awarded the B reward (for not eating). Shark 1 and 4 do not receive any points because they did not get a chance (shark 2 was picked and decided to eat). This payoff scheme resembles that of investors who do not trade very often or are not able to trade due to market liquidity. A game stops, as it did in Round 4, when

Table 7.1: Shark reward structure for every round

Rewards	Shark 1	Shark 2	Shark 3	Shark 4
Round 1	0	A	B	0
Round 2	A	A	B	0
Round 3	A+A	A	B	B
Round 4	A+A	A+Z	B	B

the toxic part is eaten (since afterwards points are, on average, distributed equally) or when all the sharks decide not to eat. If a game reaches a round where nobody eats anymore then all sharks receive payoff B and the game ends. When a game ends the sharks receive their payoffs. These payoffs accumulate during multiple games.

We have designed and build a computer simulated environment where programs simulate the decisions of sharks. Sharks, like investors do, use simple thumb-rules to make their decisions. Each shark k has an array of rules, each formed by two parameters $Rule_k(H, S)$. H represents the maximum number of portions of meat the shark can eat. S represents the maximum number of pieces of meat available when the shark should stop eating. During a game, a shark's only information is the number of portions of meat left to eat. A shark using $Rule(3, 2)$ implies he will eat a maximum of 3 portions of meat and will stop eating when there are at most 2 pieces of meat still available.

A simulation lasts for K number of games (usually $K > 100$). At the beginning of each simulation all sharks start with all the $(M + 1)^2$ possible rules: from $(0,0), (0,1)$ to $(M,M)^2$. The sharks assign probabilities to each rule. At the start of each simulation, all sharks' rules have the same probability

$$P(Rule(x, y)) = \frac{1}{(M + 1)^2}. \quad (7.1)$$

After each game sharks receive their payoffs and update the probability of their rules. For example, shark k has used $Rule_k(3, 2)$ and received a payoff of G_k . The shark will add the payoff G_k to the rule probability and normalize the probabilities of all of his rules.

$$P(Rule_k(3, 2)) = P(Rule_k(3, 2)) + G_k \quad (7.2)$$

$$S = \sum_{x,y} P(Rule_k(x, y)), \forall x, y \quad (7.3)$$

$$P(Rule_k(x, y)) = \frac{P(Rule_k(x, y))}{S}, \forall x, y \quad (7.4)$$

At the start of a new game all sharks choose a rule that they use. This rule is randomly drawn using the shark's distribution of rule probabilities. The more

² M is the total number of portions of meat available in each game

payoff a rule wins the more the probability of the rule will increase and the consequently the more often the shark will use that rule. Because probabilities cannot be 1 or 0 all the rules are taken into consideration. Therefore, using this type of reinforcement-learning, sharks are able to adapt to extreme situation (like sudden changes in the payoff structures). When making a decision, a shark has a single piece of information: the number of pieces of meat left. The sharks, like investors, do not even know the model of the game. What they observe is the result of their actions at the end of the game. It is worth stressing that sharks cannot distinguish between toxic and toxic-parts. For all they know, the meat could have more than one piece of toxic meat or the rewards could change every game (or even every round). Sharks can only observe the outcome of their actions (through the end-of-game score G). The method of learning they use is reinforcement-learning, and it is model-free. The goal of this work is to see if competition for more profit, like it real financial markets, drives out inferior strategies and makes the best us out of the available resources (all good meat is eaten).

The intuitive solution to our model is to compute the expected value of eating the next portion of meat. This expected value E of eating, calculated if a shark has the chance to choose to eat, uses the payoff structure (A good meat, Z toxic meat and B for not eating) and the number (L) of remaining portions of meat. E is computed as the sum of expected value of eating the toxic piece and the expected value of eating a good portion less the payoff for not eating.

$$E = \frac{Z}{L} + \frac{((L-1) * A)}{L} - B \quad (7.5)$$

If the expected value of eating ³ (E) is positive then the shark should eat the current portion of meat, otherwise he should choose no to eat. The E value can be computed only if we know the payoff structure (A, B, Z). As in real financial markets, these payoffs are unknown to sharks. A point can be made that the payoff structure (the model) can be learned with a Bayesian-type method (Payzan-LeNestour, 2012). Such an approach would be fast but it would require a formal description of the model. When the game model changes (e.g. financial innovation) the sharks have to learn a new description and cannot adapt by themselves.

We choose reinforcement-learning because it mimics better the conditions of real financial markets. Investors know the outcome of their actions (returns) and also know their investment horizons (the portions of meat left). Moreover, investors know that an asset can turn out to be inflated and can loose value fast. Therefore the payoff Z for the toxic meat can equate to a risk premium investors place on risky asset. Now we will see if, after sufficient repetitions, sharks are able to reach an equilibrium (eating the most meat without the toxic part).

³In this case we assume that there is only a single piece of toxic meat. Generalisation is possible

7.3 Results

In our first tests we use simple payoff structures with clear equilibrium solutions and best strategies. Every simulation is preceded by a table, like the one below, containing the parameters of the test.

This first test sets the last portion of meat as the toxic one. Therefore there is an obvious solution that all sharks should adopt: eat the maximum portions of meat possible and stop eating if there is one piece left. The optimal rule, that sharks are best off choosing, is rule (3,1). In this case, the majority choice is the best and it assures a Nash equilibrium which coincides with all resources being used.

Using our computer-based simulator we run this scenario and observe the behavior of sharks. As depicted in Figure 7.2, during the first few hundred rounds, sharks do not avoid the toxic part (emphasized as the red line). Because the payoff of the toxic part is -10 the average cumulative payoff of sharks decreases. Every time a sharks bites the toxic part he penalizes the corresponding rule he used. As described in Figure 7.3, after some time the cumulative payoff of sharks rises and continues to rise.

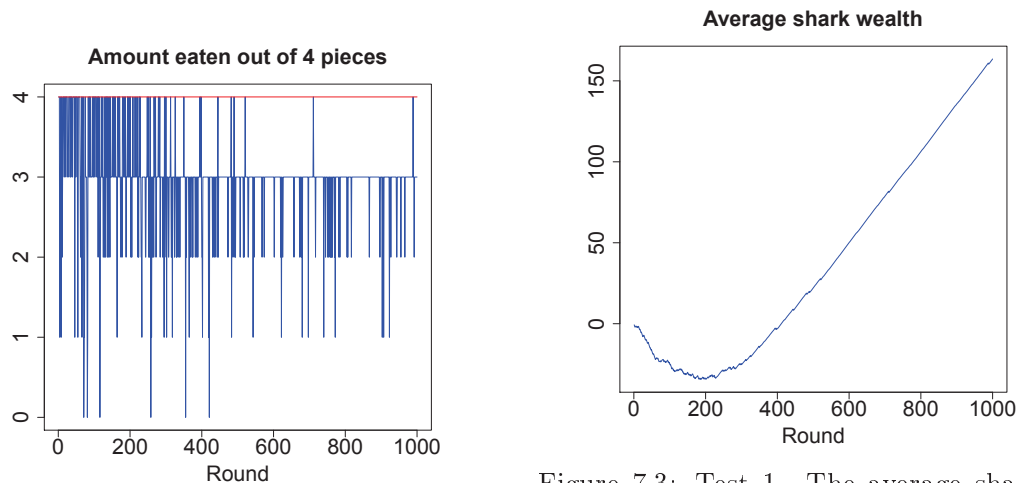


Figure 7.2: Test 1. Portions of meat eaten

Figure 7.3: Test 1. The average sharks' cumulative payoff

Eventually all sharks arrive at a distribution of probabilities where the probability of rule (3,1) approaches exponentially to 1 (and the all other probabilities approach zero). We point out that sharks learn to avoid rules that make them eat

Table 7.2: Test 1. Context and payoff structure

Nr. sharks	Nr. of meat portions	Position of toxic piece	Payoff A	B	Z
10	4	4	1	0	-10

the toxic part, like $Rule(\forall, 0)$. Moreover they also discard suboptimal rules, like $(1, \forall)$ or $(2, \forall)$, that even though make them avoid the toxic part are limiting their possibilities of wealth increase. In Figure 7.4 we observe the long-term results of Test 1. Convergence to the optimal solution (where the system stays in equilibrium) is evident.

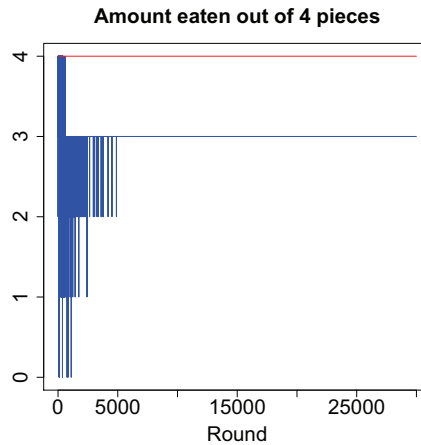


Figure 7.4: Test 1. Portions of meat eaten (30000 rounds)

The second test takes all the parameters of Test 1 and changes only the payoff for not-eating, from $B=0$ to $B=1$. This time, the choices of eating or not produce the same payoff. Intuitively we think that a strategy of eating everything (except the toxic piece) or not eating at all are equivalent. Therefore the optimal rules would be $(0, \forall)$ for never eating and $(3, 1)$ for eating everything. The choice between these rules is not evident and it creates an aggregate dynamic where all or none of the pieces of meat are eaten. We remind that a shark actually gets the not-eating payoff only when he has the chance to decide not to eat. Let us imagine a situation where 9 of the sharks choose to not eat, $Rule(0, \forall)$, and 1 shark chooses to eat everything up to the toxic part, $Rule(3, 1)$. Every round the greedy shark will eventually get his turn and will eat a portion of meat. Therefore, his payoff will be of 3 in every game (3 pieces of meat each with payoff $A=1$). The sharks that decide to abstain will receive the payoff $B=1$ only if they get picked, and refuse to eat, before the greedy shark gets his turn. Since the draws are uniform and there is only a shark that eats, a not-eating shark will have, in every round, the chance to make his choice with probability of 50% (in all the other cases the eating shark will be picked before him and the round will be over). Therefore the expected value of the payoff received by a non-eating shark, in a game, will be of $3 * B * 50\% = 1.5$. Thus, a non-eating shark has the incentive to choose the eat-all rule. When a shark switches to the eat-all rule two things happen:

1. More sharks compete for every portion of meat. The expected value of the eat-all sharks decreases (but is still bigger than that the expected payoff of

the non-eating sharks).

2. Because two sharks can eat a portion of meat, the probability of a non-eating shark to get the payoff be will drop below 50%. The expected value of the non-eating sharks will also drop.

The more sharks switch to $Rule(3,1)$ the more competition there is for eating the meat. The more competition there is the less is the expected payoff. Eventually the initial situation reverses and 9 sharks will use $Rule(3,1)$ and only one shark will use the not-eat rule $Rule(0,\forall)$. When this happens the expected value for an eat-all shark will be so low that they will again start using no-eat rules. We observe in Figures 7.5 and 7.6 that even after a considerable number of rounds sharks still choose, as a group, to eat nothing.




Figure 7.5: Test 2. Portions of meat eaten (5000 rounds)




Figure 7.6: Test 2. Portions of meat eaten (30000 rounds)

This payoff structure creates an environment with a normative equilibrium (all good meat parts are eaten and the toxic part avoided) but it also offers the possibility for sharks to depart from this equilibrium (in their search for more profit). The aggregate outcome fluctuates between eating all the meat and not eating at all. In their search for a maximum payoff it is possible that sharks leave untouched possible profits. When they all choose not to eat they each get B points and the untouched meat is worth $3*B$. This situation shows how a locally optimal solution (the sharks decision not to eat) produces an aggregate suboptimal solution (resources are not used, all pieces of good meat are untouched).

