



Davi D'Andréa Baccan

Contributions of Computational Cognitive Modeling to the Understanding of the Financial Markets

Tese de Doutoramento
do Programa de Doutoramento em Ciências e Tecnologias da Informação
orientada pelo Professor Doutor Luís Miguel Machado Lopes Macedo
e apresentada ao departamento de Engenharia Informática
a Faculdade de Ciências e Tecnologia da universidade de Coimbra

novembro de 2017



UNIVERSIDADE DE COIMBRA

Contributions of Computational Cognitive Modeling to the Understanding of the Financial Markets

Davi D'Andréa Baccan

Dissertation submitted to the University of Coimbra
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

November 2017



Department of Informatics Engineering
Faculty of Science and technology
University of Coimbra

This research has been developed as part of the requirements of the Doctoral Program in Information Science and Technology of the Faculty of Sciences and Technology of the University of Coimbra. This work was carried out in the Cognitive and Media Systems Group of the Center for Informatics and Systems of the University of Coimbra (CISUC). Funding for this work was partially provided by the Portuguese Foundation for Science and Technology through the contract SFRH/BD/60700/2009, and by the entrepreneurship project TribeCA (Trading and investing with behavioral-emotional Cognitive Agents), funded by the “Fundo Europeu de Desenvolvimento Regional” (FEDER) through the “Programa Operacional Regional do Centro”.

This work has been supervised by Professor **Luís Miguel Machado Lopes Macedo**, Assistant Professor of the Department of Informatics Engineering of the Faculty of Sciences and Technology of the University of Coimbra.

To my beloved wife, Tarícia

When the mountain is high,
Just look up to the sky,
Ask God to teach you,
Then persevere with a smile.

Lenny Kravitz

Abstract

The financial markets are complex, dynamic, and strategic socio-economic systems in which a great number of heterogeneous market participants interact by essentially buying and selling assets of different types. The financial markets such as stock markets serve many functions. For instance, the financial markets help market participants and companies in improving the capital allocation process. Additionally, the behavior of the financial markets is assumed to be an important gauge for helping the understanding of the current and future state of companies and, ultimately, of the whole economic and financial system. However, the importance of the financial markets might be better noticed when they do not fulfill their primary functions, specifically and most dramatically when the financial markets crash. On September 2008, for example, a series of events threatened the stability of the world's financial system. Some gigantic financial services companies had unexpectedly failed and had to be rescued by governments while others simply filed for bankruptcy. The world's financial system came close to a meltdown. Although disaster had somehow been averted due to a series of actions, the Crash of 2008 had immediate, profound, vast and long-lasting consequences for the world economy.

The true understanding of the financial markets are indeed quite difficult. Several hypotheses have been proposed to try to explain the behavior of the market participants individually as well as the behavior of markets as a whole. On the one hand, traditional economic theories such as the Efficient Market Hypothesis tend to assume that market participants are rational as well as that markets are efficient. However, behavioral economics research has been providing extensive and vast evidence that market participants have what is known as behavioral biases, i.e., deviations from the so-called rational behavior. In addition to the behavioral economics evidence, disciplines such as cognitive neuroscience and neuroeconomics have been clarifying the role of emotions (e.g., happiness, unhappiness, surprise) for the human reasoning, memory system and processes, and decision-making process. For instance, emotions play a very important role in the memory processes of encoding and retrieving as well as are the basis of a sort of learning system. On the other hand, recent theories such as the Adaptive Market Hypothesis tries to reconcile market efficiency with behavioral economics by acknowledging the importance of emotions, the existence of behavioral biases, and the occurrence of interesting phenomena and anomalies such as bubbles.

The Crash of 2008 together with new evidence provided by different research areas have been stressing the need for novel and interdisciplinary approaches for the study of economic and financial systems and problems. One of these approaches is the use of Agent-based Financial Markets. Agent-based Financial Markets allows researchers to depart from classical assumptions in order to test different hypotheses, concepts, ideas, etc, making it possible the design and realization of more realistic and behavioral plausible experiments. This thesis is in line with this context. In this exploratory work we aim to investigate which contributions the application of a cognitive modeling approach might bring to the understanding of the financial markets, specifically to the behavior of both market participants (individually) and the financial markets (globally). The starting point is to empower artificial agents with mechanisms similar to or inspired in those used by humans so that we have artificial cognitive agents, i.e., artificial agents with different memory systems and processes, the capacity of recognizing, simulating and expressing emotions, decision-making processes, the ability to receive and process different kinds of information, and the ability to learn.

To this end, we first conceive a generic novel cognitive model of individual market participants (human agents) named TribeCA (Trading and investing with behavioral-emotional Cognitive Agents). TribeCA is based on the Belief-Desire Theory of Emotions (BDTE), on the cognitive-psychoevolutionary model of surprise proposed by Meyer and colleagues, and on the artificial surprise model proposed by Macedo and Cardoso. Then we provide an implementation of the proposed model which is later integrated into two tools used in the context of agent-based financial markets. The resulting platform allows the design and realization of a variety of economic and financial experiments with artificial cognitive agents. We carried out three agent and multi-agent based experiments to address some fundamental aspects regarding the financial markets such as efficiency and rationality. Additionally, we carried out two case studies on comparing the traditional (economics and finance) perspective with the cognitive science perspective on modeling and computing surprise in economics and finance.

This thesis provides contributions to the advance in the design and realization of interdisciplinary approaches to the study of economic and financial systems or problems. Our generic conceptual cognitive model and implementation might be used both to explore other aspects of the financial markets in addition to those addressed in this work and to other agent-based models. We consider this work opens up a set of novel possibilities for investigations in the academia and in the industry. In the end, we

may have a better understanding of the behavior of market participants individually as well as of the financial markets globally. It has the potential to result in, for instance, the development of novel (potentially better and highly lucrative) financial services to support the market participants' decision-making process based on his/her emotions, behavior, etc.

Keywords:

Computational cognitive modeling; Financial Markets; Emotions; Decision-making; Agent-based Computational Economics; Cognitive Emotion Theories; Adaptive Market Hypothesis; Behavioral Economics; Efficient Market Hypothesis

Resumo

Os mercados financeiros são sistemas socioeconômicos complexos, dinâmicos e estratégicos nos quais um grande número de heterogêneos participantes interagem por meio da compra e venda de diferentes tipos de ativos. Os mercados financeiros tais como os mercados de ações possuem múltiplas funções. Por exemplo, os mercados financeiros propiciam meios para que os participantes e as companhias façam um melhor processo de alocação de capital. Para além disso, o comportamento dos mercados financeiros são geralmente considerados como importantes medidas/sinais para a compreensão do estado atual e futuro das companhias e, em última análise, de todo o sistema econômico e financeiro. Entretanto, a importância dos mercados financeiros poderá ser melhor compreendida quando os mercados não cumprem suas funções primordiais, mais especificamente e de forma dramática quando ocorre um *crash* nos mercados financeiros.

Em Setembro de 2008, por exemplo, uma série de eventos ameaçou a estabilidade do sistema financeiro mundial. Gigantescas empresas dos mercados financeiros inesperadamente falharam e tiveram de ser resgatadas pelos seus respectivos governos, enquanto que outras simplesmente entraram com pedido de falência. O sistema financeiro mundial esteve próximo do colapso. Embora o desastre tenha, de certa forma, sido evitado, o *Crash* de 2008 teve consequências imediatas, profundas, e duradouras para a economia mundial. A verdadeira compreensão dos mercados financeiros é de fato difícil. Diferentes hipóteses têm sido propostas para explicar o comportamento dos participantes dos mercados financeiros de forma individual bem como o comportamento dos mercados de forma global. Por um lado, teorias econômicas tradicionais como, por exemplo, a Hipótese do Mercado Eficiente, tendem a considerar que os participantes dos mercados financeiros são racionais e que os mercados são eficientes. Entretanto, pesquisas na área de economia comportamental têm fornecido extensa e vasta evidência que demonstra que os participantes dos mercados financeiros possuem desvios comportamentais do chamado comportamento racional. Para além das evidências da economia comportamental, disciplinas tais como a neurociência cognitiva e a neuroeconomia têm clarificado a função das emoções (e.g., felicidade, tristeza, surpresa) no processo de raciocínio e tomada de decisão, aprendizado, bem como a importância das emoções no âmbito da memória humana, particularmente para os processos de codificação e recuperação de memórias. Por outro lado, teorias recentes como a Hipótese do Mercado Adaptativo tentam reconciliar a ideia de mercados eficientes com a economia comportamental, ao

reconhecer a importância das emoções, a existência de desvios comportamentais, e a ocorrência de fenômenos e anomalias como as bolhas.

O *Crash* de 2008 conjuntamente com novas evidências fornecidas por diferentes áreas têm salientado a necessidade de novas e interdisciplinares abordagens para o estudo de sistemas e problemas econômicos e financeiros. Uma destas abordagens é o uso de *Agent-based Financial Markets*. Esta abordagem permite aos investigadores se distanciarem das tradicionais crenças a fim de testar novas hipóteses, conceitos, ideias, etc, tornando possível o projeto e realização de experimentos mais realistas e mais plausíveis em termos comportamentais. Esta tese está em linha com este contexto. Neste trabalho exploratório, nosso objetivo é investigar quais contribuições a aplicação de uma abordagem de modelação cognitiva pode trazer para a compreensão dos mercados financeiros, especificamente para o comportamento dos participantes dos mercados financeiros (individualmente) e dos mercados financeiros (globalmente). O ponto de partida é a criação de agentes artificiais com mecanismos similares aos ou inspirados nos usados pelos seres humanos de modo que seja possível conceber agentes artificiais cognitivos, i.e., agentes artificiais com diferentes sistemas de memórias e processos, com a capacidade de reconhecer, simular, e expressar emoções, diferentes processos de tomada de decisão, com a habilidade de receber e processar diferentes tipos de informação, e com a habilidade de aprender.

Para este fim, nós primeiro concebemos um modelo cognitivo genérico individual dos participantes dos mercados financeiros (agentes humanos) intitulado TribeCA (*Trading and investing with behavioral-emotional Cognitive Agents*). O modelo cognitivo proposto é baseado na *Belief-Desire Theory of Emotions (BDTE)*, no modelo *cognitive-psychoevolutionary* de surpresa proposto por Myer e colegas, e no modelo de surpresa artificial proposto por Macedo e Cardoso. De seguida nós fornecemos uma implementação do modelo proposto a qual foi posteriormente integrada a duas ferramentas utilizadas no contexto dos *agent-based financial markets*. A plataforma resultante permite o projeto e realização de uma variedade de experimentos econômicos e financeiros com agentes artificiais cognitivos. Nós realizamos neste trabalho três experimentos com agentes e multi-agentes a fim de endereçar alguns aspectos fundamentais dos mercados financeiros tais como eficiência e racionalidade. Adicionalmente, nós realizamos dois estudos de casos a fim de comparar a perspectiva tradicional (econômica e financeira) com a perspectiva da ciência cognitiva na modelação e computação da surpresa na economia e finança.

Esta tese fornece contribuições para o avanço no projeto e realização de abordagens interdisciplinares para o estudo de sistemas ou problemas económicos e financeiros. Nosso modelo cognitivo genérico e sua implementação podem ser utilizados a fim de que sejam explorados outros aspectos dos mercados financeiros, para além dos que foram endereçados neste trabalho, e em outros modelos baseados em agentes. Nós consideramos que este trabalho abre um novo conjunto de possibilidades para investigações quer na academia quer na indústria. Ao final, nós poderemos obter uma melhor compreensão e entendimento sobre o comportamento dos participantes dos mercados financeiros (individualmente) bem como dos mercados financeiros (globalmente). Estas investigações poderão resultar, por exemplo, no desenvolvimento de novos (potencialmente melhores e altamente lucrativos) serviços financeiros para suportar o processo de tomada de decisão dos participantes dos mercados financeiros baseados nas suas emoções, comportamentos, etc.

Palavras-chave:

Modelação cognitiva computational; Mercados Financeiros; Emoções; Processo de tomada de decisão; *Agent-based Computational Economics*; *Agent-based Financial Markets*; Teorias Cognitivas de Emoções; Hipótese do Mercado Adaptativo; Economia Comportamental; Hipótese dos Mercados Eficientes

Acknowledgements

Getting a PhD is a truly extraordinary journey but, at the same time, very hard to achieve. Getting a PhD abroad is a journey even more extraordinary.

To my beloved wife and best friend, Tarícia Marques Toledo Baccan. Dear Tarícia, the love of my life, there are no words to express my gratitude. We know how difficult and hard this journey was. I will be forever and ever thankful for you for all the unconditional and endless support you gave me. Without you this journey would not have been possible. Without you this journey would not make any sense at all. Thank you for believing in me and in my dreams. Thank you for being here with me. I hope we can stay together forever. Thank you for making me so happy. I hope we can continue bringing happiness to each other in our future journeys: you and me, you and me and ... :) With all my love and care. I love you so much!!!

To my father, my mother, and my sister, thank you for everything. Thank you for believing in my dreams too and for giving me all the necessary support you always gave me. We are physically apart but our hearts and minds are certainly together. I love you!

To God, Lord of all things, for giving me the faith, strength and health to successfully conclude this journey.

To my supervisor, Professor Luís Miguel Machado Lopes Macedo, thank you for giving me the freedom to pursue the research area and topic that truly fascinated me. Thank you for all your support and help during all these years, your orientation and guidance.

To my great friend and Masters supervisor, Professor Felipe Afonso de Almeida (in memoriam), thank you for having shared so many things with me. Thank you for having taught me so much about academics and the scientific method. I wish you were here Master. I miss our conversations. I will always miss you my friend.

To Professors Carlos Henrique Costa Ribeiro, Franklina Maria Bragion de Toledo, and Jauvane Cavalcante de Oliveira, thank you for giving me support in the beginning of this journey.

To my dear friends, especially Carlos Nuno Bizarro e Silva Laranjeiro, Naghmeh Ramezani Ivaki, Naaniel Vicente Mendes, and Ivano Irrera, thank you for your friendship, for all the good moments we spent together. You will always be in my heart and in my mind.

To my friend Ricardo Manuel da Conceição Rodrigues, thank you for proofreading the thesis.

List of Publications

This thesis relies on the published scientific research presented in the following peer reviewed papers:

- Baccan, D. Macedo, L., Sbruzzi, E. On the cognitive surprise in risk management: an analysis of the Value-at-Risk (VaR) historical. 17th Portuguese Conference on Artificial Intelligence (EPIA), ISBN 978-3-319-23485-4, Coimbra, Portugal, 2015.
- Baccan, D. Macedo, L., Sbruzzi, E. Towards modeling surprise in economics and finance: a cognitive science perspective. 7th European Starting AI Researcher Symposium (STAIRS), 21st European Conference on Artificial Intelligence (ECAI), *Frontiers in Artificial Intelligence and Applications*, volume 264, pp. 31-40, ISBN: 978-1-61499-420-6, Czech Republic, Prague, 2014.
- Baccan, D. Macedo, L., Sbruzzi, E. Is the El Farol more efficient when cognitive rational agents have a larger memory size?. IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 39-44, San Diego California, USA, 2014.
- Baccan, D. Macedo, L. Are markets more efficient when cognitive agents act rationally? Evidence from the El Farol. 16th Portuguese Conference on Artificial Intelligence (EPIA), pp. 438-449, ISBN: 978-989-95489-1-6, Azores, Portugal, 2013.
- Baccan, D. Macedo, L. Revisiting the El Farol problem: a cognitive modeling approach. *Multi-Agent-Based Simulation XIII, MABS 2012, Lecture Notes in Computer Science 7838*, pp. 56-68, ISBN: 978-3-642-38858-3, Springer Berlin Heidelberg, 2013.
- Baccan, D. Macedo, L. The Role of Surprise in Agent-based Computational Economics (ACE). *Doctoral Symposium on Artificial Intelligence (SDIA), 15th Portuguese Conference on Artificial Intelligence (EPIA)*, pp. 885-896, ISBN: 978-989-95618-4-7, Lisbon, Portugal, 2011.

- Baccan, D. Macedo, L. Endowing artificial agents with emotional autobiographical memories. Doctoral Symposium, IEEE Conference on Affective Computing and Intelligent Interaction (ACII), pp. 73-80, ISSN 09290672, Amsterdam, Netherlands, 2009.
- Baccan, D. Macedo, L. Emotional autobiographical memories for artificial agents. Doctoral Symposium on Artificial Intelligence (SDIA), 14th Portuguese Conference on Artificial Intelligence (EPIA), Aveiro, Portugal, 2009.

Table of Contents

List of Figures	xxiii
------------------------	--------------

List of Tables	xxix
-----------------------	-------------

1 Introduction	1
1.1 Goals	5
1.2 Approach	6
1.3 Contributions	8
1.4 Thesis structure	9
2 Background Concepts	11
2.1 Computational Cognitive Modeling	12
2.1.1 Emotions: definition, representation and functions	12
2.1.2 Cognitive Emotion Theories	20
2.1.3 Surprise: definition, functions and process	23
2.1.4 Artificial Surprise	26
2.1.5 Computational cognitive modeling of the coin tossing game	29
2.2 Financial Markets	31
2.2.1 Stock Markets	31
2.2.2 Investment Strategies	37
2.2.3 Efficient Markets Hypothesis	41
2.2.4 Behavioral Economics	43
2.2.5 Adaptive Markets Hypothesis	50
2.2.6 Surprise in economics and finance: importance, evidence and computation	53
2.3 Agent-based Financial Markets	57
2.3.1 Questions in designing artificial agents in the context of Agent- based Financial Markets	59
2.3.2 Agent-based Financial Markets platforms	60
2.3.3 Agent-based Models for studying financial markets	61
2.4 Related Work	67

3	TribeCA: a Conceptual Cognitive Model and Implementation	71
3.1	TribeCA: a Conceptual Cognitive Model	72
3.2	A cognitive modeling approach to the El Farol	74
3.2.1	Cumulative Efficiency, Cumulative Happiness, and Happiness Dis- tribution	78
4	Traditional Modeling versus Cognitive Modeling of Surprise in Eco- nomics and Finance	81
4.1	Case Study I: Surprise in Market Consensus	82
4.1.1	Results and discussion	86
4.2	Case Study II: Surprise in Risk Management	96
4.2.1	Cognitive surprise of using Value-at-Risk historical	98
4.2.2	Results and discussion	106
5	Agent-based Experimental Evaluation with the El Farol	115
5.1	Agent-based Experiments with the El Farol	116
5.2	Experiment I: Revisiting the El Farol	117
5.3	Experiment II: The impact of rationality on market efficiency	134
5.4	Experiment III: The impact of memory size on market efficiency	146
6	Conclusion	157
6.1	Contributions	157
6.2	Future work	161
Appendix A TribeCA: design and implementation		167
A.1	JABM classes	167
A.2	Cognitive model classes	173
Bibliography		177

List of Figures

Figure 2.1	Schematic representation of the valence and arousal model (extracted from Russell and Barrett (1999)).	14
Figure 2.2	Schematic representation of the BBC (Belief-Belief Comparator) and the BDC (Belief-Desire Comparator) (adapted from Reisenzein (2009)).	22
Figure 2.3	Schematic representation of the cognitive-evolutionary model of surprise process proposed by Meyer and colleagues (adapted from Macedo et al. (2009)).	25
Figure 2.4	Surprise as a function of the difference between $P(E_h)$ and $P(E_g)$, i.e., Equation 2.1 that is $S(E_g) = \log_2(1+P(E_h)-P(E_g))$ (Macedo et al., 2004).	29
Figure 2.5	Daily evolution of the S&P500 index from October 2005 to January 2015. Figure generated with the Economic Research Federal Reserve Bank of St. Louis tool - http://research.stlouisfed.org/ - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.	34
Figure 2.6	Daily volatility of the S&P500 from October 2005 to January 2015. Figure generated with the Economic Research Federal Reserve Bank of St. Louis tool - http://research.stlouisfed.org/ - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.	34
Figure 2.7	Evolution of the S&P500 index, U.S. Real Gross Domestic Product (GDP), and U.S. Civilian Unemployment Rate, from October 2005 to January 2015. Figure generated with the Economic Research Federal Reserve Bank of St. Louis tool - http://research.stlouisfed.org/ - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.	36
Figure 2.8	<i>S&P500</i> index, and corporate business profits before tax, from July 2005 to July 2015, quarterly adjusted, generated with the Economic Research Federal Reserve Bank of St. Louis site - http://research.stlouisfed.org/ - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.	55
Figure 2.9	El Farol bar problem: an example of the dynamics of attendance over time for a same run.	64
Figure 3.1	Architecture of the proposed artificial cognitive agent.	73

Figure 4.1	Unemployment rate (left) versus real data versus number of forecasts (right).	86
Figure 4.2	Unemployment rate forecasts: some statistical indicators.	87
Figure 4.3	Unemployment rate: real data versus mean forecast.	87
Figure 4.4	Unemployment rate: real data to mean forecast ratio.	88
Figure 4.5	Unexpected Unemployment Rate (Equation 4.1).	88
Figure 4.6	Scaled Unexpected Unemployment Rate (Equation 4.2).	89
Figure 4.7	Standardized Unexpected Unemployment Rate (Equation 4.3).	89
Figure 4.8	Comparison of the dynamics of Equations 4.1, 4.2 and 4.3.	90
Figure 4.9	Unemployment rate: mean forecast, standard deviation, and real data.	90
Figure 4.10	Real data is higher either than +2 standard deviation or -2 standard deviation of the mean forecast, and therefore there is a violation of the 0.954 confidence level.	91
Figure 4.11	Surprise intensity according to the artificial surprise model proposed by Macedo and colleagues (Equation 4.4).	92
Figure 4.12	Comparison of the dynamics of Equations 4.1, 4.2 and 4.3 with the surprise intensity (vertical dashed line) according to the artificial surprise model proposed by Macedo and colleagues (Equation 4.4).	92
Figure 4.13	Example of a $VaR_{historical}$ computation considering daily returns, a begin as of 01-03-2004, four window sizes namely 50 (from begin to 10-05-2004), 125 (from begin to 26-08-2004), 250 (from begin to 24-02-2005), and 500 (from begin to 22-02-2006) values, and two confidence levels (p) equal 0.99 and 0.95.	98
Figure 4.14	Whole period: S&P500, daily close.	99
Figure 4.15	Whole period: S&P500, daily return.	100
Figure 4.16	Calm period: S&P500, daily close.	100
Figure 4.17	Calm period: S&P500, daily return	101
Figure 4.18	Crash period: S&P500, daily close.	101
Figure 4.19	Crash period: S&P500, daily return.	102
Figure 4.20	Rolling window size approach: a schematic representation for the Calm period for a given simulation.	103
Figure 4.21	Calm and Crash period: S&P500, density of daily returns.	106
Figure 4.22	Calm and Crash period: S&P500, box plot of daily returns.	107
Figure 4.23	Rolling window size approach: treatments for the Calm period.	107
Figure 4.24	Alpha approach: treatments for the Calm period.	108
Figure 4.25	Rolling window size approach: treatments for the Crash period with VaR_{95} and VaR_{99}	108
Figure 4.26	Alpha approach: treatments for the Crash period.	109
Figure 4.27	Comparing different Rolling window size treatments with VaR_{95} for the Calm and Crash period.	110
Figure 4.28	Comparing different Rolling window size treatments with VaR_{99} for the Calm and Crash period.	110