The game cycles continuously between the two extreme situations: all sharks try to eat everything and all sharks do not eat. Their behavior is rational and profit driven. We equate this behavior to that of investors who migrate from asset to asset, market to market or strategy to strategy in search for higher profits.

We now extend Test 2 and increase the payoff for not eating to $B=2$. The payoff

for not-eating is now higher than the payoff for eating. Like in the previous test, if all the sharks decide not to eat they can all get $B=2$ points. Yet if one shark switches to eating he will receive a payoff of 3 (because he will eat by himself all 3 portions of meat). The market will cycle between eating and not eating, between global optimal and suboptimal states. This behavior is rational and it is due to the fact that sharks base their decision on expected payoff. Expected payoffs depend on the strategies chosen by other players. Therefore, a shark's best strategy is ipso factum a function of the strategy of other sharks.

With financial markets in mind, we view the toxic piece as a price bubble. More clearly, a shark eating the toxic piece is similar to the 'biggest fool' investor, who buys an asset exactly when it's price is at a peak. As we observed, sharks are able to quickly learn the position of a fixed toxic meat portion. Therefore we run a simulation, with parameters from Table 7.3, where the toxic portion is placed in a random position at the beginning of each new game.

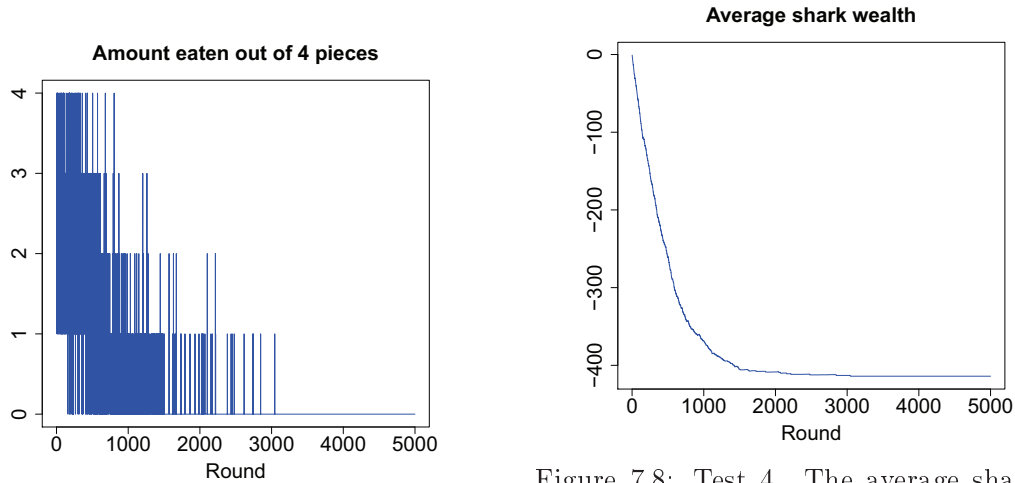


Figure 7.7: Test 4. Portions of meat eaten

Figure 7.8: Test 4. The average sharks' cumulative payoff

Because the toxic portion of meat can be anywhere sharks rapidly discount their expected payoffs. They first learn to avoid the 4th piece, than the 3rd until they resolve to not eating at all. We observe, in Figures 7.7 and 7.8, that after 3000 rounds the average cumulative payoff stabilizes and all sharks choose not to eat. In this situation sharks reach a stable equilibrium.

We now increase the payoff A to 5 and eating becomes very profitable. Sharks soon learn that eating the first piece of meat has a higher expected value than eating the second piece, and so on. As observed in Figure 7.9, sharks stabilize the

Table 7.3: Test 4. Context and payoff structure

Nr. sharks	Nr. of meat portions	Position of toxic piece	Payoff A	B	Z
10	4	Random between 1 and 4	1	0	-10

probability distributions of the rules they use and will choose each rule proportionate to it's expected value. In decreasing order of probability the rules sharks use are:

1. Eat nothing
2. Eat 1 piece (until 3 are left). This rule means eat only the first portion of meat.
3. Eat the first 2 portions
4. Eat the first 3 portions

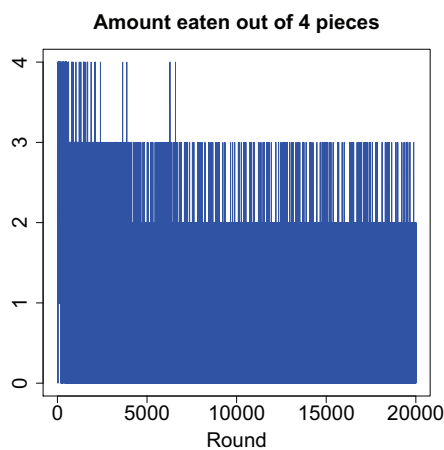


Figure 7.9: Test 5. Portions of meat eaten (30000 rounds)

If we tweak the model even more, using stochastic payoff values or multiple toxic parts, we obtain similar qualitative results. Sharks will update their probabilities, on the long run, as a function of the mean payoff structure. Using 'The Shark Game' we show that a group of agents, who do not understand models (they do not know the structure of payoffs, the number of toxic parts or that a toxic part even exists), and use reinforcement-learning, manage to learn profitable behaviors. Moreover we point out that individual profit-driven solutions do not always generate optimal global outcomes.

7.4 Conclusions

Using a toy model of a financial market we have showed how bounded rational sharks can learn, without ever knowing the game model, to discover optimal strategies. This is an argument for the possibility of market efficiency even if none of the investors knows (or are able to learn) the fundamental models of the assets.

The main driver of the sharks' actions, maximum payoff, will induce them to choose actions that are only locally optimal (on a particular round the shark gets a good payoff). When such actions are adopted by a large number of sharks, by chance or imitation, the aggregate result is suboptimal (meat, which could have brought more payoff, is left untouched).

Closer to financial markets, this phenomena was partially explained by Grossman and Stiglitz (Grossman and Stiglitz, 1980). When markets are close to efficiency, information-based trading becomes less profitable and informed investors leave the market (their local optimal solution). The more inefficient a market becomes (sub-optimal aggregate outcome) the more attractive it is for informed investors. In the light of this proof, financial markets can be viewed as minority games where the rational choices to be made are: trade actively or not trade at all. When active trading is abundant two situations can occur:

1. Prices integrate information better and faster (efficiency). The market becomes more efficient and the incentives for active trading decrease. Index (passive) trading becomes more profitable (which leads again to inefficiency). This is the paradox of informationally efficient markets.
2. Prices are disconnected from their fundamental values and incentives for active trading increases even more. This situation appears only in special circumstances⁴ and can lead to price bubbles. During these periods (of high market inefficiency) there is no objective best strategy. Going with the trend or against it could yield similar levels of returns.

Therefore, an investor can recognize that markets are in a continuous transition between an efficient state and an inefficient state. Since an investor's goal is to maximize profits he should adopt a strategy as a function of the current market state. In an inefficient market it is, up to a point, best to go with the crowd (thus explaining positive-feedback price bubbles). In an efficient market it is best to: either not trade at all or choose a passive portfolio and profit from fundamental growth (if sufficient). This mixed investment strategy is more profitable than any mono-strategy approach (always active or always passive) since it capitalizes gains during all market states. The key issue a rational investor poses next is: when should I switch strategies, from a technical strategy to value investment? The answer is: it depends on the choice of other investors. This rational realization is the seed for 'beauty-contest' behaviors. We observed the same behavior in our model, where sharks unwillingly play beauty contests (see Test 2). This rational behavior, motivated by profit maximization, generates locally optimal solutions for investors (maximum profits) but it does not assure that markets fulfill their normative role of allocating capital for economic growth.

⁴Technological innovations and cheap credit can create hypes that attract active but uninformed investors.

Conclusions

Contents

8.1 Our contribution	185
8.2 Main findings	187

Conclusion in one phrase: *Our results validate the idea that a financial market is a beauty-contest - the winner is the one that manages to predict the market's opinion and not the correct objective fundamental value - and so we put into question all asset pricing methods (along with their consequences) that do not rely on "objectively-irrelevant" information like peer opinion or strategy.*

8.1 Our contribution

We review here our models (or model improvements) or tools that contribute to the existence research body. The theoretical approaches we envisioned and tested are meant to respect the spirit of parsimony: explain the most things possible with the least amount of details.

Most information modeling processes, in regards to financial markets, are based on dividend (or dividend proxy) information. This requires investors to discount this information using discounting rates which are subject to debate. We proposed an informational model, see equation 5.2 from chapter 5, which is based on the difference between the last market price and the new fundamental value. Such information consists of a price differential which is an indication of discounted future cash flows. Therefore investors have access to a information stream that provides information, in a single series, of both cash flows and discount rates. As we shall see (in the next paragraph) this does not imply that all investors actually make future returns estimations using homogeneous discount rates.

Investors are known to generally have a risk-averse behavior. In most current models of financial markets, investors have to evaluate the riskiness of a stock. Using traditional asset pricing models, like the CAPM of (Sharpe, 1964), investors need to estimate multiple variables: risk-free rate, risk level of the asset and the market return. In our model we use a single investor specific measure, called minimum expected return r_{min} . An investor's minimum expected return, relative to an asset,

combines his risk-aversion, his risk estimation of that asset (or sensitivity and market expected returns) and also his reference risk-free rate. Analytically we can express a minimum expected return of investor X as:

$$R_{min}(X) = R_f(X) + RiskPremia(X) * EstimationOfAssetRisk(X) \quad (8.1)$$

By using a directly aggregated measure, instead of two or three different investor specific variables, we maintain investor heterogeneity while avoiding debates about risk premia or various risk estimations.

In our work we have presented an array of investor strategy models (see chapter 5). These models described in a parsimonious way (using only two parameters) a wide range of behaviors from optimists, conservators to pessimists and other variations. Moreover, our model was design and specifically checked as to not contain any mathematical methods, like use of historical volatility, that can implicitly generate "stylized facts" in market returns.

We have presented two new models of investment strategy. The first, called SITH, can arbitrage assets or choose go with the trend depending on his information. Different from other existing switching strategy models, a SITH does not arbitrarily change strategies but rather tries to maximize his profits depending on the variable distribution of market strategies. The second model we propose, called GOD, represents a simple benchmark for information value and semi-strong efficiency. A GOD investor is an arbitrageur that has perfect access to public information. By observing his portfolio performance, relative to other less informed investors, we can start drawing conclusions about the semi-strong efficiency of the studied market.

Our results, see chapter 6, show that GOD investors, as a group, are not always the best performers (hence markets are not efficient). We observed that SITH investors (which are actually well informed investors with a more complex strategy) can make better risk-adjusted returns. SITH investors have this advantage because they "observe" market efficiency (via a proxy distance between market prices and intrinsic value) and adjust their strategy accordingly. Therefore, we can say a SITH investor "arbitrates" investment strategies based on perceived market efficiency levels. From these observations we propose a new risk measure for assets which should incorporate not only fundamental (economic) related risks but also efficiency-level risk (which should not be confused with market risk).

We define efficiency-level risk as the risk faced by an investor, when buying/selling an asset, that the asset's price is highly under/over priced because of market inefficiency. Moreover, because (as we saw with the SITH investors) above market returns can be generated, at similar information level, by varying strategies we propose a new form of market efficiency: **strategy-strong efficiency**¹.

¹It is arguable that strategy-strong efficiency can be a reformulation of strong-form informational efficiency which can be tested easily.