Figure 4.29 Comparing different Alpha treatments for the Calm and Crash period. 111

Figure 5.1 *GBS* (Global Belief in Strategy) for treatments of the first group. 120

Figure 5.2 *CuE* (Cumulative Efficiency) for treatments of the first group. . . 120

Figure 5.3 *CuH* (Cumulative Happiness) for treatments of the first group. . 121

Figure 5.4 Cumulative *GSu* (Global Surprise) for treatments of the first group. 121

Figure 5.5 Comparing the mean of treatments T1(*MS*=60), T2(*MS*=100), T3(*MS*=500), and T4(*MS*=2500) of the first group, i.e., all strategies, changing of strategy not allowed, in terms of *GBS*, *CuH*, Cumulative *GSu*, and *CuE* (clockwise, from top left). 122

Figure 5.6 *GBS* (Global Belief in Strategy) for treatments of the second group. 123

Figure 5.7 *CuE* (Cumulative Efficiency) for treatments of the second group. 123

Figure 5.8 *CuH* (Cumulative Happiness) for treatments of the second group. 124

Figure 5.9 Cumulative *GSu* (Global Surprise) for treatments of the second group. 124

Figure 5.10 Comparing the mean of treatments T5(*MS*=60), T6(*MS*=100), T7(*MS*=500), and T8(*MS*=2500) of the second group, i.e., all strategies, changing of strategy allowed, in terms of *GBS*, *CuH*, Cumulative *GSu*, and *CuE* (clockwise, from top left). 125

Figure 5.11 *GBS* (Global Belief in Strategy) for treatments of the third group. 126

Figure 5.12 *CuE* (Cumulative Efficiency) for treatments of the third group. . 126

Figure 5.13 *CuH* (Cumulative Happiness) for treatments of the third group. . 127

Figure 5.14 Cumulative *GSu* (Global Surprise) for treatments of the third group. 127

Figure 5.15 Comparing the mean of treatments T9(*MS*=60), T10(*MS*=100), T11(*MS*=500), and T12(*MS*=2500) of the third group, i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed, in terms of *GBS*, *CuH*, Cumulative *GSu*, and *CuE* (clockwise, from top left). 128

Figure 5.16 Comparing *GBS* (Global Belief Strategy) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed). 129

Figure 5.17 Comparing *CuE* (Cumulative Efficiency) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed). 130

Figure 5.18 Comparing *CuH* (Cumulative Happiness) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed). 131

Figure 5.19 Comparing Cumulative GSu (Global Surprise) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., NTS (Noise Trader Strategy), changing of strategy not allowed).	132
Figure 5.20 Attendance for four runs of Treatment 1: $BT=0.20$	136
Figure 5.21 Attendance for four runs of Treatment 2: $BT=0.25$	136
Figure 5.22 Attendance for four runs of Treatment 3: $BT=0.30$	137
Figure 5.23 Attendance for four runs of Treatment 4: $BT=0.35$	137
Figure 5.24 Attendance for four runs of Treatment 5: $BT=0.40$	138
Figure 5.25 Attendance for four runs of Treatment 6: $BT=0.45$	138
Figure 5.26 Attendance for four runs of Treatment 7: $BT=0.50$	139
Figure 5.27 CuE (Cumulative Efficiency) over time, zoomed in.	140
Figure 5.28 CuE (Cumulative Efficiency), box plot.	140
Figure 5.29 CuH (Cumulative Happiness) over time, zoomed in.	141
Figure 5.30 CuH (Cumulative Happiness), box plot.	141
Figure 5.31 HaD (distribution of happiness among agents): density.	142
Figure 5.32 HaD (distribution of happiness among agents), box plot.	142
Figure 5.33 Frequency of strategies for each treatment, from Treatment 1 (BT_20) to Treatment 7 (BT_50), over time.	143
Figure 5.34 Number of changes of strategies per round over time.	144
Figure 5.35 Number of changes of strategies per round, box plot.	144
Figure 5.36 CuE (Cumulative Efficiency) over time, zoomed in.	148
Figure 5.37 CuE (Cumulative Efficiency), box plot.	149
Figure 5.38 CuH (Cumulative Happiness) over time, zoomed in.	149
Figure 5.39 CuH (Cumulative Happiness), box plot.	150
Figure 5.40 HaD (distribution of happiness among agents), density.	150
Figure 5.41 HaD (distribution of happiness among agents), box plot.	151
Figure 5.42 Frequency of strategies of some treatments, namely Treatment 1 (MS_10), Treatment 10 (MS_100), Treatment 19 (MS_1000), Treatment 29 (MS_2000), and Treatment 34 (MS_2500).	152
Figure 5.43 Frequency of the NTS (Noise Trader Strategy) compiled from all treatments.	153
Figure 5.44 Number of changes of strategy per round over time.	153
Figure 5.45 Number of changes of strategy per round, box plot.	154
Figure A.1 JABM main classes.	168
Figure A.2 JABM: Simulation interface.	168
Figure A.3 JABM: Population class.	169
Figure A.4 JABM: SimulationController class.	170
Figure A.5 JABM: Agent interface.	171
Figure A.6 JABM: AbstractAgent class.	171
Figure A.7 JABM: AgentList class.	172
Figure A.8 Conceptual model: main classes of the artificial cognitive agent.	174

List of Figures

Figure A.9 Conceptual model: BBC and BDC classes. 175

List of Tables

Table 2.1	Belief-desire theory of emotions, qualitative formulation (adapted from (Reisenzein, 2009)). Notation: $Bel(p, t)$ believes p at time t ; $Certain(p, t)$ firmly believes p at time t ; $Uncertain(p, t) \equiv Bel(p, t) \wedge \neg Certain(p, t) \wedge \neg Certain(\neg p, t)$; $Desire(p, t)$ desires p at time t ; $Desire(\neg p, t)$ desires not- p at t (\approx is aversive against p at time t).	23
Table 2.2	Belief-desire theory of emotions, quantitative formulation (adapted from (Reisenzein, 2009)). Notation: $b(p, t)$ represents the strength of belief in p at time t , with 1.0 denoting certainty that p , 0.5 maximal uncertainty, and 0.0 certainty that not- p ; $d(p, t)$ represents the direction and strength of the desire for p at time t , with values > 0 denoting positive desire in favor of p , 0 indifference, and values < 0 aversion against p	23
Table 4.1	Unemployment Rate: date, number of forecasts, consensus estimate (mean), standard deviation of the forecasts, and real Unemployment Rate date, in a total of 25 observations.	84
Table 4.2	Statistical data about the Calm and Crash period.	106
Table 5.1	Set of available strategies for all agent-based experiments and treatments.	116
Table 5.2	Experiment with different strategies available (Strategies), the ability of changing current strategy or not (Fixed), and memory size (MS). For all treatments, belief threshold (BT) = 0.5.	119
Table 5.3	Experiment with different belief thresholds (BT). For all treatments, memory size (MS) = 100.	135
Table 5.4	Experiment with different memory sizes (MS). For all treatments, belief threshold (BT) = 0.5.	147

Chapter 1

Introduction

The financial markets play a very important role in society. Perhaps the most prominent example of the financial markets are the stock markets. The financial markets can be defined as complex and dynamic socio-economic systems wherein a large number of heterogeneous market participants interact by essentially buying or selling assets of different types (Bodie et al., 2004; Graham, 1986). Market participants like retail (individual) and institutional investors, investment banks are heterogeneous in the sense that they have, for instance, different beliefs, goals, and investment or trading strategies. Additionally, they can be either a human, artificial (e.g., fully automated), or hybrid (e.g., semi-automated) participant (Treleaven et al., 2013). In the financial markets, market participants need to cope with uncertainty as well as with a variety of different types of risks, facing the inevitable risk-return trade-off. As a strategic environment, the market in itself can be thought of as a superaggregator of beliefs, and prices result from the interactions of all market participants. As a result, markets are generally considered more accurate than the beliefs of any market participant individually, in a sort of wisdom of the crowd.

The financial markets serve many functions. They help improving the capital allocation process and therefore contribute to economic growth (Frank, 2008; Wurgler, 2000). In addition, they generate a feedback effect from market prices to the real economy, affecting a series of variables, e.g., market prices affect consumer confidence that, in turn, affects consumption that ultimately affects market prices. The behavior of the financial markets is generally assumed to be an important gauge for helping the understanding of the current and future state of the whole economic and financial system.

Perhaps the importance of the financial markets might be better noticed when they do not fulfill their primary functions. More specifically, there are historical and empirical evidence about the significant and dramatic consequences of the occurrence of crashes in the financial markets.

As an example, a series of remarkable events threatened the stability of the world's financial system on September, 2008 (Brunnermeier, 2009; Farmer et al., 2012; Gorton, 2009; Ivashina and Scharfstein, 2010; Lewis, 2011; Melvin and Taylor, 2009). To quote but a few examples (Hall, 2010; Ohanian, 2010): the United States (US) Government stepped in mortgage giants Fannie Mae and Freddie Mac by promising to bail out bad loans; the financial services firm Lehman Brothers, a 158-year-old company, filed for Chapter 11 bankruptcy protection, holding more than \$600 billion in assets in the largest bankruptcy filing in US history; the AIG, one of the world's biggest insurers, was saved by an \$85 billion loan from the US Federal Reserve. These widespread failures resulted in a worldwide large-scale crisis of confidence. Market participants were unwilling to lend in the financial markets, unsure about the quality of the assets their counterparties were holding. This crisis of confidence turned rapidly into global financial panic, where uncertainty and fear reign in the financial markets. As a consequence, in an attempt to try to avoid the effects of similar shocks that occurred in the past, like the financial markets Crash of 1929 that lead later to the Great Depression (Bernanke, 1983), on October 2008 the US government approved the Troubled Asset Relief Program (TARP), an act granted the Secretary of the Treasury authority to either purchase or insure up to \$700 billion in troubled assets owned by financial institutions. Additionally, on the same month, central banks from Britain, US, EU, Canada and Switzerland took the coordinated action of cutting interest rates to try to restore confidence. The world's financial system was on the brink of disaster, coming close to a financial meltdown.

The Crash of 2008 had immediate, profound, vast and long-lasting consequences for the world economy (Reinhart and Rogoff, 2009b, 2014). To briefly illustrate the impact on the US economy, where the crisis have originated, the Gross Domestic Product (GDP) fell 1.97% in the fourth quarter of 2008, in the worst fall since 1947 and; from September, 2008 to June, 2009 (the end of the recession period at least in formal terms) the Unemployment Rate rose from 6.1% to 9.5%. It is not an overstatement to say that the consequences has somewhat been felt to the present time, even almost eight years after the unforgettable and remarkable events of 2008. Two months later on November,

2008 during a briefing by academics on the turmoil on the financial markets Britain's Queen Elizabeth asked: "Why did no one see it coming?"

The financial markets are indeed systems very difficult to be studied and truly understood. In this context, there are different perspectives and hypotheses that try to explain how the financial markets work, both on micro (individual) and on macro (global) perspective. On the one hand, traditional economic theories such as the Efficient Market Hypothesis (EMH) (Fama, 1970) tend to consider market participants are rational, that is to say that they are able to form expectations rationally, and to make optimal decisions, have stable and well-defined preferences. As a consequence, markets are efficient in the sense that market prices fully and instantaneously incorporate all the available information and expectations of market participants. However, behavioral economics (Kahneman and Tversky, 1979b; Shiller, 2003), a research field resulted from the combination of the psychology and economics disciplines, has been providing extensive experimental evidence that market participants decision-making under uncertainty has deviations from the so-called rational behavior, what is known as behavioral biases. Such deviations seem to be indissociable aspects of the human nature. Furthermore, market participants (human beings) have bounded rationality, meaning essentially that they do not always have enough time or the cognitive ability to process all the related information with accuracy.

In addition to the behavioral economics findings, different disciplines such as cognitive neuroscience and neuroeconomics (Camerer et al., 2005; Loewenstein et al., 2008; Tom et al., 2007) have been clarifying, particularly over the last decade, the relevance of emotions for the human reasoning and decision-making process. It is currently well-accepted that emotions such as happiness, unhappiness, and surprise have been generally considered the basis for a sort of learning system, and are essential to abilities and mechanisms usually associated with basic rational and intelligent behavior. Emotions also have influence on which and how memories are encoded, stored and also on how memories are later retrieved (Dalglish, 2004; Kensinger, 2009b). For instance, memories associated with either positive or negative emotions appear to gain privileged status in memory so that they tend to persist over years, decades and lifetime.

On the other hand, the Adaptive Market Hypothesis (AMH) (Lo, 2004, 2012), a relatively recent and still under development hypothesis, reconciles market efficiency with behavioral biases by applying the principles of evolution. The AMH acknowledges the

importance of emotions, the existence of behavioral biases, and the occurrence of interesting, though complicated, phenomena such as bubbles, cycles, panics, and trends. It represents a shift of thinking from the physics to biological sciences.

Despite of novel hypotheses and relevant findings, considering the nature and complexity of the financial markets together with the sophisticated and complex human decision-making mechanism (which is influenced by a myriad of intertwined factors), it can be claimed that we are all still far away from understanding the behavior of market participants individually and, as a result, from understanding the behavior of the financial markets globally. The failure of preventing, or at least predicting, the Crash of 2008 had once again exposed the allegedly flaws of traditional economic thinking. For instance, policymakers and regulators do not have sufficient tools to cope with the complexity of the financial markets with its hidden causalities rooted in sophisticated technological developments (Buchanan, 2009).

In the aftermath of the Crash of 2008 several transformations and improvements have been proposed, specifically in terms of how complex economic and financial systems can (should) be conceived and modeled. The crisis has been perceived by many as the appropriate moment to carefully investigate and explore novel approaches and ideas (Bouchaud, 2008; Caballero, 2010; Farmer and Foley, 2009; Gatti et al., 2010; Kirman, 2010; Shiller, 2013).

One of these novel approaches is given by the use of Agent-based Financial Markets, a subfield of a broader research area known as Agent-based Computational Economics (ACE) (Tesfatsion, 2002). Researchers in the context of Agent-based Financial Markets can depart from classical assumptions claimed by traditional economic theories and hypotheses, and then test new hypotheses, assumptions, concepts, and ideas. Furthermore, the agent-based models make possible the design, implementation, and realization of a variety of different, novel and realistic experiments. It offers the possibility of novel insights into the understanding of fundamental aspects of a wide range of different types of economic and financial systems, particularly the financial markets. From the agent-based perspective, market participants are represented by artificial agents, while the financial markets can be modeled from bottom up as well as be simulated as a multi-agent system composed of a large number of heterogeneous, adaptive artificial agents (Holland and Miller, 1991; Russell and Norvig, 2009; Wooldridge, 2002).

1.1 Goals

The failures of the traditional economic thinking together with the complexity of the human-decision making process have been stressing the difficulty of the problem. The challenge of better understanding the behavior of market participants individually and, ultimately, of truly understanding the financial markets as a whole requires the use of novel interdisciplinary approaches (e.g., (Blume et al., 2015; Challet, 2016; Lo, 2016a)). This work is in line with this context. We address the problem of understanding the financial markets from the cognitive science perspective. We consider that the use of relatively simple cognitive agents together with agent and multi-agent-based simulation, case studies, and time series analysis, constitute powerful tools, presenting a rich and novel set of possibilities for investigation, having the potential to shed a new light on the behavior of market participants (human agents) in economic and financial contexts, as detailed in the next sections.

The main goal of our exploratory work is to investigate which contributions the application of a cognitive modeling approach, where artificial agents are empowered with human-like characteristics, might bring to the understanding of the financial markets, specifically to explain and understand the behavior of market participants (human agents) in a micro (individual) perspective as well as the behavior of the financial markets in a macro (global) perspective.

We aim to empower artificial agents with mechanisms similar to or inspired in those used by humans so that ultimately we have cognitive agents. This empowerment includes, among other things, that an artificial agent must be equipped with the following characteristics and capabilities: different memory systems (e.g., short-term memory and long-term memory) as well as various memory processes (i.e., encoding, storage, and retrieval); the capacity of recognizing, simulating and expressing emotions (e.g., surprise, happiness, unhappiness); several adjustable decision-making processes; the ability to receive and process endogenous (internal) and exogenous (from the environment) information; and the ability to learn new knowledge, eventually reusing past knowledge.

It is important to notice that our goal is neither to discover relevant relations (e.g., causality) between economic and financial variables with the purpose of developing more profitable investment nor trading strategies nor we have the ambition of providing

an exhaustive, complete, and final explanation of the behavior of human agents in the financial markets.

1.2 Approach

In this work we are particularly influenced by the following three concepts which we consider relevant toward the understanding of how human agents behave individually in the financial markets and, ultimately, of how the financial markets work globally. First, we consider that emotions are inherent and important aspects of the human decision-making and reasoning processes as well as we acknowledge the fact that emotions have a significant impact on the memory processes of encoding, storage, and retrieval. Second, we consider that the findings from the behavioral economics research, particularly the behavioral biases, must be taken into account. Third, we consider that the Adaptive Markets Hypothesis (AMH) (Lo, 2012) is one of the best hypotheses to explain and understand how the financial markets work by reconciling market efficiency with behavioral biases, acknowledging the importance of emotions, the existence of behavioral biases, and the occurrence of interesting phenomena such as bubbles. These three concepts offer a framework for the design and realization of more behaviorally plausible and realistic investigations.

In this context, **one of the main important issues in this work is the construction of the cognitive model of individual market participants (human agents)**. To this end, we conceived the TribeCA (Trading and investing with behavioral-emotional Cognitive Agents) model, which builds on the following three major concepts: the cognitive emotion theory known as Belief-Desire Theory of Emotions (BDTE) (Reisenzein, 2009), the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997), and the artificial surprise model conceived by Macedo et al. (2004).

The first concept is the cognitive emotion theory known as Belief-Desire Theory of Emotions (BDTE) (Reisenzein, 2009). The BDTE consists of propositions, beliefs, desires, new beliefs, and two hard-wired comparator mechanisms, namely the Belief-Belief Comparator (BBC) and the Belief-Desire Comparator (BDC). The conceptual framework of the BDTE is the same as the belief-desire theory of action which inspired the BDI (Belief-Desire-Intention) (Bratman et al., 1988; Rao and Georgeff, 1995, 1998)

approach to artificial agents. BDTE defines emotions as products or signals produced by the BBC and BDC.

The second concept is the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997). They claim surprise-eliciting events elicit a four-step sequence of processes. The first step is the appraisal of an event as unexpected or schema-discrepant. In the second step, surprise is experienced, ongoing mental processes are interrupted, and resources are reallocated towards the unexpected event. The third step is the analysis and evaluation of the unexpected event. Finally, the fourth step is the schema update. However, neither the BDTE nor the model of surprise proposed by Meyer et al. (1997) detail how surprise can be computed.

The third concept is the artificial surprise model proposed by Macedo et al. (2004). They carried out an empirical study with humans with the goal of investigating how to compute the intensity of surprise in an artificial agent. They proposed several alternative functions for computing the surprise intensity based on the assumption that the surprise “felt” by an agent elicited by an event E_g is proportional to the degree of unexpectedness of the event E_g . Their study suggests that the intensity of surprise about an event E_g , from a set of mutually exclusive events E_1, E_2, \dots, E_m , is a nonlinear function of the difference, or contrast, between its probability/belief and the probability/belief of the highest expected event (E_h) in the set of mutually exclusive events E_1, E_2, \dots, E_m .

With the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997) as well as with the artificial surprise model proposed by Macedo et al. (2004), core components of the TribeCA, we carried out two case studies on comparing the traditional (economics and finance) perspective with the cognitive science perspective on modeling surprise in economics and finance. The first case study is in the context of the so-called market consensus (Baccan et al., 2014b), while the second is in the context of risk management (Baccan et al., 2015).

We subsequently integrated the implementation of the proposed cognitive model of individual human agents (TribeCA) into two sophisticated tools known as Java Auction Simulator API (JASA) (Phelps, 2007) and Java Agent-Based Modeling (JABM) toolkit (Phelps, 2012) used in the field of Agent-based Financial Markets. We carried out three agent and multi-agent-based simulation experiments (Baccan and Macedo, 2013b) with the proposed cognitive model in the El Farol bar problem, revisiting some fundamental

issues regarding the understanding of the financial markets such as efficiency (Baccan and Macedo, 2013a) and rationality (Baccan et al., 2014a). The El Farol bar problem is a relatively simple but powerful model that can be used as a good starting point for modeling and understanding the financial markets both individually and globally.

1.3 Contributions

The main contribution of our work is the cognitive model of human agents developed and applied to the context of economics and finance, particularly to the financial markets. This encompassed several steps, each giving rise to an important contribution. Thus, we can list the following contributions of our work:

- An interdisciplinary work that is in line with a broader context that has been stressing the need of new, particularly multidisciplinary, approaches for the study of economic and financial complex systems, such as the financial markets, in the light of the Crash of 2008 as well as empirical evidence.
- A generic, novel cognitive model of human agents (TribeCA) based on the following three pillars: the cognitive emotion theory known as the Belief-Desire Theory of Emotions (BDTE), the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997), and the artificial surprise model proposed by Macedo et al. (2004).
- The presentation of two case studies that demonstrate how the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997) and the artificial surprise model proposed by Macedo et al. (2004), two core components of the TribeCA, can be applied to compute artificial surprise in economic and financial contexts.
- A generic implementation of the proposed cognitive model of human agents (TribeCA) and integration of this implementation into two sophisticated tools used in the field of Agent-based Financial Markets known as Java Auction Simulator API (JASA) and Java Agent-Based Modelling (JABM), allowing the realization of a wide variety of different and new experiments in the context of economics and finance, taking into account the cognitive science perspective.

- Considerations and conclusions, some of them non trivial and counter intuitive, resulted from the realization of three different multi-agent-based experiments with our artificial cognitive agents in the El Farol bar problem (Baccan and Macedo, 2013b) regarding some fundamental questions to the understanding of the financial markets such as efficiency (Baccan and Macedo, 2013a) and rationality (Baccan et al., 2014a).