This new research perspective implies that a market can be called strategy-strong efficient, provided above-market risk-adjusted returns are not possible even if the distribution of investor strategies is known. This type of efficiency implies that it is not profitable to develop profitable active investment strategies because we would be better off with a passive buy and hold strategy. It is worth mentioning that **a strategy-strong efficient market does not imply it is also informationally efficient**².

Our last theoretical contribution consists in a minority-type model, called The Shark Game (see chapter 7). This model provides a parsimonious toy-model of a financial market that can be used for didactic as well as for research purposes. The richness of the model consists in its few parameters (A,B and Z) which generate a wide range of financial settings.

Finally, we mention the design and creation of our financial market simulation tool as a practical contribution. The LUMA software tool, available online³ at [LUMA website](http://cerag.org/lucianstanciu/LUMA/index.html), is the first simulator of its type developed in a french doctoral school of management⁴ (and the second in France after the ATOM simulator, which we described in chapter 4).

8.2 Main findings

We summarize our research conclusions in a few simple propositions and afterwards we address them in more detail:

1. Financial markets can be efficient, even in the overwhelming presence of biases. Efficiency can be achieved if the average information bias (relative to the fundamental value) is null. This implies that biases are balanced and, on the aggregate, cancel each other out. When biases have a skewed distribution the market prices are biased (e.g. more investors are optimists or optimist investors control more wealth than other investors), meaning market prices diverge persistently from the rationally expected price.
2. Stylized facts are observed even when a market is efficient (market prices fluctuate, with a zero mean, around the fundamental value). A simple lag in the information used by investors can generate these statistical effects. The presence of "stylized facts" does not offer much information on the efficiency of a market or about which investor strategies are active in the markets.
3. When market prices are biased (e.g. persistent overpricing), informed investors have all the necessary incentives to sell their assets. Their joint selling action

²We redefine efficiency while respecting the (Grossman and Stiglitz, 1980) paradox on the impossibility of informationally efficient markets).

³<http://cerag.org/lucianstanciu/LUMA/index.html>

⁴Ecole Doctorale de Sciences de Gestion, EDSG 275, Grenoble

is not always sufficient to re-establish unbiased prices. If informed investors exhaust their resources then the market moves into a non-efficient state and the market prices are dominated by non-informed strategies. During these "non-informed" market periods we observe higher than usual volatility because the non-informed investors know they take higher risks and therefore they demand higher returns. Price bubbles usually occur and their amplitude and duration is limited by the wealth (and credit options) of the non-informed strategies. During such periods of market inefficiency, lowering costs of money can enhance non-informed speculation and can create bigger price bubbles.

4. When market prices are persistently biased it is profitable for informed investors to engage in non-informed (technical type) strategies (try to go with the "crowd"). The profitability of such pure speculative behavior, even when knowing that the assets are mispriced, depends on the speed at which market prices revert to fundamental values and also on the gap created between these values (maximum mispricing level). If markets rest biased for a longer time than a pure-speculator's investment horizon then it is profitable to engage in technical trading.
5. All else being equal, we say that a market with the following characteristics has a high chance of being inefficient and of generating price bubbles (or other anomalies):
 - (a) The market has an obscure asset (difficult to understand or lacking in information), that has showed promises of growth
 - (b) Money for investing is borrowed cheaply
 - (c) Markets, in general, are steady or growing. This condition attracts optimistic investors (and generates skewed a bias distribution which trigger more market mispricing).

We now go through a short review of the work presented. First, we detail the important factors that guide this research: the assumptions we make, the limitations of our methodology, our contribution to the scientific debate and specific results. In the end we offer a few hints about possibilities of extension to this work.

In the context of the recent financial crisis, regulators, researchers and the general public are questioning the role and safety of financial markets. In the academic finance field, the theory of behavioral finance has offered a scientific basis for analysing the market realities which are inconsistent with classical financial theory. The political, economic, social and academic context and motivations of our work are detailed in Chapter 1.

In Chapter 2 we go into more detail about our specific object of study: formation of prices in financial markets starting from the individual investor's behavior. In this part of our study we describe the flow of actions that transform economic information and profit motivations into investment orders which merge and form the

market price. Due to the complexity of such models of reality, classical mathematical language is not sufficient to create an easily comprehensible model. We choose to use a different language, namely computer language, to describe and model a generic financial market.

In Chapter 3 we present the relevant literature and offer our view about the academic works on the subjects of modelling, building and using computer-based market simulators for research in finance. We explain each part of a computer-based financial market model, from information to price formation. The chapter addresses the main critics that are faced by our methodology of research. Furthermore we show that a computer simulator is equivalent to a mathematical function - therefore all advantages and disadvantages of pure mathematical modelling also apply to computer language modelling.

Next, in Chapter 4, we presented the LUMA simulator, that was constructed specifically for our work. Our main contributions are found in this model. We proposed a unique type of economic information. Instead of revealing information directly about the economic value of an asset, which is very difficult and costly to do in real markets, we proposed a model where information about a financial asset is revealed in relation to the market price (too high, too low) - see for more details in chapter , page 80. Moreover, we contributed by modelling investor biases using two parameters which create an affine transformation of the perfect economic information. By changing this simple set of two parameters we obtain a wide range of investor behaviors, from pessimism to hindsight and optimism. We limit the scope of our study to extra-day trading periods. While this may be considered a though limitation on our study, we believe that intraday trading periods are not relevant to be studied in an efficient-behavioral context. We make this assumption considering that it is difficult to offer a solid economic model for the evolution of the value of a company in the duration a day (or a trading period).

Another limitation of our study is the fact that we model a single risky asset. We consider that a group of investors has a much harder time to gather, interpret and speculate on multiple assets than it does by focusing on a single stock. Therefore we implicitly assume that if a group of investors can (or cannot) create an efficient market with a single stock then they have (don't have) the possibility to create an efficient market with more than one stock. In regards to the common assumption of classical financial (and economic) theories, of homogeneity of investor strategy, we attest to the soundness of this assumption only because of the following simple fact: in a market overwhelmed with informed investors a non-informed investor cannot make superior profits. While this is true in theory it does not mean that real markets are all the time dominated by informed investors. In our tests we show that a market populated by mostly (in terms of financial weight) non-informed investors can create inefficient markets and also provide incentives for informed investors to not use economic information anymore. Moreover, we show that even informed investors have incentives to trade against their information, using technical strategies, in an

attempt to generate superior profits.

In Chapter 5 we put our market simulator to the test and provide answers to the five research questions. We start by performing benchmark tests that use clear input-output data sets aimed to validate the correct functionality of our research tool. Our tests show that a market with persistently biased prices can be created by a population of informed investors which have skewed biases (more optimists, more pessimists, etc.). Furthermore, "stylized" facts are showed to be generated by a number of causes, out of which the simplest is the fact the informed investors can act using outdated information (newspapers, old analyst reports, delayed news feeds, etc.). We underline that we contribute to the literature, concerning stylized facts, by proposing a realistic model of investor behavior that does not implicitly use autocorrelated functions (such as future price expectations as a function of past volatility).

After proving that persistently biased market prices are possible we focus on the survival of different investor strategies. By answering research questions 2 and 3 we show that biased informed investors can continue trading when market prices are themselves biased. Therefore these biased investors face weaker competition for the wealth of uninformed investors (noise or chartist traders) because well informed investors exhaust their resources when trying to re-establish correct price levels. Simply put, the investor strategy that dominates prices, for extended periods of time, can survive and prosper in financial markets.

From a philosophical point of view, our results can be viewed as a natural result of human adaptation. Investors tend to adopt the most profitable strategy even if it is a 'non-economic information based' strategy. Price bubbles do not imply a cost for the speculative investors that profited from the price rise. Such adaptive behavior can also be seen in political systems. Let us imagine that a democratic system is like an efficient financial market. When some citizens brake the laws of the democracy (steal, incite to violence, etc.) then the system will punish this behavior and restore order in the democracy. This is similar to when market prices depart shortly from their fundamental value. If enough informed investor react quickly than market prices can regain their economic values. Getting back to a democratic system, there are moments when the power in the state is seized by people with totally antidemocratic views (like in a coup d'état). When such events happen, citizens still have the same 'rationally-expected' behavior: try to punish the guilty and apply the laws of the democratic system. In extreme cases like in a totalitarian system, it is possible that the state of law is not respected anymore. Therefore, individuals that try to combat and protest against the new undemocratic regime can end up losing (imprisoned or even killed). In such situations, citizens have more than one choice: they can flee the country (to continue fighting the new regime and avoid imprisonment) or they can choose to support (or not openly disapprove) the new regime. History shows both choices are adopted by citizens and the results are as follows:

1. If the non-democratic system persists then people that have chosen to support it will prosper (some will never face trials for their behavior)
2. If the system topples quickly than the citizens that supported the new system will suffer the consequences.

We can imagine that a financial market where prices depart abruptly (and intensely) from economic fundamentals can resemble a non-democratic system that is coming to power. In such a situation, a rational informed investor is faced with two options (like a citizen in the new political system):

1. If the investor believes the market price will fall quickly to its fundamental value, than it is profitable to go against (short) the market.
2. If the investor believes the mispricing can persist (long enough) than it can be more profitable to go with the market (adopt a technical strategy). It is important for the investor to close his position before the prices readjust (and thus not face the downside of the market). Of course, after prices readjust (if ever) the investor can continue trading using a normal information based strategy.

In our results presentation, chapter 6, the formation and destruction of market price bubbles is also investigated. We show that price bubbles are created when a significant mass of investors, with positive future expectations of returns and lacking in economic information, can concentrate enough financial resources on a stock. Because it is clear that a market can lose its efficiency we have devised a new investor strategy (generically called SITH). A SITH investor intentionally switches between an informed and non-informed strategy. We showed how such a SITH investor can earn more than a pure informed investor in situation where the market loses efficiency for sustained periods of time.

Our research results tend to validate the 32 years old proof by Grossman and Stiglitz, see (Grossman and Stiglitz, 1980). This proof showed, in a mathematical language, that a market cannot be informational efficient. We believe a future research direction is to study an alternative risk measure for investment that will replace the classical volatility measure. This new investment risk measure will consist in **a function of the level of efficiency** of a certain traded asset. Therefore, instead of computing an expected return using volatility levels (or beta levels) we will use the level of efficiency of a market:

1. The more efficient a market is the less chances we have to earn an excess return (to the fundamental return of the asset) because prices are closely tied to the fundamental reality of the asset
2. The less efficient a market is the more possible it is to make excess returns (relative to the fundamental returns of an asset) because prices will vary depending on investor biases, mentality or crowd psychology.

These results add up to a unifying general theory about financial markets that we call '*sometimes efficient markets*'. Financial markets constantly move between an unstable efficient state and an large number of inefficient states (positive/negative bubbles, high/low volatility periods, etc). This perpetual transition is due to investors that observe diminishing returns on their strategies and choose to imitate (or innovate) the better-performing strategies of other investors⁵. When a market is efficient most of the investors use a fundamental based informed strategy and when a market is inefficient investors have a variety of competing strategies (and fundamental information is scarcely used). Therefore, when markets are efficient it can be more profitable to either exit that market or engage in non-fundamental strategies (like the case of our SITH investors). Vice-versa, in an inefficient market it may be profitable to adopt a fundamental strategy.

Measuring directly the level of 'efficiency' of a market is not simple because we do not have good models for the fundamental value of real assets. Instead, using the results from our simulation, we can envision methods for determining the relative strengths of different investor strategy types in a market (informed, not-informed, trend following etc.) in order to know if a market is approaching or distancing itself from an efficient state.

⁵A subfield in game theory, called minority games, offers good models of games where the minority strategies wins. We refer you to (Brandouy, 2005b), (Moro, 2004) or (Stanciu-Viziteu, 2013) for examples and reviews.

List of Figures

2.1	Global view of methodology	23
3.1	Global view of methodology	33
3.2	Estimation of correct information with central limit theory	36
3.3	Transformation of information into market prices	37
3.4	(black) Distribution of CAC40 returns, (green) Normal distribution	43
3.5	Autocorrelogram of absolute CAC40 returns	44
3.6	US residential real-estate price index (Source: S&P Case-Shiller Home Price Indices)	45
3.7	In 250 days time-periods: (above) Volatility of returns of the CAC40, (below) Signal to Noise Ratio of CAC40 returns	46
4.1	Representation of generic complex system	52
4.2	(A) Centralized (B) Scale-free (C) Closely distributed - Graphs with different topologies of connections between different entities of a system	57
4.3	Simulation of the flow of people evacuating a building	59
4.4	Figure from (Levy et al., 2000), chapter 7 page 162, showing the market generated price with 95% rational investors and 5% "efficient market believers".	66
5.1	Flow of information inside the LUMA simulator	80
5.2	Trading periods and information flows	82
5.3	Fundamental value and market price for an asset and the perfect information stream	84
5.4	Information interpretation for investment decision making	92
5.5	Typology of investment strategies	95
5.6	Typology of investment strategies (extended)	97
5.7	Speculators and the evolution of fundamental value (green) and market price (red)	100
6.1	Representative result - Market price (red) and fundamental value (black) for Test 1 - benchmark without investors	107

6.2	Representative result for Test 2: Fundamental (black) and market price (red) in a market with perfectly informed investors	108
6.3	Price time series (red) for a market with informed ($r_{min} > 0$) and noise traders (parameters in Table 6.4)	111
6.4	Wealth evolution for a market with informed ($r_{min} > 0$) and noise traders (parameters in Table 6.4)	112
6.5	Price time series (red) for a market only with chartist traders (parameters in Table 6.5)	114
6.6	Price time series (red) for a market only with chartist traders with highly volatile return expectations. Market price in red and fundamental value in black	114
6.7	Graph of the difference between fundamental value and market price (market parameters from Table 6.6)	116
6.8	Graph of the average portfolio Sharp value for each investor type (market parameters from Table 6.6)	116
6.9	Graph of the difference between fundamental value and market price with optimist biased investors)	118
6.10	Graph of the average portfolio value for each investor type with optimist biased investors	118
6.11	Dynamics of market price and fundamental value with biased investors	119
6.12	Excess volatility in a market where informed investors are optimists, a_p in $[0.9, 1.4]$	119
6.13	Example time series variation with optimistic biased investors (market price in red and fundamental value in black)	121
6.14	Test 7 - Example price series in a market where pessimistic investors represent 90% of the market	123
6.15	Test 7 - Evolution of wealth for different types of investors in a market where pessimist investors dominate (stocks are undervalued)	124
6.16	Distribution of returns for the fundamental value resulting from a simulation using the specifications in Table 6.12	126
6.17	Distribution of returns for the market price resulting from a simulation using the specifications in Table 6.12	127
6.18	Autocorrelation of absolute market returns resulting from a simulation using the specifications in Table 6.12 and a maximum lag of 8	128

6.19	Autocorrelation of absolute market returns resulting from a simulation using the specifications in Table 6.12 and a maximum lag of 20	128
6.20	Example market price behavior with 40% chartists, 60% perfectly informed (market price in red and fundamental value in black) - simulation parameters from Table 6.14	131
6.21	Another example market price behavior with 30% chartists, 70% perfectly informed (market price in red and fundamental value in black) - simulation parameters from Table 6.14	132
6.22	Example market price return series with market simulation parameters from Table 6.14	132
6.23	Probability distribution of market returns - simulation parameters from Table 6.14	134
6.24	Probability distribution of fundamental value returns - simulation parameters from Table 6.14	134
6.25	Test 9 Autocorrelation diagram for absolute market returns - simulation parameters from Table 6.14	135
6.26	Test 9 Autocorrelation diagram for normal market returns - simulation parameters from Table 6.14	135
6.27	Example simulation results in a market with 31% chartist (see table 6.17 for the market parameters)	142
6.28	Example simulation results in a market with 35% chartist (see table 6.17 for the market parameters)	142
6.29	Portfolio holdings of investors in simulation with a market with 35% chartist (see table 6.17 for the market parameters)	143
6.30	Excess volatility of market returns in simulation with a market with 35% chartist (see table 6.17 for the market parameters)	143
6.31	Trading volume in simulation with a market with 35% chartist (see table 6.17 for the market parameters)	144
6.32	Example simulation results in a market with 31% chartist (see table 6.17 for the market parameters)	146
6.33	Average cash holding portfolio for each trader type inside the simulation example from 6.32 - with parameters found in Table 6.17	146
6.34	Example simulation with bubbles and crash from a simulation with parameters like in Table 6.17, but changing $\mu_C = 8\%$ and memory length of 2 (instead of $\mu_C = 7\%$ and memory length of 3)	148
6.35	Presentation of abstract trading strategies	150

6.36	Market price movement - simulation parameters 10 a) in Table 6.19 .	151
6.37	Average wealth evolution per investor type - simulation parameters 10 a) in Table 6.19	151
6.38	Market price movement in a market where all investors have positive minimal required returns - simulation parameters 10 a) in Table 6.19 and $r_{min} > 0$	153
6.39	Evolution of ownership of stocks - simulation parameters 10 a) in Table 6.19 and all investors have positive minimal required returns, $r_{min} > 0$	153
6.40	Market and fundamental value - test 11 simulation parameters in Table 6.22	157
6.41	Average wealth per investor type - test 11 simulation parameters in Table 6.22	157
6.42	Average cash holdings per investor type - test 11 simulation parameters in Table 6.22	158
6.43	Average size of stock portfolio per investor group - test 11 simulation parameters in Table 6.22	158
6.44	Trading volume - test 12 simulation parameters in Table 6.22	162
6.45	Difference between fundamental and market value - test 12 simulation parameters in Table 6.22	163
6.46	Average wealth per investor type - test 12 simulation parameters in Table 6.22	163
6.47	Average portfolio holdings per investor type - test 12 simulation parameters in Table 6.22	164
6.48	Evolution of fundamental value and market price - test 13 simulation parameters in Table 6.26	167
6.49	Evolution of average wealth per investor type - test 13 simulation parameters in Table 6.26	167
6.50	Evolution of excess volatility - test 13 simulation parameters in Table 6.26	168
6.51	Evolution of average cash for each investor type - test 13 simulation parameters in Table 6.26	168
6.52	Evolution of average stock portfolio for each investor type - test 13 simulation parameters in Table 6.26	169
6.53	Theoretical transition of market from efficiency to inefficiency	171

7.1	Sharks and a piece of meat with 6 portions (where portion 4 is toxic)	176
7.2	Test 1. Portions of meat eaten	179
7.3	Test 1. The average sharks' cumulative payoff	179
7.4	Test 1. Portions of meat eaten (30000 rounds)	180
7.5	Test 2. Portions of meat eaten (5000 rounds)	181
7.6	Test 2. Portions of meat eaten (30000 rounds)	181
7.7	Test 4. Portions of meat eaten	182
7.8	Test 4. The average sharks' cumulative payoff	182
7.9	Test 5. Portions of meat eaten (30000 rounds)	183

List of Tables

5.1	The informational components of a limit order	85
5.2	Example limit-order book	88
5.3	Market price determination - computing cumulative quantities at all price levels	88
5.4	Step one of price determination - maximum volume criteria	89
5.5	Step two and three of price determination - minimum surplus and market pressure	90
5.6	Transaction resolution after price discovery	91
5.7	Bias parameters and associated behavior trait	93
6.1	Parameters for Test 1, first benchmark test (no agents)	106
6.2	Simulation parameters for Test 2 with perfectly informed investors that do not speculate	107
6.3	Simulation parameters for a test with perfectly informed investors that do not speculate	110
6.4	Parameters for a test with risk-averse perfectly informed investors and noise traders	111
6.5	Simulation parameters for a market having only technical traders	112
6.6	Scenario with informed investors, with balanced biases and without biases	115
6.7	Summary statistics for 100 trial simulation with input parameters from Test 4 with biased a_p parameters	117
6.8	Scenario with informed investors, with and without biases	121
6.9	Summary statistics for test 6 - shows that market price is consistently overestimating the fundamental value prices (in each test we take the average excess volatility of the average distance between fundamental and market prices and we average these measures across all the simulations	122
6.10	Test 7 - Market with dominant pessimistic informed investors	122
6.11	Summary statistics for test 7 - shows that market price is consistently underestimating the fundamental value	122
6.12	Market with perfectly informed investors, with lag in information	126

6.13	Statistical parameters for excess kurtosis and long memory of returns for 100 tests with the parameters from Table 6.12	127
6.14	Simulation parameters for Test 9 - volatility clustering with technical traders	130
6.15	Robustness test for simulation with parameters from Table 6.14	133
6.16	Hypothesis tests for the presence of price bubbles	137
6.17	Simulation parameters for Test 10: we vary the proportion X from 100% downwards	140
6.18	Summary statistics for results of multiple simulations on different levels of relative proportions of chartists to informed investors.	140
6.19	Simulation with rational investors, biased and not biased (which are dominant)	150
6.20	Sharpe ratios for test 10 a) (biased investors dominated) when investors are and investors have $r_{min} = 0\%$	154
6.21	Summary statistics for test 10 b) - perfectly informed investors dominate	154
6.22	Test 11 - Simulation of a market dominated by chartist traders and with perfectly informed traders	156
6.23	Summary statistics for multiple runs of Test 11	159
6.24	Test 12 - Simulation with rational investors, biased (dominant), not biased perfectly informed investors (GOD) and also with SITH investors.	162
6.25	Summary statistics for Test 12 with dominating biased investors	164
6.26	Test 13 - simulation with rational investors, biased, not biased perfectly informed investors (GOD or Yoda investors), SITH investors and a dominant proportion of chartist investors	165
6.27	Summary statistics for Test 13 simulations	169
7.1	Shark reward structure for every round	177
7.2	Test 1. Context and payoff structure	179
7.3	Test 4. Context and payoff structure	182

Appendix

A.1 Are financial markets ethically efficient?

In this article I will share with you my thoughts about ethics in finance. My aim is to show that, fortunately, the world of finance has some internal mechanisms for becoming more "ethical," just like a living organism has the capacity to evolve and adapt. To start with I would like to challenge you, the reader, to a mental exercise: "Imagine a world where all finance professionals are 100% ethical and highly qualified". What next? Will markets be safer? Will crises disappear? Let's explore the answers together. The subprime crisis affected the global financial system. Its effects have also been visible in many world economies..... Who is being blamed: the market, financial consultants, Madoff, Enron, global invaders and other common financial scapegoats?

Does this sound familiar? Every time something bad happens in financial markets, "investors" lose confidence in markets and they start blaming others (especially finance professionals). This natural defence mechanism then motivates regulatory bodies to start taking new actions like creating tighter regulations, more transparency, better ethical codes or other new measures to help protect the weak and exposed investor.

This line of thought is not entirely complete, because what we don't hear is that investors themselves are among the main causes of price bubbles and the ensuing crises. I strongly believe that even if all financial professionals were "100% ethical" and highly qualified, price bubbles and crashes would still happen and it would be entirely the fault of investors. Let me explain why this is true.