We would like to add that part of the work presented in this thesis has also been published in several scientific papers, presented and/or published in national and international venues, such as MABS (Multi-Agent-Based Simulation) workshop at AAMAS (Autonomous Agents and Multi-Agent Systems), STAIRS (Starting AI Researcher Symposium) at ECAI (European Conference on Artificial Intelligence), EPIA (Portuguese Conference on Artificial Intelligence), and SMC (IEEE International Conference on Systems, Man, and Cybernetics) as well as in international workshops, such as WEHIA (Annual Workshop on the Economic Science with Heterogeneous Interacting Agents), and CCFEA (Centre for Computational Finance and Economic Agents) Workshop.

1.4 Thesis structure

In this first chapter we introduced the problem addressed and the main contributions of the thesis. The rest of this thesis is organized as follows.

Chapter 2 presents background on the Computational Cognitive Modeling, Financial Markets, and Agent-based Financial Markets, three major areas related to this work. Some of the addressed topics include computational cognitive modeling and computation of emotions, different financial market hypotheses ranging from traditional economic theories to recent and novel theories, and agent-based modeling and simulation of the financial markets. We conclude this Chapter by presenting a brief survey of the most representative related work, stressing the main differences and similarities to our work.

Chapter 3 presents our conceptual cognitive model (TribeCA), the agent architecture and its main modules, namely Sensors, Reasoning, Goals, Memory, and Actions. Additionally, we present a cognitive modeling approach to the El Farol bar problem

using the TribeCA and an implementation of our proposal that make it possible the realization of the agent-based experimental evaluation with artificial cognitive agents.

Chapter 4 presents two different perspectives on computing surprise in economics and finance, namely the traditional (economic and financial) perspective and the cognitive perspective. To this end, we carried out two case studies. In the first case study we compare the computation of surprise from the perspective of economics and finance to the perspective of cognitive science in respect to the idea of market consensus. In the second case study we compute the cognitive surprise “felt” by an artificial agent relying on a popular risk management tool known as Value-at-Risk (VaR) historical under two different market scenarios, specifically a Crash period and a Calm period.

Chapter 5 presents the agent-based experimental evaluation we carried out with the implementation of our conceptual cognitive model in the context of the El Farol bar problem, an agent-based model for studying the financial markets. In the first experiment, we present the initial exploration to observe how the artificial cognitive agents we modeled behave. In the second experiment, we focus on the impact of rationality on market efficiency. In the third experiment, we investigate the impact of the memory size of the artificial cognitive agents on market efficiency.

Chapter 6 concludes the thesis and proposes topics for future research directions by pointing out and detailing some possible future works.

Chapter 2

Background Concepts

This thesis is essentially build on ideas, concepts, definitions, and assumptions provided by the following three major areas: Computational Cognitive Modeling, Financial Markets, and Agent-based Financial Markets.

First, we present in Section 2.1 the major concepts on computational cognitive modeling, in particular, how emotions (including surprise) can be defined as well as qualitatively and quantitatively computed in the context of cognitive emotion theories. We give particular attention to the surprise emotion by presenting the surprise process proposed by Meyer and colleagues and the artificial surprise model proposed by Macedo and colleagues. We then illustrate how these concepts may be applied to cognitive modeling the coin tossing game, a familiar and well-known game.

Second, we overview in Section 2.2 the general concepts regarding the financial markets, and present some traditional and non-traditional hypotheses that try to explain how the financial markets work. We also present some findings from behavioral economics research as well as survey on how the term surprise has been used in economics and finance, particularly in the context of the financial markets.

Third, we present in Section 2.3 a relatively recent field of research known as Agent-based Financial Markets that make it possible the design, implementation, and execution of different and novel experiments with artificial agent and multi-agent-based simulations in the context of economics and finance, especially the financial markets.

Finally, in Section 2.4 we survey the most representative related work, make a comparison of our work to those representative studies by stressing the main differences and similarities.

2.1 Computational Cognitive Modeling

In this section we start by presenting some definitions of what emotions are and how they can be represented. We then illustrate the importance of emotions by presenting how emotions influence the human decision-making process as well as the relationship between emotions and memories, especially how emotions influence the processes of encoding, storage and retrieval of memories. Subsequently, we present a cognitive emotion theory named Belief-Desire Theory of Emotions (BDTE). Finally, we focus on research regarding the surprise emotion, present the surprise process proposed by Meyer and colleagues, and conclude with the presentation and discussion of the artificial model of surprise proposed by Macedo and colleagues.

2.1.1 Emotions: definition, representation and functions

Firstly, it is important to take into account that **there is no universal definition of what emotions are** (for examples of different forms of defining emotions see (Boehner et al., 2007; Frijda et al., 2000; James, 1884; Reisenzein, 2007; Scherer, 2005)). Additionally, emotions are not a human being peculiarity, actually it is believed that other animals also possess emotions. For example, Darwin (1872) pointed out resemblances between the expressions of the emotions in humans and animals.

Three concepts usually related to emotions are emotional state, emotional expression, and emotional experience (Mauss and Robinson, 2009). **Emotional states** consists in the mental and physical state of the one who felt the emotion. It means that emotions essentially take place within and therefore cannot be directly observed. However, emotions can be **expressed** voluntarily or involuntary to others through mechanisms such as facial expressions (Ekman and Friesen, 1971) as well as by means of physiological changes that generally modify the physical state such as a growth in breathing, sweating, and heart rate. Based on the signals provided, others try to infer which emotions one felt. Last but not least, **emotional experience** refers to everything that the one

who felt the emotion consciously perceives of the felt emotional state and is believed to be a human being peculiarity (Damasio, 1999).

The lack of agreement about what exactly emotions are, results in a different number of emotions attributed to humans. For instance, resulting from the work on how emotions are expressed through facial muscular patterns, Ekman and colleagues (Ekman, 1992; Ekman and Rosenberg, 1998) claims that humans possess six “basic emotions” namely anger, disgust, fear, happiness, sadness, and surprise. Others researchers (e.g., see (Lazarus, 1991; Ortony and Turner, 1990)), and even Ekman in a later review (Ekman, 1999), present a more sophisticated list of emotions in addition to the six “basic” ones.

Representing emotions

Regardless of how many emotions we have, there is considerably acceptance on representing the emotional state in terms of two orthogonal dimensions namely **valence and arousal** (Mauss and Robinson, 2009; Russell, 1980; Russell and Barrett, 1999; Russell et al., 1989). The **valence** dimension contrasts states of pleasure (positive) with states of displeasure (negative) with neutral often considered an intermediate value. The **arousal** dimension contrasts states of low arousal (calm) with states of high arousal (excitement). We illustrate in Figure 2.1 how these two orthogonal dimensions accommodate different emotions. As an example of how the arousal and valence dimensions accommodate different emotions, consider the case of happiness, sadness, and surprise. Happiness has positive valence, while sadness has negative valence. In both cases the level of arousal varies depending on how desirable or undesirable the consequences of the event that elicited the emotion are. Surprise, in turn, has neither positive nor negative valence (it has neutral valence), what actually is positive or negative is the event that elicits it. Surprise has also a high level of arousal.




FIGURE 2.1: Schematic representation of the valence and arousal model (extracted from Russell and Barrett (1999)).

Some functions of emotions: emotions and decision-making, and emotions and memories

Emotions serve humans in **many functions**. From an evolutionary perspective (Darwin, 1859; Plutchik, 1980, 1982) it is believed that emotions have been of fundamental importance by contributing to maintain our physical well-being, and essentially helping us to preserve the survival (Lazarus, 1991). For example, consider the function of surprise, fear, and happiness. When one is confronted with a stimulus that is evaluated as representing a threat to survival, the first emotion that may be elicited is probably surprise so that her/his attention is immediately shifted to that stimulus. Subsequently, anger may be elicited and her/his capabilities are temporarily enhanced, preparing her/him to execute a possible set of actions in response to that threat, a kind of response known as “fight-or-flight”. Finally, if the threat is successfully eliminated, she/he can feel happiness as a result of the accomplishment of a goal or fulfillment of task.

Emotions and decision-making

Different disciplines such as cognitive neuroscience and neuroeconomics (Brand et al., 2006; Camerer et al., 2005; Fehr and Rangel, 2011; Gigerenzer and Gaissmaier, 2011; Glimcher and Fehr, 2013; Krajbich et al., 2014; Loewenstein et al., 2008; Mellers et al., 1998; Ochsner and Gross, 2005; Sanfey et al., 2006; Sherry and Schacter, 1987; Tom et al., 2007) have been clarifying the **relevance of emotions for the human reasoning and decision-making process**. It is currently well-accepted that, in opposition to the popular idea, emotions are essential to abilities and mechanisms we tend to associate with basic rational and intelligent behavior. Emotions have been generally considered the **basis for a sort of learning mechanism known as the reward-and-punishment system** (e.g., (Peterson, 2007; Quartz, 2009)). We constantly receive feedback of different modalities resulting from our interaction both with other humans and with the environment. Interactions that elicit positive feedback or reward (e.g., pleasure or happiness) tend to be repeated in the future, whereas interactions that elicit negative feedback or punishment (e.g., pain or sadness) tend to be avoided. Additionally, the reward-and-punishment system provides humans with a sort of “cost-benefit analysis” (Lo, 2004) which contributes to the selection of advantageous behavior when we are confronted with a set of possible actions. Human decision-making is believed to depend on different areas of the brain related to emotions (e.g., amygdala, prefrontal cortex) and memory (e.g., hippocampus). Such different areas provide humans with the neural and structural basis so that emotions can effectively serve in a sort of a reward-and-punishment system of learning.

An influential conceptualization of how emotions affect decision-making is the idea of the **Somatic Marker Hypothesis (SMH)** provided by Damásio, Bechara and colleagues (Bechara and Damasio, 2005; Bechara and Reimann, 2010; Damasio, 1994, 1996; Damasio et al., 1991). The SMH originated as a response to observations that patients with brain damage, especially in a region known as prefrontal cortex, demonstrated severe impairments in personal and social decision-making. For instance, patients make financial decisions that generally lead them to financial losses. The SMH states that “marker” signals, which are somatic in the sense that they are related to body-state structure and regulation, i.e., homeostasis, help to make advantageous decisions in real-life situations. When one needs to make decisions under circumstances of complexity and uncertainty, and whose outcomes could be either positive or negative, somatic signals (e.g., heart rate) are used to mark the whole experience, including both

the chosen decision and the outcome. When a situation that is similar to a previously experienced situation occurs, information concerning the possible courses of action and their likely outcomes are “reactivated” so that it serves as a signal that tells she/he to either accept or reject some particular action.

Empirical studies of the SMH generally involve the use of a specific task known as the **Iowa Gambling Task (IGT)** (Bechara, 2004). The IGT was designed to simulate real life scenarios by confronting participants with situations in which they must make decisions under conditions involving reward, punishment, and uncertainty. Participants need to select one from four card decks namely A, B, C, and D. Two decks, A and B, yield high gains and high losses, whereas the other two decks, C and D, yield small gains and small losses. It has been demonstrated that, in accordance with the SMH, damage in the prefrontal cortex impairs decision-making. Assuming that the task is continuously applied to participants, on the one hand, healthy participants tend to realize and correctly learn that card decks A and B are disadvantageous in the long run, in contrast to card decks C and D that are advantageous. On the other hand, patients with brain damage tend to continuously select more cards from decks A and B. Regardless of some critics about the SMH (e.g., see (Dunn et al., 2006; Maia and McClelland, 2004, 2005)), the IGT is a well-accepted and valid way of measuring decision-making under uncertainty. The IGT has been widely used in a range of studies with patient groups such as amnesic patients (Gupta et al., 2009) and pathological gambling patients (Cavedini et al., 2002).