Imagine two intelligent portfolio managers, Pierre and John. They both follow a highly regarded code of ethics and have strong professional qualifications. In 2007, Pierre and John had the opportunity to invest in a new financial product called CDO (collateralized debt obligation). To make a good investing decision (both ethical and in accordance with their client's objectives) the two managers engaged in a due diligence process using the best of their knowledge and skills. John arrived at the conclusion that investing in CDOs could offer a good return and an acceptable level of risk for his clients (given his research and also the good ratings displayed by widely-accepted rating agencies). So John bought these products. Like John, Pierre tried to understand, as best he could, these financial products yet he could not arrive at an acceptable conclusion about their risk and return profiles. Being

cautious by nature, Pierre decided not to buy any CDOs and continued to seek other opportunities. What happened next is the most interesting part. The following year, in 2008, investors saw the results of their portfolios. Some of Pierre's clients were not very happy with the results and pointed out that John's portfolio had a much greater return (due mainly to CDO growth). Pierre, of course, defended his decision, pointing out that a CDO is very complex and it is difficult to properly assess the associated risk. Some investors believed Pierre, but others stopped doing business with him and moved all of their wealth to John's care.

In reality, in 2007 and 2008, many finance professionals like Pierre went out of business because they did not go with the CDO growth wave. Their clients left them for a "better" manager, one that would invest in a CDO and thus gain bigger "returns" with the same level of (perceived) risk. The end of the story, the subprime crisis and its effects, is history. Expanding on this recent event and the story of John and Pierre, I believe that even if all advisers and portfolio managers were cautious like Pierre and had refused to continue investing in CDOs, investors would have found independent ways to buy those high-return/low-risk securities, without the help of professionals. History shows us that we can replace the CDO security with anything else (petrol, Enron stock, tulips) because the same (il)logical behavior will apply. After the crash at the end of 2008, when CDOs and related securities lost most of their value, investors started looking for answers and explanations for their losses.

Who is to blame they asked? Yes of course, John is to blame because he invested in CDOs. But John worked to the best of his knowledge and followed all ethical guidelines. His fault was maybe that he chose to invest in a security he did not understand very well (yet people like John probably did not know that they did not properly understand CDO risks). So, rating agencies are probably to blame. They were misleading and lacked ethics. Following the crash, there were legislators who wanted to punish rating agencies and make them adopt tighter rules. They proposed a law that would forbid rating agencies to announce "bad news", at least during turbulent financial periods. I found (from an anonymous internet forum comment) a simple explanation that mirrors the logic behind this kind of proposition: "It's cold outside. Let's break the thermometer!" Rating agencies, by definition, operate in a state of conflict of interest. In 2007 and 2008 there were, however, companies, associations and professionals that sent warning messages about the dangers of CDO investing. Financial markets always send a mix of reassuring and warning messages and it is up to the investor to make the final choice. Unfortunately, before the subprime crisis, many investors chose to listen only to the good signals. This is not unlike people who smoke even though packs of cigarettes come with messages and pictures warning of the dangers of smoking. In short, the recent crisis, like many others, has a principle cause found in investor behavior and choices. These behaviours and choices are not backed up by a careful reasoning process and are often motivated by a blind desire for wealth (sometimes called greed). This irrational behavior and its consequences are not limited to the field of finance and the

behavior is enabled by all democratic systems that give people, regardless of their competencies, the privilege of making their own decisions.

With the story of John and Pierre we illustrated how market bubbles and crashes can appear even if all finance professionals are ethical. Therefore, I consider that ethical behavior is necessary for all finance professionals but is not sufficient to create crisis-free markets. Let's first look at what should be done to achieve a maximum level of ethics in finance and then we will point out what is needed to complement ethics.

At present we observe the existence of numerous codes of ethics adopted by companies, professional associations and even educational institutions. These codes are most often a very detailed list of things to do (good things) and things not to do (bad things) whilst exercising a finance related function. With a few exceptions (due to the complexity of financial dealings) most of the do's and don'ts in finance can, in my view, be summarized in one simple phrase: "Don't do things to others that you would not like to have done to you!" I believe that this simplification is possible because most ethical guidelines are perceived intuitively as being correct. Said otherwise, all professionals know that it is wrong to do things like give false interpretation of information or mislead clients into taking actions that go against their best interests. Nevertheless, there are professionals that bend this simple rule. Society's reaction to these misdeeds is simple and straightforward: some receive fines and penalties or lose their jobs or face justice. Yet, unethical behavior is a special kind of misdeed requiring a larger and more global reaction to limit the potential for future abuse. I believe that knowing the causes of these acts can greatly improve future ethical developments.

Why do finance professionals engage in unethical behavior? Leaving aside obvious answers, like "they are bad and greedy", I believe that the field of finance provides a special context for the application of ethics because of the intersection of two unique conditions. Ethical codes impose a moral way of thinking that is in contradiction with the for-profit mind set of all finance professionals. The second unique condition is the difficulty (and sometimes impossibility) of observing the causal link between one's actions and their consequences.

Ethics involves thinking and acting in an empathetic way. In practice, ethics involves letting go of some benefits (or opportunities) to help (or not hurt) others. In total opposition to this, we have the driving force of capitalism - the maximization of self-interest. Maximization of self-interest ignores the benefits of ethical behavior because there is no widely accepted value for "being good". This is probably why many people see finance and ethics as incompatible. Finance professionals are trained and are supposed to maximize their clients (utility of) wealth. As long as there is no price (and no market) for ethical behavior, this contradiction in objectives will be a restricting condition to the application of ethics in finance.

In addition to the conflict between ethical and financial objectives, I believe

that ethics are not easily applied in the field of finance because the causal link between one's unethical behavior and its consequences is obscure. For example, a doctor when faced with his patient is very aware of the consequences of performing unnecessary surgery or over prescribing drugs. Because the causal link between medically unethical behavior and its consequences (illness or even death) is clear, most doctors are motivated to closely adhere to widely held ethical codes. On the other hand, finance professionals often see the results of their actions only in the bottom line of their portfolios. It is not easy to see the link between one's investing actions and people losing their pension funds or their homes. In my opinion, stating the logical links between unethical behavior and loss of jobs, homes and wealth is infinitely more effective than any abstract moral rhetoric. These links exist and their global recognition is critical to developing ethical behavior.

We have seen conflict between objectives and the obscurity of causal links as two conditions that make it difficult for ethics to be embraced by finance professionals. Working towards eliminating these hindering conditions is necessary for further ethical advancement. The good thing about finance in general is that it has a subtle method for adopting ethical behavior: ethical efficiency. As with informational efficiency, I believe that unethical behavior ultimately has a positive impact on future ethical behavior. When inefficiency is discovered, it is exploited by market participants until it disappears. Similarly, since the discovery of the "Madoff ethical inefficiency", investors and finance professionals are double-checking their environment for similar "inefficiencies". We can quite safely say that Ponzi-like schemes have a very slim chance of working in the future (of course, provided that future generations will learn and understand how to avoid this "ethical inefficiency" and its consequences).

In this article we have pointed to some of the obstacles to ethics in finance. Moreover, we have explained that in the future an ethical professional body will not be sufficient to create crisis-free markets. It is important to continue refining and improving the financial environment so that ethics can flourish, especially by eliminating the above-mentioned restricting conditions. In my opinion, however, legislators and professionals should also invest more energy in educating all market participants present and future. This can be achieved by teaching children about ethics, risk and finances from the beginning of their school careers. It is through education that future investors, today's children, will be conscious of and understand risks and returns and will be able to make better informed decisions. Likewise, education will help professionals see the links between their actions and possible consequences. It is only through education that markets will eventually become ethically efficient.

A.2 Central Limit Theory for number estimations

To solve the problem describe in 3.1 we provide below an R code. We first generate the estimations of people trying to guess the number of balls. In this case the estimates are uniformly distributed but any distribution can be used. Afterwards, we take random samples of these estimates and compute the mean of each sample. Finally we draw the histogram of all the sample means and observe they follow a normal distribution with the good average (of exactly the number of balls in the jar). This code is executable in the R Statistical software from <http://www.r-project.org/>

```

N<-100 # real number of balls in the jar
popSize<-20000 # how many people try to guess the number
estimate<-NA

for (i in (1:(popSize/2))*2)
  estimate[i]<-N+runif(1,-10,0); # estimations above N
for (i in ((1:(popSize/2))*2)-1)
  estimate[i]<-N+runif(1,0,10); # estimations below N

mi<-NA
size = 1000;
nb_samples = popSize/4;
for (i in (1:popSize/4))
  mi[i]=mean(sample(estimate,size)); # averages of samples of estimations

hist(mi, nclass=100) # histogram of all sample averages

```

A.3 Signal to noise Ratio

The Signal to Noise Ratio of CAC40 indices was computed using the code below. The price quotes of CAC40 can be found on Yahoo Finance. We draw the attention that the quotes from Yahoo Finance start from the last quote, so they need to be inverted when computing returns. The code below reads price quotes (where the first quote in the file is the more recent quote) and then inverts them when computing returns. This code is executable in the R Statistical software from <http://www.r-project.org/>. A version of CAC40 price quotes can be found at <http://cerag.org/lucianstanciu/table.csv>

```

#compute returns from prices
A=read.table('table.csv',sep=',');

```

```
#inversion of price quotes.
#remove if first quote is the oldest (not the newest)
P=A[rev(1:length(A[,1])),1];

R=rep(0,length(P)-1);
for (i in 2:length(P))
  R[i-1]=log(P[i]/P[i-1]);

#calcul signal-to-noise and periodic volatility

window=250; # how many returns to get in a sample mean and volatility
C=rep(0,length(R)-window);
StN=rep(0,length(R)-window);
for (i in window+1:length(R)-1) {
  StN[i-window] = mean(R[i-window:i])/sd(R[i-window:i]);
  C[i-window] = sd(R[i-window:i]);
}
affichage_sans_premiere_obs=250;
plot(StN[affichage_sans_premiere_obs:length(StN)],...
type="l",main="CAC40 sd (periods window)");
```

Bibliography

- Abreu, Margarida, and Victor Mendes. "Information, Overconfidence and Trading: Do the Sources of Information Matter?" Working Papers 2011/25, Department of Economics at the School of Economics and Management (ISEG), Technical University of Lisbon., 2011. <http://ideas.repec.org/p/ise/isegwp/wp252011.html>. (Cited on page 34.)
- Alchian, Armen A. "Uncertainty, Evolution, and Economic Theory." The Journal of Political Economy 58, 3: (1950) 211–221. <http://www.jstor.org/stable/1827159>. (Cited on page 28.)
- Allen, Franklin, and Douglas Gale. "Asset Price Bubbles and Monetary Policy." Center for Financial Institutions Working Papers 01-26, Wharton School Center for Financial Institutions, University of Pennsylvania, 2000. <http://ideas.repec.org/p/wop/pennin/01-26.html>. (Cited on page 130.)
- Amat, Oriol, and Catherine Gowthorpe. "Creative accounting: Nature, incidence and ethical issues." Economics working papers, Department of Economics and Business, Universitat Pompeu Fabra, 2004. <http://EconPapers.repec.org/RePEc:upf:upfgen:749>. (Cited on page 37.)
- Anderson, P. W. "More Is Different." Science 177, 4047: (1972) 393–396. <http://www.sciencemag.org/content/177/4047/393.short>. (Cited on pages 51 and 52.)
- Arthur, W Brian. "Designing Economic Agents that Act Like Human Agents: A Behavioral Approach to Bounded Rationality." American Economic Review 81, 2: (1991) 353–59. <http://ideas.repec.org/a/aea/aecrev/v81y1991i2p353-59.html>. (Cited on page 174.)
- B., O. Brandouy, P. Beaufils, et al. "Une analyse de la complexite des dynamiques financieres a l'aide de modeles multi-agents." Qu'appelle-t-on aujourd'hui les sciences de la complexite 185–212. (Cited on page 59.)
- Bachelier, L. Theorie de la speculation. Gauthier-Villars, 1900. <http://books.google.fr/books?id=ANJUAAAAYAAJ>. (Cited on page 30.)
- Barberis, Nicholas, and Richard Thaler. "A survey of behavioral finance." In Handbook of the Economics of Finance, edited by G.M. Constantinides, M. Harris, and R. M. Stulz, Elsevier, 2003, volume 1 of Handbook of the Economics of Finance, chapter 18, 1053–1128. <http://ideas.repec.org/h/eee/finchp/2-18.html>. (Cited on page 30.)
- Barrevel, D.J. The Enron Collapse: Creative Accounting, Wrong Economics Or Criminal Acts? : a Look Into the Root Causes of the Largest Bankruptcy in

- U.S. History. Writers Club Press, 2002. <http://books.google.ca/books?id=uzDhZ5V1pZgC>. (Cited on page 37.)
- Barry, Christopher B., and Robert H. Jennings. "Information and Diversity of Analyst Opinion." Journal of Financial and Quantitative Analysis 27, 02: (1992) 169–183. http://EconPapers.repec.org/RePEc:cup:jfinqa:v:27:y:1992:i:02:p:169-183_00. (Cited on page 34.)
- Beatrice, Vladu Alina, and Adrian Grosanu. "SOME INSIGHTS REGARDING CREATIVE ACCOUNTING IN ROMANIAN ACCOUNTING ENVIRONMENT - REGULATORS, FINANCIAL AUDITORS AND PROFESSIONAL BODIES OPINION." Annals of Faculty of Economics 1, 1: (2011) 661–668. (Cited on page 37.)
- Beaufils, B., and O. Brandouy. "Artificial Economics: Agent-Based Methods in Finance, Game Theory & Their Applications." Lecture notes in Economics & Mathematical Systems 564: (2005) 02. (Cited on page 59.)
- Beaufils, B., O. Brandouy, L. Ma, and P. Mathieu. "Simuler pour comprendre: Une explication des dynamiques de marches financiers par les systemes multi-agents." Systemes d'Information et Management 14, 4: (2009) 1–26. (Cited on page 56.)
- Bebczuk, R.N. Asymmetric Information in Financial Markets: Introduction and Applications. Cambridge University Press, 2003. <http://books.google.fr/books?id=dEPt1Rz4z6EC>. (Cited on page 34.)
- Beltratti, A.E., and S. Margarita. Evolution of Trading Strategies Among Heterogenous Artificial Economic Agents. Nota di lavoro. 1992. <http://books.google.fr/books?id=bNQRtwAACAAJ>. (Cited on page 73.)
- Bernoulli, Daniel. "Exposition of a New Theory on the Measurement of Risk." Econometrica 22, 1: (1954) 23–36. <http://dx.doi.org/10.2307/1909829>. (Cited on page 47.)
- Black, Fischer. "Noise." The Journal of Finance 41, 3: (1986) pp. 529–543. <http://www.jstor.org/stable/2328481>. (Cited on page 110.)
- Blume, Larry, and David Easley. "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." Cowles Foundation Discussion Papers 1319, Cowles Foundation for Research in Economics, Yale University, 2001. <http://ideas.repec.org/p/cwl/cwldpp/1319.html>. (Cited on page 29.)
- Bo, Yang, Wu Yong-gang, and Wang Cheng. "A Multi-agent and GIS Based Simulation for Emergency Evacuation in Park and Public Square." In Computational Intelligence and Security, 2009. CIS '09. International Conference on. 2009, volume 1, 211–216. (Cited on page 58.)

- Bottazzi, Giulio, and Mikhail Anoufriev. "Price and Wealth Dynamics in an Agent-Based Model with Heterogeneous Evolving Strategies." *Computing in Economics and Finance 2004* 227, Society for Computational Economics, 2004. <http://EconPapers.repec.org/RePEc:sce:scecf4:227>. (Cited on page 59.)
- Bottazzi, Giulio, Giovanni Dosi, and Igor Rebesco. "Institutional architectures and behavioral ecologies in the dynamics of financial markets." *Journal of Mathematical Economics* 41, 1-2: (2005) 197–228. <http://ideas.repec.org/a/eee/mateco/v41y2005i1-2p197-228.html>. (Cited on page 73.)
- Brandouy, O. "Complexite et phenomenes critiques en finance." . (Cited on page 55.)
- . "Stock markets as Minority Games: cognitive heterogeneity and equilibrium emergence." *Physica A: Statistical Mechanics and its Applications* 349, 1: (2005b) 302–328. (Cited on pages 57 and 192.)
- Brandouy, O., P. Barneto, and L.A. Leger. "Insider trading, imitative behaviour and price formulation in a stimulated double-auction stock market." *Economics Working Research Paper* . (Cited on page 37.)
- Brandouy, O., A. Corelli, I. Veryzhenko, and R. Waldeck. "A re-examination of the "zero is enough" hypothesis in the emergence of financial stylized facts." *Journal of Economic Interaction and Coordination* 1–26. (Cited on page 44.)
- Brandouy, O., and P. Mathieu. "A Conceptual Framework for the Evaluation of Agent-Based Trading and Technical Analysis." *Artificial Markets Modeling* 63–79. (Cited on page 61.)
- Brandouy, O., P. Mathieu, and I. Veryzhenko. "Risk Aversion Impact on Investment Strategy Performance: A Multi Agent-Based Analysis." *Managing Market Complexity* 91–102. (Cited on page 71.)
- Brandouy, O., P. Mathieu, et al. "Large Scale investigation of EMH with virtual agents." . (Cited on page 59.)
- Brandouy, O., et al. "Les marches financiers artificiels." *Revue francaise de gestion* , 6: (2006b) 161–180. (Cited on page 56.)
- Brock, William A., and Cars H. Hommes. "Heterogeneous beliefs and routes to chaos in a simple asset pricing model." *Journal of Economic Dynamics and Control* 22, 8-9: (1998) 1235–1274. <http://ideas.repec.org/a/eee/dyncon/v22y1998i8-9p1235-1274.html>. (Cited on page 64.)
- Campbell, John Y. "Asset Pricing at the Millennium." *Journal of Finance* 55, 4: (2000) 1515–1567. <http://ideas.repec.org/a/bla/jfinan/v55y2000i4p1515-1567.html>. (Cited on page 60.)

- Challet, Damien, Matteo Marsili, and Yi-Cheng Zhang. “Minority games and stylized facts.” Physica A: Statistical Mechanics and its Applications 299, 1: (2001) 228–233. <http://ideas.repec.org/a/eee/phsmap/v299y2001i1p228-233.html>. (Cited on pages 174 and 175.)
- Chang, Eric C., Joseph W. Cheng, and Ajay Khorana. “An examination of herd behavior in equity markets: An international perspective.” Journal of Banking & Finance 24, 10: (2000) 1651–1679. <http://ideas.repec.org/a/eee/jbfina/v24y2000i10p1651-1679.html>. (Cited on pages 35 and 37.)
- Charness, Gary, and Dan Levin. “When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect.” University of California at Santa Barbara, Economics Working Paper Series qt7g63k28w, Department of Economics, UC Santa Barbara, 2003. <http://ideas.repec.org/p/cdl/ucsbec/qt7g63k28w.html>. (Cited on page 173.)
- Chiarella, Carl, and Giulia Iori. “A simulation analysis of the microstructure of double auction markets.” Quantitative Finance 2, 5: (2002) 346–353. <http://ideas.repec.org/a/taf/quantf/v2y2002i5p346-353.html>. (Cited on pages 24 and 74.)
- Cipriani, M., and A. Guarino. “Herd behavior in financial markets: an experiment with financial market professionals.” Open Access publications from University College London <http://discovery.ucl.ac.uk>, University College London, 2009. <http://ideas.repec.org/p/ner/ucllon/http--discovery.ucl.ac.uk-16552-.html>. (Cited on page 36.)
- Cont, Rama. “Empirical properties of asset returns: stylized facts and statistical issues.” Quantitative Finance 1: (2001) 223–236. (Cited on page 47.)
- Cootner, P. The Random Character of Stock Market Prices. MIT, Press, Cambridge, MA, 1964. (Cited on page 28.)
- Coval, Joshua D., and Tyler Shumway. “Do Behavioral Biases Affect Prices?” Journal of Finance 60, 1: (2005) 1–34. <http://ideas.repec.org/a/bla/jfinan/v60y2005i1p1-34.html>. (Cited on page 17.)
- De Bondt, Werner F. M., and Richard H. Thaler. “Further Evidence on Investor Overreaction and Stock Market Seasonality.” The Journal of Finance 42, 3: (1987) 557–581. <http://www.jstor.org/stable/2328371>. (Cited on page 17.)
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. “The Survival of Noise Traders in Financial Markets.” Scholarly Articles 3725470, Harvard University Department of Economics, 1991. <http://ideas.repec.org/p/hrv/faseco/3725470.html>. (Cited on page 28.)

- Derveeuw, J., B. Beaufils, O. Brandouy, P. Mathieu, et al. “L’apport des smaa la modelisation des marches financiers.” Revue d’Intelligence Artificielle 21, 5-6: (2007a) 617–641. (Cited on page 56.)
- Derveeuw, J., B. Beaufils, P. Mathieu, O. Brandouy, et al. “Un modele d’interaction realiste pour la simulation de marches financiers.” . (Cited on pages 58 and 63.)
- Ding, Zhuanxin, Clive W. J. Granger, and Robert F. Engle. “A long memory property of stock market returns and a new model.” Journal of Empirical Finance 1, 1: (1993) 83–106. <http://ideas.repec.org/a/eee/empfin/v1y1993i1p83-106.html>. (Cited on page 44.)
- D.Sornette, and J.V. Andersen. “The Dollar Game.” THEMA Working Papers 2003-30, THEMA (THéorie Economique, Modélisation et Applications), Université de Cergy-Pontoise, 2003. <http://EconPapers.repec.org/RePEc:ema:worpap:2003-30>. (Cited on page 174.)
- Dudley, R.M. Uniform Central Limit Theorems. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999. <http://books.google.fr/books?id=DtLp7Y76h-gC>. (Cited on page 36.)
- Dupre D., Jimenez-Garces S., Girerd-Potin I., and Louvet P. “Les investisseurs paient pour l’ethique. Conviction ou prudence ?” Banque et finance 84, 1: (2006) 45–60. (Cited on page 39.)
- Edwards, W. “Conservatism in human information processing.” Formal representation of human judgment 17–52. (Cited on page 40.)
- Enke, Stephen. “On maximizing profits: A distinction between Chamberlin and Robinson.” American Economic Review 41: (1951) American Economic Review. (Cited on page 28.)
- Epstein, J. M., and R. Axtell. Growing Artificial Societies: Social Science from the Bottom Up. Cambridge: MIT Press, 1996. (Cited on page 59.)
- Epstein, Joshua M. “Remarks on the Foundations of Agent-Based Generative Social Science.” In Handbook of Computational Economics, edited by Leigh Tesfatsion, and Kenneth L. Judd, Elsevier, 2006, volume 2 of Handbook of Computational Economics, chapter 34, 1585–1604. <http://ideas.repec.org/h/eee/hechhp/2-34.html>. (Cited on pages 21, 56, 59 and 73.)
- Erdo’s, P., and A Renyi. “On the Evolution of Random Graphs.” In PUBLICATION OF THE MATHEMATICAL INSTITUTE OF THE HUNGARIAN ACADEMY OF SCIENCES. 1960, 17–61. (Cited on page 57.)
- Evans D., Heuvelink A., and D. Nettle. “The evolution of optimism: a multi-agent based model of adaptive bias in human judgement.” Proceedings of the AISB’03 Symposium on Scientific Methods for the Analysis of Agent-Environment