Similar to the prefrontal cortex, the amygdala also plays an important role in decision-making (e.g., (De Martino et al., 2010; Gupta et al., 2011; Loewenstein et al., 2001; Phelps, 2006; Werner et al., 2009)). The most prominent function of the amygdala consists in affective processing, by detecting uncertainty in the environment and in triggering arousal and vigilance. Specifically, the amygdala has been recognized as an important component of the system involved in the acquisition, storage, and expression of fear (Clark et al., 2014; Davis and Whalen, 2001; LeDoux, 2000; Whalen, 2007). It has been extensively demonstrated that patients with amygdala damage do not have the ability to generate autonomic responses to reward and punishment, and consequently, cannot utilize “somatic marker” type cues to help and guide them in similar future decision-making situations.

Emotions and Memories

Human memory is a highly **sophisticated, complex and dynamic** mechanism that is composed of different memory systems and processes that is used for a great number of functions. Here we provide a brief and concise overview of the human memory, though sufficient in the context of this work (for an extensive review see (Baddeley et al., 2009; Cohen and Conway, 2008; Collins and Loftus, 1975; Eysenck, 2001; Kahana et al., 2008; Phelps, 2006; Squire, 2004, 2009)). Human beings have essentially three main types of memory systems namely sensory memory, short-term memory and long-term memory (e.g., (Cowan, 2008b)).

Sensory memory (Cowan, 2008a) is the first stage of memory. It acts as a buffer for all stimuli received through the senses, most of them received essentially by vision and hearing, retaining a copy (snapshot) of what is seen and/or heard for just very few seconds (maybe milliseconds). Each sensory channel has its own sensory memory. Attention mechanisms operate in sensory memory in order to emphasize what must be transferred to short-term memory for a further examination.

Short-term memory (also known as working memory) is an intermediary stage between the sensory memory and the long-term memory. It provides a system for temporarily storing and recalling information under process (e.g., (Cowan, 2001)). The short-term memory differs from the sensory memory once the information in short-term memory is already encoded as well as it does not fade away so quickly.

Information is transferred from short-term memory to **long-term memory** by a continuous exposure of such information in a mechanism of reinforcement or consolidation. Long-term memory allows information to be “permanently” stored, managed (e.g., “deleted” or “forgotten”), and recalled for later use. It is a dynamic system in which new information is associated with the existing ones. In addition, long-term memory is a complex mechanism generally assumed to be divided into two types namely nondeclarative (procedural) memory, and declarative memory.

Nondeclarative memory (also known as procedural memory) is the long-term memory of skills that an organism possesses and procedures which the organism is able to perform. It consists in a sort of “how to” do things, a sort of pairs of action/reaction or stimulus/response. Riding a bike is a typical example of procedural memory. Procedural memories are unconscious, generally difficult to verbalize, and very durable

over a long period of time, even in elders with some kinds of brain disease such as the Alzheimer's disease.

Declarative memory is the memory for storing knowledge/facts which are consciously available. It is divided into two types namely episodic memory, and semantic memory (e.g., (Squire, 2004, 2009)). There are fundamental differences between episodic memory and semantic memory especially in terms of the type of information represented, type of organization in memory, source of information, and focus (Binder and Desai, 2011). **Episodic memories** represent information referring to specific events, objects, and people (Conway, 2009; Tulving, 1972, 1983, 2002). They are chronologically or spatially organized in memory, focus on subjective reality namely the self, and the main source of information comes from personal experience. **Semantic memory** refers to generalized knowledge about the external world such as facts, meanings, and understandings. It is assumed to be represented in memory through knowledge structures known as schemas (Bartlett, 1932). A schema is a well-integrated chunk of knowledge or sets of beliefs related to actions and their consequences (Schank and Abelson, 1977), events, objects (Minsky, 1974), and situations. It represents all kinds of generic knowledge varying from simple to complex knowledge. The knowledge available in semantic memory derives from the episodic memory. Indeed, the interaction both with the environment and others might result in similar episodes. As those episodes accumulate, over time specific features (micro-details) of the episodic memory are generally lost, resulting in a fragment of generalized, semantic memory.

Different memory systems and memory processes serve species in many functions. Human beings rely on memory to retrieve previous knowledge in order to make plans and solve problems in the present, and to predict future events, by guiding thought and behavior (Beaman et al., 2007; Bluck and Alea, 2009; Gupta et al., 2009). Additionally, memory is important in supporting the development of the self and identify (Conway, 2005; Conway and Pleydell-Pearce, 2000) as well as helps humans in developing, maintaining, and strengthening social relationships, facilitating social interaction in general (Baron and Bluck, 2009; Kihlstrom, 2009; Thomsen, 2009).

Human memory systems, particularly the memory processes, seem to be significantly influenced by emotions (James, 1890). It has been suggested that emotions have influence on which and how memories are encoded, stored (including forgetting (Hardt et al., 2013; Schacter, 2002; Wixted, 2004)), and also on how memories are later retrieved (e.g., (Thomson and Tulving, 1970; Tulving and Osler, 1968;

Tulving and Thomson, 1973)). Memories associated with either positive or negative emotions appear to gain privileged status in memory so that they tend to persist over years, decades and lifetime (Buchanan, 2007; Dalgleish, 2004; Dalgleish and Power, 1999; Kensinger, 2009a,b; Kensinger et al., 2005; Levine et al., 2009; McGaugh, 2015; Yonelinas and Ritchey, 2015). Additionally, surprise seems to play an important role as memories in how are encoded and later retrieved (e.g., (Axmacher et al., 2010)). As an example, consider the particular case of **flashbulb memories (FBM)** (Brown and Kulik, 1977). FBM refer to detailed, long-lasting and vivid memories of the circumstances in which one first learned of a very surprising, consequential and emotionally arousing event. Details of the circumstances include not only information about the event in itself but also contextual information such as who told one about the event, where one was, and what was one doing at that time, and so forth. FBM generally occurs with public events such as the terrorist attack of 11 September 2001 (Conway et al., 2009) (for a detailed list of events used in FBM studies see (Holland and Kensinger, 2010)).

Further evidence of the relationship between emotions and memories come from the study of context-dependent memory, i.e., the finding that memory benefits when the spatio-temporal, physiological, mood, or cognitive context at retrieval matches that present at encoding. There are at least four types of context-dependent memory namely state-dependent memory, mood-congruent and mood-dependent memory, environmental context-dependent memory, and cognitive context-dependent memory (Baddeley et al., 2009). We consider that the state-dependent memory, mood-congruent and mood-dependent memory are more related to our work so that we briefly describe them below.

State-dependent memory refers to the finding that information encoded in one particular physiological state tends to be better remembered in the same state. Experiments generally involve the induction of some physiological state (e.g., excited, calm) and then the encoding of information which needs to be later retrieved either on the same or another physiological state (Eich, 1980). For instance, Miles and Hardman (1998) examined how aerobic exercise might produce state-dependent memory. Participants learned a list of words in two different physiological states namely at rest and while exercising aerobically on an ergometric bicycle. Afterwards, participants were asked to recall the words in either the same or alternative state (i.e., rest or exercising) in which the list of words was encoded. In general, participants whose encoded

and recalled states were the same had a better performance. Cavenett and Nixon (2006) carried out a similar experiment in the context of skydiving. They demonstrated preferential encoding of relevant information rather than irrelevant information once skydivers recalled significantly less irrelevant material than delayed controls. Results of similar experiments carried out with other physiological state inductors such as alcohol (Goodwin et al., 1969; White, 2003), caffeine (Kelemen and Creeley, 2003), cigars (Peters and McGee, 1982) and marijuana (Eich, 1980), have also demonstrated the state-dependent memory effect.

Mood-dependent memory is a particular form of the context-dependency effect by which information encoded in one particular mood whether positive, neutral, or negative tends to be better remembered in the same mood (Bower, 1981; Eich, 1995; Eich and Macaulay, 2000; Lewis and Critchley, 2003). Similarly, the mood-congruent memory effect states that the current mood makes memories that match the same emotional material more easily available (Mayer et al., 1995). Therefore, negative mood makes negative memories more easily available than positive memories, and vice versa. However, in contrast to mood-dependent memory, mood-congruent memory does not affect the recall of neutral memories (Baddeley et al., 2009; Lewis et al., 2005). Perhaps one of the best evidence of the mood-congruent effect is the observation of patients diagnosed with depression. When one is depressed memories associated with negative emotions are more easily available for retrieving so that the remembered memory become a sort of trigger to remember another memory associated with a negative emotion, and so on.

2.1.2 Cognitive Emotion Theories

Cognitive emotion theories (Frijda, 1986; Frijda et al., 2000; Lazarus, 1991; Ortony et al., 1998; Reisenzein, 2009; Scherer et al., 2001) rely on the assumption that emotions are mental states elicited as a result of the evaluation or appraisal of stimuli of all kinds (e.g., actions, events, interactions, sensations) and can be computed in terms of cognitions (beliefs) and motives (desires). Beliefs are mental states in which one holds a particular proposition to be true, whereas desires represent the motives or future states that one wants to accomplish.

Belief-Desire Theory of Emotions (BDTE)

The Belief-Desire Theory of Emotions (BDTE) (Reisenzein, 2009) is a cognitive emotion theory that assumes that emotions are the product of beliefs and desires. The conceptual framework of the BDTE is the same as the belief-desire theory of action which inspired the Belief-Desire-Intention (BDI) (Bratman et al., 1988; Rao and Georgeff, 1995, 1998) approach to artificial agents. BDI describes essentially autonomous agents in terms of beliefs, desires (goals), and intentions (the goals that the rational agent has committed to) (Bratman, 1987). Additionally, the central assumption incorporated in BDI agents is that actions are generated by a process of practical reasoning that comprises in the first step, the selection of a set of desires (goals), and in the second step, the determination of how these goals can be achieved by means of the available actions or plans (Macedo, 2006; Reisenzein et al., 2013). BDI has been one of the most frequently used software architectures for autonomous intelligent agents (Reisenzein et al., 2013). BDTE consists of propositions, beliefs, desires, new beliefs, and two hard-wired comparator mechanisms namely Belief-Belief Comparator (BBC) and Belief-Desire Comparator (BDC).

A **proposition** p is represented as a tuple $\langle S, B, D \rangle$ where S is the mental language expressing the proposition p , B and D are quantities representing, respectively, the agent's degree of belief and desire regarding proposition p .

The **strength of a belief in a proposition** p at time t , defined as $belief(p, t)$, is a value $\in [0.0, 1.0]$, where 1.0 denotes certainty that p , 0.5 maximal uncertainty, and 0.0 certainty that not p .

Similarly, the **strength of a desire about a proposition** p at time t , defined as $desire(p, t)$, might be a value, for instance, $\in [-100, +100]$. Positive values denote desire in favor of p , negative values denote desire against p , and 0 denotes indifference.

A **new belief is the belief or fact** that agents receive basically through its sensors. It is defined as a tuple $\langle S, B, * \rangle$, where S means the mental language expressing the proposition p , B is the belief about the proposition p , and $*$ denotes that the desire is irrelevant for new beliefs.

The **Belief-Belief Comparator (BBC)** compares each newly acquired belief with all pre-existing beliefs, looking for match versus mismatch. A match means that a pre-existing belief was confirmed by the newly acquired belief, whereas a mismatch

means that a pre-existing belief was disconfirmed. As a result, BBC yields either a belief-confirmation signal or belief-disconfirmation signal.