- Interaction 3725470, University of Wales, 2003. <http://ideas.repec.org/p/hrv/faseco/3725470.html>. (Cited on page 28.)
- Fama, Eugene F. “The Behavior of Stock-Market Prices.” The Journal of Business 38, 1: (1965) 34–105. <http://dx.doi.org/10.2307/2350752>. (Cited on page 28.)
- . “Efficient Capital Markets: A Review of Theory and Empirical Work.” Journal of Finance 25, 2: (1970) 383–417. <http://EconPapers.repec.org/RePEc:bla:jfinan:v:25:y:1970:i:2:p:383-417>. (Cited on pages 14, 16, 27, 47 and 155.)
- Farmer, J. Doyne, and Shareen Joshi. “The price dynamics of common trading strategies.” Quantitative Finance Papers cond-mat/0012419, arXiv.org, 2000. <http://ideas.repec.org/p/arx/papers/cond-mat-0012419.html>. (Cited on pages 67, 73 and 80.)
- Feichtinger, Gustav. “The economy as an evolving complex system : P.W. Anderson, K.J. Arrow and D. Pines, eds., Addison-Wesley, Redwood City, California, 1988 pp. xvii + 313, 40.95(*cloth*),21.95 (paper).” Journal of Economic Behavior & Organization 15, 2: (1991) 316–319. <http://ideas.repec.org/a/eee/jeborg/v15y1991i2p316-319.html>. (Cited on page 59.)
- Fernandez, Pablo, Javier Aguirreamalloa, and Heinrich Liechtenstein. “The equity premium puzzle: High required equity premium, undervaluation and self fulfilling prophecy.” Iese research papers, IESE Business School, 2009. (Cited on page 35.)
- Fisher, I. Mathematical investigations in the theory of value and prices. Yale University Press, 1925. <http://books.google.fr/books?id=kNlfAAAAIAAJ>. (Cited on page 30.)
- Friedman, M. Essays in Positive Economics. Phoenix books. University of Chicago Press, 1953. <http://bks7.books.google.co.in/books?id=BDBdAAAAIAAJ>. (Cited on pages 27, 28 and 148.)
- Galbiati, Marco, and Kimmo Soramaki. “An agent-based model of payment systems.” Bank of England working papers 352, Bank of England, 2008. <http://ideas.repec.org/p/boe/boeewp/0352.html>. (Cited on page 58.)
- Galbiati, Marco, and Lucian Stanciu-Viziteu. “Target2 Geography.” (soon to appear) european central bank working papers, European Central Bank, 2013. <http://ideas.repec.org/p/boe/boeewp/0352.html>. (Cited on page 58.)
- Galla, Tobias, Giancarlo Mosetti, and Yi-Cheng Zhang. “Anomalous fluctuations in Minority Games and related multi-agent models of financial markets.” Papers physics/0608091, arXiv.org, 2006. <http://ideas.repec.org/p/arx/papers/physics-0608091.html>. (Cited on page 174.)

- Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler. "Is It Always Rational To Satisfy Savage's Axioms?" Economics and Philosophy 25, 03: (2009) 285–296. http://ideas.repec.org/a/cup/ecnphi/v25y2009i03p285-296_99.html. (Cited on page 28.)
- Glosten, Lawrence R., and Paul R. Milgrom. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." Discussion Papers 570, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 1983. <http://ideas.repec.org/p/nwu/cmsems/570.html>. (Cited on page 48.)
- Graham, B., and D.L.F. Dodd. Security Analysis: The Classic 1951 Edition. McGraw-Hill Education, 1951. <http://books.google.fr/books?id=uUruZA9baowC>. (Cited on page 14.)
- Griffin, D., A. Tversky, Stanford Center on Conflict, and Negotiation. The weighing of evidence and the determinants of confidence: by Dale Griffin and Amos Tversky. Working paper. Stanford Center on Conflict and Negotiation, Stanford University, 1993. <http://books.google.fr/books?id=dGIrAQAAIAAJ>. (Cited on page 41.)
- Grossman, Sanford J, and Joseph E Stiglitz. "On the Impossibility of Informationally Efficient Markets." Levine's Working Paper Archive 1908, David K. Levine, 1980. <http://ideas.repec.org/p/cla/levarc/1908.html>. (Cited on pages 18, 155, 184, 187 and 191.)
- Hamerle, Alfred, Thilo Liebig, and Hans-Jochen Schropp. "Systematic risk of CDOs and CDO arbitrage." Discussion Paper Series 2: Banking and Financial Studies 2009,13, Deutsche Bundesbank, Research Centre, 2009. <http://EconPapers.repec.org/RePEc:zbw:bubdp2:200913>. (Cited on page 173.)
- Harris, Larry. Trading and Exchanges: Market Microstructure for Practitioners. Oxford University Press, 2002. <http://EconPapers.repec.org/RePEc:oxp:obooks:9780195144703>. (Cited on page 42.)
- Herve, Alexandre. "la Quasi Marche Aleatoire." Finance 13, 2. (Cited on page 44.)
- . "L'efficience bruitée. Une analyse non linéaire du marché français des actions." 1994. (Cited on pages 25 and 55.)
- Hirshleifer, David, and Siew Hong Teoh. "Herd Behaviour and Cascading in Capital Markets: a Review and Synthesis." European Financial Management 9, 1: (2003) 25–66. <http://ideas.repec.org/a/bla/eufman/v9y2003i1p25-66.html>. (Cited on page 30.)
- Huang, Weihong, Huanhuan Zheng, and Wai-Mun Chia. "Financial crises and interacting heterogeneous agents." Journal of Economic Dynamics and

- Control 34, 6: (2010) 1105–1122. <http://ideas.repec.org/a/eee/dyncon/v34y2010i6p1105-1122.html>. (Cited on page 57.)
- Jordan, S. “Is Momentum a Self-fulfilling Prophecy?” Working Paper Series . (Cited on page 35.)
- Kahneman, D. “On the psychology of prediction.” Psychological Review 80: (1973) 237–251. <http://books.google.fr/books?id=85FOHAAACAAJ>. (Cited on pages 26 and 92.)
- Kahneman, D., P. Slovic, and A. Tversky. Judgement under uncertainty - Heuristics and biases. Cambridge University Press, Cambridge, 1981. (Cited on pages 39, 40, 92 and 173.)
- Kahneman, D., and A. Tversky. “Choices, values and frames.” American Psychologist 39, 4: (1984) 341–350. (Cited on page 26.)
- Kahneman, Daniel, and Amos Tversky. “Prospect Theory: An Analysis of Decision under Risk.” Econometrica 47, 2: (1979) 263–91. <http://ideas.repec.org/a/ecm/emetrp/v47y1979i2p263-91.html>. (Cited on pages 30, 39 and 48.)
- Kaneko, K., and I. Tsuda. Complex Systems: Chaos and Beyond : A Constructive Approach With Applications in Life Sciences. Physics and Astronomy On-line Library. Springer-Verlag GmbH, 2001. <http://books.google.fr/books?id=71cINfgupggC>. (Cited on page 53.)
- Kirman, Alan P. “Whom or What Does the Representative Individual Represent?” The Journal of Economic Perspectives 6, 2: (1992) 117–136. (Cited on page 37.)
- Kyle, Albert S. “Continuous Auctions and Insider Trading.” Econometrica 53, 6: (1985) 1315–35. <http://EconPapers.repec.org/RePEc:ecm:emetrp:v:53:y:1985:i:6:p:1315-35>. (Cited on page 37.)
- Landier, Augustin, David Sraer, and David Thesmar. “The risk-Shifting Hypothesis : Evidence from Subprime Originations.” TSE Working Papers 11-279, Toulouse School of Economics (TSE), 2011. <http://ideas.repec.org/p/tse/wpaper/25604.html>. (Cited on page 38.)
- Lebaron, B. “Building the Santa Fe Artificial Stock Market.” Physica A . (Cited on pages 62 and 63.)
- LeBaron, Blake. “Evolution and time horizons in an agent based stock market.” Macroeconomic Dynamics 5: (2001) 225–254. (Cited on page 71.)
- Levitin, Adam, and Susan M. Wachter. “Explaining the Housing Bubble.” Mpra paper, University Library of Munich, Germany, 2012. <http://EconPapers.repec.org/RePEc:pra:mprapa:41920>. (Cited on page 45.)

- Levy, H., M. Levy, and S. Solomon. Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena. Elsevier Science, 2000. <http://books.google.fr/books?id=ffcTlVuKZtoC>. (Cited on pages 65, 66 and 193.)
- Licata, I., and A. Sakaji. Physics Of Emergence and Organization. World Scientific, 2008. <http://books.google.ca/books?id=BS01KA-5CDIC>. (Cited on page 53.)
- Lintner, John. "Security Prices, Risk, And Maximal Gains From Diversification." Journal of Finance 20, 4: (1965) 587–615. <http://ideas.repec.org/a/bla/jfinan/v20y1965i4p587-615.html>. (Cited on page 30.)
- Litzenberger, Robert, Jeff Castura, and Richard Gorelick. "The Impacts of Automation and High Frequency Trading on Market Quality." Annual Review of Financial Economics 4, 1: (2012) 59–98. <http://EconPapers.repec.org/RePEc:anr:refeco:v:4:y:2012:p:59-98>. (Cited on page 46.)
- Lo, Andrew W., Harry Mamaysky, and Jiang Wang. "Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation." Journal of Finance 40: (2000) 1705–1765. (Cited on page 35.)
- Louvet, P., and M. Dubois. "Journal of banking and finance." The day-of-the-week effect: the international evidence 20: (1996) 1463–1484. (Cited on page 43.)
- Louvet, P., and Taramasco O. "L'effet "jour-de-la-semaine" a la Bourse de Paris : une approche par les moyennes mobiles"." Finance 11: (1990) 83–109. (Cited on page 43.)
- . "L'analyse descriptive de la dependance serielle sur les rentabilites boursieres." Finance 14: (1993) 67–94. (Cited on page 44.)
- . "Estimation non parametrique de la dependance serielle sur les rentabilites boursieres d'indices internationaux." Journal de la Societe statistique de Paris 1: (1997) 53–79. (Cited on page 44.)
- Lucas, Jr, Robert E. "Asset Prices in an Exchange Economy." Econometrica 46, 6: (1978) 1429–45. <http://ideas.repec.org/a/ecm/emetrp/v46y1978i6p1429-45.html>. (Cited on page 28.)
- Mandelbrot, Benoit. "The Variation of Certain Speculative Prices." The Journal of Business 36: (1963) 394. <http://ideas.repec.org/a/ucp/jnlbus/v36y1963p394.html>. (Cited on page 43.)
- Markowitz, Harry. "Portfolio Selection." The Journal of Finance 7, 1: (1952) 77–91. <http://dx.doi.org/10.2307/2975974>. (Cited on pages 30 and 47.)
- Matassini, Lorenzo, and Fabio Franci. "On financial markets trading." Physica A: Statistical Mechanics and its Applications 289, 3-4: (2001) 526–542. <http://www.sciencedirect.com/science/article/B6TVG-4234C34-F/1/ac555d0438273685da02c43609872456>. (Cited on page 64.)

- Mathieu, P., and O. Brandouy. “A generic architecture for realistic simulations of complex financial dynamics.” Advances in Practical Applications of Agents and Multiagent Systems 185–197. (Cited on page 62.)
- Mathieu, Philippe, and Olivier Brandouy. “Introducing ATOM.” In Proceedings of the 10th International conference on Practical Applications of Agents and Multi-Agents Systems (PAAMS’2012), edited by Javier Bajo Perez, Juan M. Corchado, Emmanuel Adam, Alfonso Ortega, Maria Moreno, Elena Navarro, Benjamin Hirsch, Henrique Lopes Cardoso, Vicente Julián, Miguel A. Sánchez, and Philippe Mathieu. Salamanca: Springer, 2012, Practical Advances in Intelligent and soft computing, 269–272. <http://www.paams.net/paams2012/>. Short paper, Best Demo award. (Cited on page 63.)
- Meister, K., M. Balmer, F. Ciari, A. Horni, M. Rieser, R.A. Waraich, K.W. Axhausen, R.A. Waraich, R.A. Waraich, K.W. Axhausen, et al. Large-scale agent-based travel demand optimization applied to Switzerland, including mode choice. ETH, Eidgenössische Technische Hochschule Zürich, IVT, Institut für Verkehrsplanung und Transportsysteme, 2010. <http://books.google.fr/books?id=yBwaywAACAAJ>. (Cited on page 58.)
- Mitroi Adrian, Stancu I. “Finante comportamentale versus analiza tehnica si fundamentala.” Symposion Scientifique International, Developpement durable, echange global et ecosystems 47, 1. (Cited on page 48.)
- Moro, Esteban. “The Minority Game: an introductory guide.” Advances in Condensed Matter and Statistical Physics 154, 3: (2004) 477+. (Cited on page 192.)
- Morris, John J., and Pervaiz Alam. “Value relevance and the dot-com bubble of the 1990s.” The Quarterly Review of Economics and Finance 52, 2: (2012) 243–255. <http://ideas.repec.org/a/eee/quaeco/v52y2012i2p243-255.html>. (Cited on pages 13 and 173.)
- Mossin, Jan. “Equilibrium in a Capital Asset Market.” Econometrica . (Cited on page 30.)
- Muth, John F. “Rational Expectations and the Theory of Price Movements.” Econometrica 29, 3: (1961) 315–335. (Cited on pages 16 and 28.)
- Nagel, E. The structure of science: problems in the logic of scientific explanation. Harcourt, Brace & World, 1961. <http://books.google.fr/books?id=fIhEAAAIAAJ>. (Cited on page 51.)
- Nash, Jr., John F. “The Bargaining Problem.” Econometrica 18: (1950) 155–162. (Cited on page 175.)

- Ofek, Eli, and Matthew Richardson. "DotCom Mania: The Rise and Fall of Internet Stock Prices." NBER Working Papers 8630, National Bureau of Economic Research, Inc, 2001. <http://ideas.repec.org/p/nbr/nberwo/8630.html>. (Cited on page 38.)
- O'Hara, M. Market Microstructure Theory. Blackwell Business. John Wiley & Sons, 1998. <http://books.google.fr/books?id=udXjR2Dg7bwC>. (Cited on page 48.)
- Paranjape, R., and A. Sadanand. Multi-Agent Systems for Healthcare Simulation and Modeling: Applications for System Improvement. Premier Reference Source. Medical Information Science Reference, 2010. <http://books.google.fr/books?id=zcD0sj2JvC4C>. (Cited on page 58.)
- Payzan-LeNestour, Elise. "Learning to Choose the Right Investment in an Unstable World: Experimental Evidence Based on the Bandit Problem." Technical Report 10-28, wiss Finance Institute Research Paper, 2012. <http://ssrn.com/abstract=1628657>. (Cited on page 178.)
- Raberto, M., S. Cincotti, C. Dose, S. M. Focardi, and M. Marchesi. "Price formation in an artificial market: limit order book versus matching of supply and demand." In Nonlinear Dynamics and Heterogenous Interacting Agents, edited by T. Lux, S. Reitz, and E. Samanidou, Springer-Verlag, Berlin, 2005. (Cited on page 74.)
- Raberto, Marco, Silvano Cincotti, Sergio M. Focardi, and Michele Marchesi. "Agent-based simulation of a financial market." Physica A: Statistical Mechanics and its Applications 299, 1-2: (2001) 319–327. <http://www.sciencedirect.com/science/article/B6TVG-447MXW6-1H/1/993b52b754926a4a7732e9d46cf9803d>. (Cited on pages 63 and 74.)
- Radner, Roy. "Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets." Econometrica 40, 2: (1972) 289–303. <http://EconPapers.repec.org/RePEc:ecm:emetrp:v:40:y:1972:i:2:p:289-303>. (Cited on page 28.)
- Ramnath, Sundaresh, Steve Rock, and Philip B. Shane. "Financial Analysts' Forecasts and Stock Recommendations: A Review of the Research." Foundations and Trends in Finance 2, 4: (2006) 311 – 421. (Cited on page 35.)
- R.E, Hodel. An introduction to mathematical logic. PWS Publishing Company, Boston, 1995. (Cited on page 76.)
- Rubinstein, Mark E. "Rational Markets: Yes or No? The Affirmative Case." Financial Analysts Journal 57, 3. (Cited on page 25.)
- S., Dupre D. Girerd-Potin I. Jimenez-Garces, and Louvet P. "Influence de la notation éthique sur l'évolution du prix des actions : un modele theorique." Revue Economique 60. (Cited on page 39.)

- Savage, L. J. The foundations of statistics. JohnWiley and Sons, 1954. (Cited on page 25.)
- Sharma, V. Herding in Financial Markets During Internet Bubble. VDM Publishing, 2009. <http://books.google.fr/books?id=NYfYPwAACAAJ>. (Cited on page 35.)
- Sharpe, William F. “Capital Asset Prices: A Theory Of Market Equilibrium Under Conditions Of Risk.” Journal of Finance 19, 3: (1964) 425–442. <http://ideas.repec.org/a/bla/jfinan/v19y1964i3p425-442.html>. (Cited on pages 15, 30, 47 and 185.)
- Shen, Zhen, Kai Wang, and Fenghua Zhu. “Agent-based traffic simulation and traffic signal timing optimization with GPU.” In Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on. 2011, 145 –150. (Cited on page 58.)
- Shiller, Robert J. “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?” Working Paper 456, National Bureau of Economic Research, 1980. <http://www.nber.org/papers/w0456>. (Cited on pages 40 and 131.)
- Silver, N. The Signal and the Noise: Why So Many Predictions Fail-but Some Don't. Penguin Group US, 2012. http://books.google.fr/books?id=SI-VqAT4_hYC. (Cited on page 46.)
- Smaldino, Paul, Cynthia Pickett, Jeffrey Sherman, and Jeffrey Schank. “An Agent-Based Model of Social Identity Dynamics.” Journal of Artificial Societies and Social Simulation 15, 4: (2012) 7. <http://jasss.soc.surrey.ac.uk/15/4/7.html>. (Cited on page 58.)
- Stanciu-Viziteu, LucianDaniel. “The shark game: equilibrium with bounded rationality.” In Managing Market Complexity, edited by Andrea Teglio, Simone Alfarano, Eva Camacho-Cuena, and Miguel Ginés-Vilar, Springer Berlin Heidelberg, 2013, volume 662 of Lecture Notes in Economics and Mathematical Systems, 103–111. http://dx.doi.org/10.1007/978-3-642-31301-1_9. (Cited on page 192.)
- Stiglitz, Joseph E. “Symposium on Bubbles.” Journal of Economic Perspectives 4, 2: (1990) 13–18. <http://EconPapers.repec.org/RePEc:aea:jecper:v:4:y:1990:i:2:p:13-18>. (Cited on page 44.)
- Subrahmanyam, Avanidhar. “Behavioural Finance: A Review and Synthesis.” European Financial Management 14, 1: (2008) 12–29. <http://ideas.repec.org/a/bla/eufman/v14y2008i1p12-29.html>. (Cited on page 30.)
- Svenson, Ola. “Are we all less risky and more skillful than our fellow drivers?” Acta Psychologica 47, 2: (1981) 143–148. [http://dx.doi.org/10.1016/0001-6918\(81\)90005-6](http://dx.doi.org/10.1016/0001-6918(81)90005-6). (Cited on page 40.)

- Tesfatsion, Leigh. “Agent-Based Computational Economics: A Constructive Approach to Economic Theory.” In Handbook of Computational Economics, edited by Leigh Tesfatsion, and Kenneth L. Judd, Elsevier, 2006a, volume 2 of Handbook of Computational Economics, chapter 16, 831–880. <http://ideas.repec.org/h/eee/hecchp/2-16.html>. (Cited on pages 61 and 84.)
- . “Agent-Based Computational Modeling And Macroeconomics.” Staff General Research Papers 12402, Iowa State University, Department of Economics, 2006b. (Cited on page 59.)
- Thurner, S., J.D. Farmer, and J. Geanakoplos. “Leverage causes fat tails and clustered volatility.” Quantitative Finance 12: (2012) 695–707. (Cited on pages 47 and 71.)
- Tobin, James. “Liquidity Preference as Behavior Towards Risk.” Cowles Foundation Discussion Papers 14, Cowles Foundation for Research in Economics, Yale University, 1956. <http://ideas.repec.org/p/cwl/cwldpp/14.html>. (Cited on page 30.)
- Valdez, S. An Introduction to Global Financial Markets, Fifth Edition. Palgrave Macmillan, 2007. http://books.google.com.au/books?id=_MzEQgAACAAJ. (Cited on page 14.)
- Veryzhenko, I., O. Brandouy, and P. Mathieu. “Agent’s minimal intelligence calibration for realistic market dynamics.” Progress in Artificial Economics 3–14. (Cited on pages 61, 62 and 63.)
- Veryzhenko, I., P. Mathieu, and O. Brandouy. “Key points for realistic agent-based financial market simulations.” Proceedings of ICAART (3rd conf.) 2: (2011) 74–83. (Cited on pages 61 and 62.)
- Veryzhenko, Iryna. “A reexamination of modern finance issues using Artificial Market Frameworks.” 2012. (Cited on pages 56, 59 and 63.)
- Williams, J.B. The theory of investment value. Harvard University Press, 1938. <http://books.google.fr/books?id=cIhCAAAIAAJ>. (Cited on page 30.)
- Yan, Hongjun. Natural Selection in Financial Markets: Does it Work? 2008. <http://EconPapers.repec.org/RePEc:yism:somwrk:amz2648>. (Cited on page 29.)

The influence of investor biases on the formation of prices in financial markets

Abstract: We construct an agent-based computer simulated financial market. Trading in this market is not continuous. The market price is formed using a limit-order book. The modelled investors receive biased information and they attempt to maximize their wealth. Different traders, from noise to chartist and informed, coexist in the same market. We show how stylized facts can be formed by the presence of chartists or a simple lag in investor information. Price bubbles can arise when market prices are dominated by technical traders. Interestingly we show that well informed investors can earn more if they adopt, in special situations, a technical strategy. Using our results we propose a new model for market dynamics called "sometimes efficient markets". Moreover, we define the concept of "strategy-strong efficient markets".

Keywords: computational finance, stock markets, efficiency, multi-agent, simulation, bubbles, stylized facts, risk management