Similarly, the **Belief-Desire Comparator (BDC)** compares each newly acquired belief with all pre-existing desires, looking for match versus mismatch. A match means that a desire was “fulfilled”, whereas a mismatch means that a desire was “frustrated”. As a result, BDC yields either a desire-fulfillment signal or desire-frustration signal.

Figure 2.2 shows a schematic representation of the BBC and BDC comparators. BDTE defines **emotions as products or signals produced by the BBC and BDC**. The computation of non-neutral emotions (e.g., happiness, sadness) depends on the direction of the desire (i.e., positive or negative) of agents regarding p , whereas the computation of neutral emotions (e.g., surprise) depends only on the belief the agent has in p prior to receiving a new belief corresponding to p .

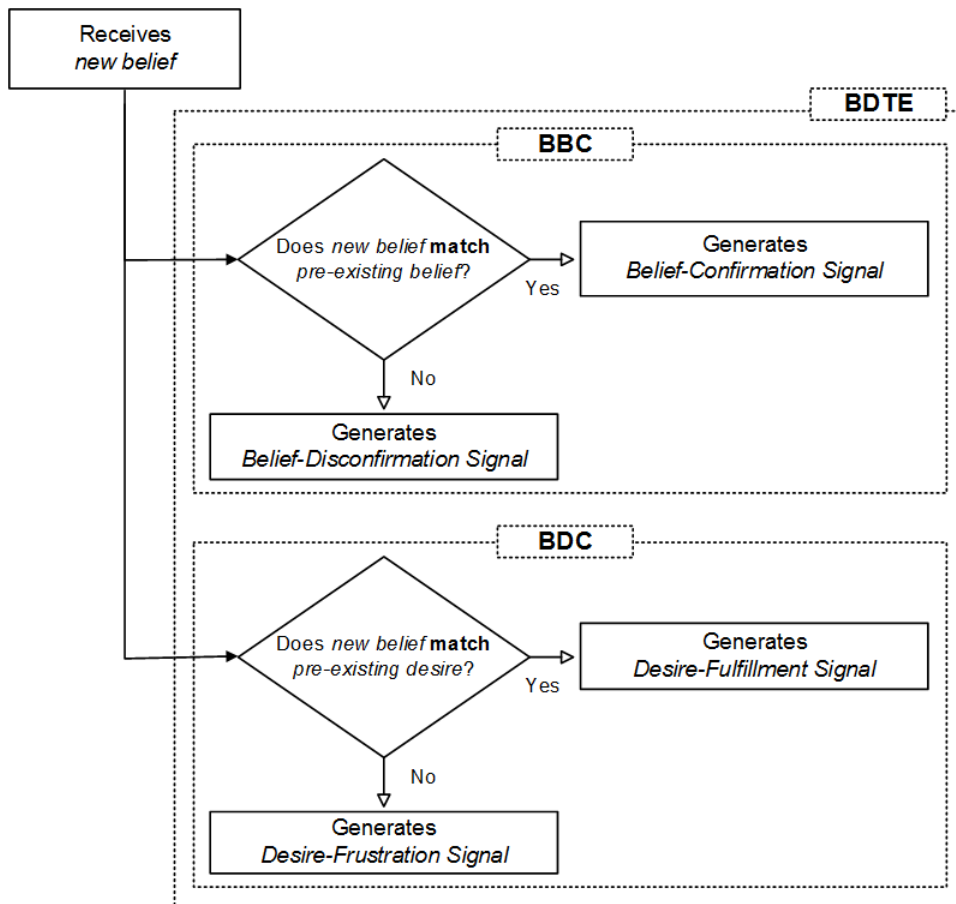


FIGURE 2.2: Schematic representation of the BBC (Belief-Belief Comparator) and the BDC (Belief-Desire Comparator) (adapted from Reisenzein (2009)).

TABLE 2.1: Belief-desire theory of emotions, qualitative formulation (adapted from (Reisenzein, 2009)). Notation: $Bel(p, t)$ believes p at time t ; $Certain(p, t)$ firmly believes p at time t ; $Uncertain(p, t) \equiv Bel(p, t) \wedge \neg Certain(p, t) \wedge \neg Certain(\neg p, t)$; $Desire(p, t)$ desires p at time t ; $Desire(\neg p, t)$ desires not- p at t (\approx is aversive against p at time t).

Emotion	if	Belief at t	Desire at t	Belief at $t-1$
$happy(p, t)$		$Certain(p, t)$	$Desire(p, t)$	
$unhappy(p, t)$		$Certain(p, t)$	$Desire(\neg p, t)$	
$hopes(p, t)$		$Uncertain(p, t)$	$Desire(p, t)$	
$fears(p, t)$		$Uncertain(p, t)$	$Desire(\neg p, t)$	
$surprised(p, t)$		$Certain(p, t)$	- (irrelevant)	$Bel(\neg p, t-1)$

TABLE 2.2: Belief-desire theory of emotions, quantitative formulation (adapted from (Reisenzein, 2009)). Notation: $b(p, t)$ represents the strength of belief in p at time t , with 1.0 denoting certainty that p , 0.5 maximal uncertainty, and 0.0 certainty that not- p ; $d(p, t)$ represents the direction and strength of the desire for p at time t , with values > 0 denoting positive desire in favor of p , 0 indifference, and values < 0 aversion against p .

Emotion Intensity	= function of d and/or b	for domain subset (else emotion intensity = 0)
$Happiness(p, t)$	$= \phi_{ha} [d(p, t)]$	$b(p, t) = 1 \wedge d(p, t) > 0$
$Unhappiness(p, t)$	$= \phi_{uh} [d(p, t)]$	$b(p, t) = 1 \wedge d(p, t) < 0$
$Hope(p, t)$	$= \phi_{ho} [b(p, t) \times d(p, t)]$	$0 < b(p, t) < 1 \wedge d(p, t) > 0$
$Fear(p, t)$	$= \phi_{fe} [b(p, t) \times d(p, t)]$	$0 < b(p, t) < 1 \wedge d(p, t) < 0$
$Surprise(p, t)$	$= \phi_{su} [b(p, t-1)]$	$b(p, t) = 0 \wedge b(p, t-1) > 0$

We present in Tables 2.1 and 2.2 the qualitative and quantitative formulations, respectively, of some emotions relevant to our work.

2.1.3 Surprise: definition, functions and process

In the complex and dynamic environment that surround us, humans are continuously, either voluntarily or involuntarily, receiving a vast amount of different types of information. A significant part of this information became available through vision and hearing. Such incoming information often contradicts previous beliefs or expectations. For example, during the financial crisis of 2008 (e.g., (Geithner, 2012; Gorton, 2009, 2008)), for the **surprise and shock of possibly the majority of market participants**, financial companies that many believed to be in good financial health have

ended up having to be rescued by governments (i.e., “too big to fail” companies like the AIG (Sorkin, 2010)) while others unexpectedly have failed (i.e., the bankruptcy of the investment bank Lehman Brothers¹). The experience of living in such an environment has provided humans with, at least in part, a general and intuitive idea of what surprise is. For instance, verbal expressions like “... caught by surprise ...” or “... what a good surprise” are generally used to manifest surprise about something.

From the cognitive science perspective (e.g., (Miller, 2003)), there is no consensus on whether surprise has either no valence or negative valence. For instance, Miceli and Castelfranchi (2015), have recently claimed that rather than being neutral, surprise is negatively valenced, since it represents an interference and, according to their view, interference is in general unpleasant and in conflict with the desire for predictability and structure (Noordewier and Breugelmans, 2013). Nevertheless, we follow (Macedo et al., 2009; Meyer et al., 1997; Reisenzein, 2009) by considering that surprise can be thought of as a neutral valence emotion, what actually is either positive or negative is the resulting event that elicited surprise. As a result, **surprise can be defined as the peculiar state of mind, usually of brief duration, caused by unexpected events, or proximally the detection of a contradiction or conflict between newly acquired and pre-existing beliefs**. Beliefs are assumed to be stored in our semantic memory, presented in Section 2.1.1. Semantic memory, our general knowledge and concepts about the world, is accepted to be represented in memory through knowledge structures known as schemas, which represent sets of beliefs related to actions and their consequences, events, objects, and situations. From the perspective of the BDTE, presented in Section 2.1.2, surprise can be thought of as mismatch or a belief-disconfirmation signal yielded by the BBC (Belief-Belief Comparator) as a result of the comparison between a newly acquired belief and a pre-existing belief or, in other words, the detection of a schema-discrepancy.

Surprise has many functions. It has been suggested that, among other functions, surprise is related to adaptation, attention (Ortony and Partridge, 1987; Ranganath and Rainer, 2003), creativity (Boden, 1995), and learning (Wasserman and Castro, 2005), playing an important role with respect to survival in a rapidly changing environment.

¹<http://www.nytimes.com/2008/09/15/business/15lehman.html> [Accessed: July 2017]

Surprise process

One of the most accepted models of surprise is the cognitive-psychoevolutionary model proposed by Meyer et al. (1997). Meyer and colleagues proposed a cognitive-psychoevolutionary model of surprise. They claim **surprise-eliciting events elicit a four-step sequence of processes**. Figure 2.3 shows a schematic representation of this process.

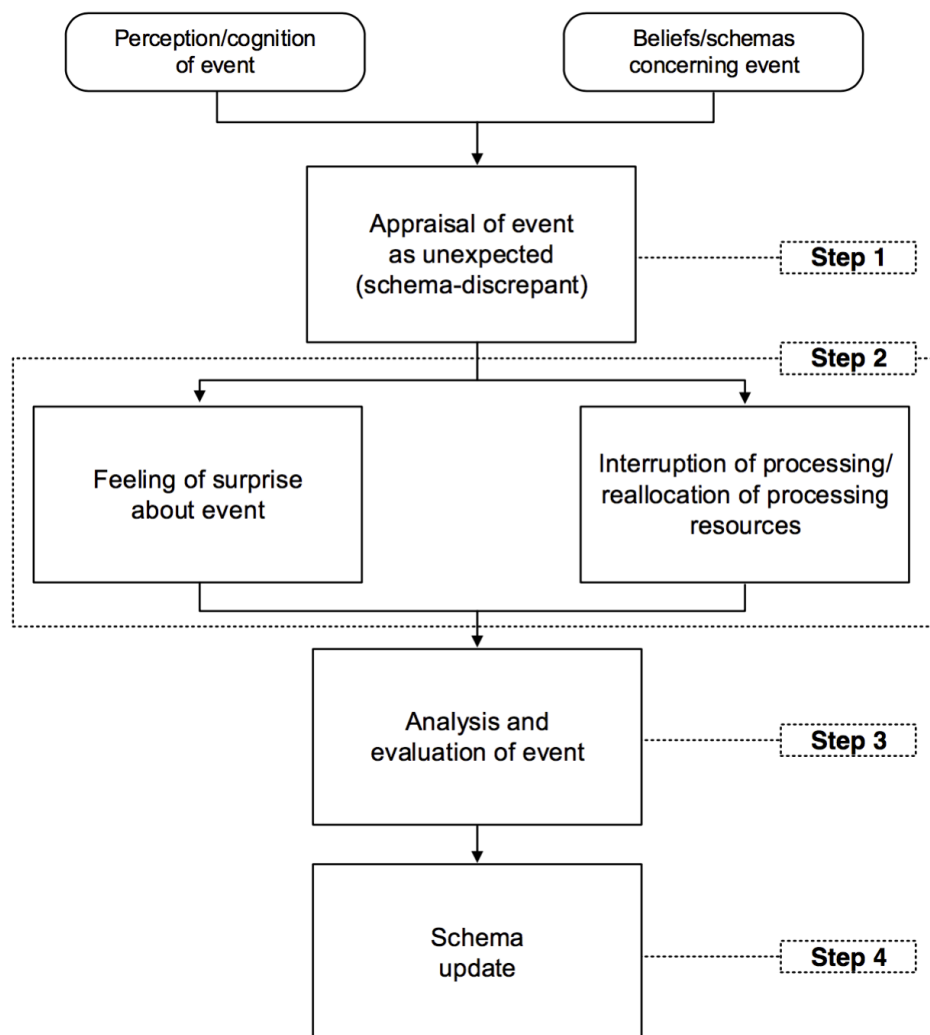


FIGURE 2.3: Schematic representation of the cognitive-evolutionary model of surprise process proposed by Meyer and colleagues (adapted from Macedo et al. (2009)).

The **first step** is the appraisal of an event as unexpected or schema-discrepant (Schutzwohl, 1998; Schutzwohl and Reisenzein, 1999). For instance, this appraisal is carried

out by the Belief-Belief Comparator (BBC) in the case of the Belief-Desire Theory of Emotions (BDTE), as mentioned earlier in this Section.

If the degree of unexpectedness or schema-discrepancy exceeds a certain threshold then, in the **second step**, surprise is experienced, ongoing mental processes are interrupted, and resources such as attention are reallocated towards the unexpected event.

The **third step** is the analysis and evaluation of the unexpected event. This step generally includes a set of subprocesses namely the verification of the schema discrepancy (e.g., is this information correct?), the analysis of the causes of the unexpected event (e.g., why did it happen?), the evaluation of the unexpected event's significance for well-being (e.g., is the event good or bad?), and the assessment of the event's relevance for ongoing action (e.g., can one ignore the event or does one need to respond to it?). It is assumed that some aspects of the analysis concerning the unexpected event (e.g., causes and relevances) are stored as part of the schema for this event so that in the future analysis of similar events can be significantly reduced both in terms of time and cognitive effort.

The **fourth step** is the schema update (Friedman and Halpern, 1997; Gardenfors, 1978, 1992, 1994). It involves producing the immediate reactions to the unexpected event (if it is the case), and/or operations such as the update, extension, or revision of the schema or sets of beliefs that gave rise to the discrepancy. The schema change (belief update process) ideally enables one to some extent to predict and control future occurrences of the schema-discrepant event and, if possible, to avoid the event if it is negative and uncontrollable, or to ignore the event if it is irrelevant for current action.

2.1.4 Artificial Surprise

The research on artificial surprise can be thought of as a natural extension of the broader field of affective computing (Picard, 1997, 2003; Picard et al., 2002). The main motivation of both fields is essentially the assumption that artificial agents empowered with emotional mechanisms similar to or inspired in those used by humans may take advantage in several dimensions possibly including intelligence (Martinho and Paiva, 2006).

In the context of artificial surprise for artificial agents two models can be stressed: the model proposed by Macedo and colleagues (Macedo, 2006; Macedo and Cardoso, 2001,

2004; Macedo et al., 2004) and the model proposed by Lorini and Castelfranchi (2006, 2007). Both models were mainly inspired by the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997), presented in Section 2.1.3, and have influence of the analysis of the cognitive causes of surprise from a cognitive science perspective proposed by Ortony and Partridge (1987). For a detailed description of the similarities and differences of the models, written by Macedo, Cardoso, Reizenzein, Lorini, and Castelfranchi, please see (Macedo et al., 2009).

Other computational approaches of artificial surprise have been proposed for other context rather than artificial agents. For instance, Baldi and Itti (Baldi and Itti, 2010; Itti and Baldi, 2009) define surprise as a measure of how data affects an observer by specifically computing the difference between the posterior and prior beliefs about the world. Like the artificial surprise model proposed by Macedo and colleagues, the Baldi and Itti model rely on the Bayesian or subjectivist framework of probability theory in which the degrees of belief are associated with different hypotheses or outcomes.

The empirical tests we performed provide evidence in favor of using the model proposed by Macedo and colleagues in our work.

Artificial surprise model proposed by Macedo and colleagues

Macedo and colleagues carried out an empirical study (Macedo et al., 2004) with the goal of investigating how to compute the intensity of surprise in an artificial agent. Their model can be considered as an extension of the BDTE or at least in line with the BDTE, particularly with respect to the computation of the surprise intensity.

They proposed, consistent with probability theory, several alternative functions for computing the surprise intensity based on the assumption that the surprise “felt” by an agent elicited by an event E_g is proportional to the degree of unexpectedness of the event E_g . They examined the functions by carrying out a two-step experiment. In the first step, they collected ratings of probability and surprise intensity provided by human participants in two domains namely political elections and sports games. Then, in the second step, they empowered artificial agents with the alternative functions as well as with the ratings of probability provided by human participants so that the artificial agents were able to compute the surprise intensity values. Finally, the values obtained by the artificial agents were compared with the actual surprise intensity given

by human participants. Their study suggests that the intensity of surprise about an event E_g , from a set of mutually exclusive and exhaustive n events, $\{E_1, E_2, \dots, E_n\}$, is a nonlinear function of the difference, or contrast, between its probability/belief and the probability/belief of the highest expected event (E_h) in the set of mutually exclusive events $\{E_1, E_2, \dots, E_n\}$.

Formally, let (Ω, A, P) be a probability space where Ω is the sample space (i.e., the set of possible outcomes of the event), $A = \{A_1, A_2, \dots, A_m\}$, is a σ -field of subsets of Ω (also called the event space, i.e., all the possible events), and P a probability measure which assigns a real number $P(F)$ to every member F of the σ -field A . Let $E = \{E_1, E_2, \dots, E_n\}$, $E_i \in A$, be a set of mutually exclusive events in that probability space with probabilities $P(E_i) \geq 0$, such that $\sum_{i=1}^n P(E_i) = 1$. Let E_h be the highest expected event from E . **The intensity of surprise about an event E_g** , defined as $S(E_g)$, is calculated as

$$S(E_g) = \log_2(1 + P(E_h) - P(E_g)) \quad (2.1)$$

where E_h is the event with the highest probability in the set (Macedo et al., 2006; Teigen and Keren, 2003). In each set of mutually exclusive events, there is always at least one event whose occurrence is unsurprising, namely E_h .

Although the model of Macedo and colleagues was first used in artificial agents whose central function was to explore unknown environments, Macedo has been extending his ideas to other fields such as intelligent transportation system (Macedo, 2010), working specifically on how their model of surprise can be used as a component/dimension of a selective attention mechanism (Macedo, 2013).

Figure 2.4 shows the surprise intensity about event E_g as a function of the difference between $P(E_h)$ and $P(E_g)$.




FIGURE 2.4: Surprise as a function of the difference between $P(E_h)$ and $P(E_g)$, i.e., Equation 2.1 that is $S(E_g) = \log_2(1 + P(E_h) - P(E_g))$ (Macedo et al., 2004).

2.1.5 Computational cognitive modeling of the coin tossing game

To illustrate how previous concepts presented so far throughout Section 2.1, specifically the BDTE concepts and the artificial surprise model proposed by Macedo and colleagues, can be used to model and compute the happiness, unhappiness, and surprise emotions, consider the scenario provided by the familiar and well-known coin tossing game. A coin tossing game generally consists in a two-player game and two different moments of playing. In a first moment one player chooses one of the two faces of a coin, heads or tails, so that the other face is automatically assigned to the other player. Then, the coin is flipped, possibly thrown into the air such that it rotates several times, until it is caught in the air and inverted, or allowed to land on the ground, depending on the rules agreed by the players in advance. When the coin comes to rest (or is intentionally stopped), the toss is complete and the player who picked or was assigned to the face-up side is declared the winner. Let us assume in this example, for the sake

of the illustration, a fair game (equally likely outcomes) and that the loser needs to pay €10 to the winner. In the coin tossing game there will be always a winner and a loser.

Based on the BDTE, coin tossing game players, in this case Agent A and Agent B, might be cognitive modeled as follows. A proposition p can be defined as “I have picked the right coin face therefore I will win €10”. The strength of a belief in proposition p at time t can be defined as $b(p, t) = 0.5$, possibly meaning that the player is “rational” in the sense that she/he is assuming the mathematical probability of the outcomes as her/his subjective belief, i.e., there is no bias such as optimism or pessimism. Similarly, the strength of a desire about proposition p at time t , can be defined as $d(p, t) = +100$ (with counter domain -100), meaning that players have a strong desire in favor of p , i.e., they strongly desire to win €10. Agent A and Agent B share the same $b(p, t)$ and $d(p, t)$ values.

After flipping the coin, agents receive a new belief regarding proposition p , i.e., the face-up side is either heads or tails. Assume that Agent A has beforehand picked heads, Agent B was assigned to tails, and the face-up side outcome is also heads. In this case, Agent A will be declared winner, whereas Agent B will be declared loser.

The Belief-Belief Comparator (BBC) compares the newly acquired belief (face-up side is heads) with all pre-existing beliefs. Considering the artificial surprise model proposed by Macedo and colleagues, the two mutually exclusive and exhaustive possible outcomes can be defined as E_g (heads) and E_h (tails) both with $P(E_h)$ and $P(E_g)$ equal to 0.5. Therefore, **in this case there is no surprise** (surprise with 0 intensity) regardless of the outcome, i.e., $S(E_g) = \log_2(1+0.5-0.5)$ or $S(E_h) = \log_2(1+0.5-0.5)$ equal 0. As a result, BBC yields a belief-confirmation signal. Neither Agent A nor Agent B would “feel” surprise, given the fact that there is no difference between $P(E_i)$ of all outcomes.

Similarly, **the Belief-Desire Comparator (BDC)** compares the newly acquired belief with all pre-existing desires. At this moment, Agents A and B are aware of which outcome actually happened (either heads or tails). On the one hand, Agent A’s BDC detects a match between a newly acquired belief (face-up side is heads) and its pre-existing desire (heads) so that it yields a desire-fulfillment signal. From a quantitative perspective, presented in Table 2.2, Agent A had $d(p, t) = +100$, $b(p, t) = 1$ and $d(p, t) > 0$ and therefore the desire-fulfillment signal is the happiness emotion with intensity of 100, i.e., $Happiness(p, t) = \phi_{ha}[d(p, t)]$. On the other hand, Agent B’s BDC detects a mismatch

between a newly acquired belief (face-up side is heads) and its pre-existing desire (tails) so that it yields a desire-frustration signal. Similar to Agent A, Agent B had $d(p,t) = +100$, $b(p,t) = 1$ and $\wedge d(p,t) > 0$ and therefore the desire-frustration signal is the unhappiness emotion with intensity of 100, i.e., $Unhappiness(p,t) = \phi_{uh}[d(p,t)]$.

2.2 Financial Markets

We start by presenting in Section 2.2.1 some basic definitions in the context of the financial markets. Then we briefly describe some statistical features, known as stylized facts, found in a series of different financial markets. We also describe three investment strategies commonly used by market participants in Section 2.2.2.

We then present different and somewhat opposing perspectives on how the financial markets work. First, we describe in Section 2.2.3 the Efficient Markets Hypothesis (EMH) as an example of how the traditional economic theories tend to consider the behavior of market participants as rational and markets as efficient. Subsequently, we show in Section 2.2.4 the main findings from behavioral economics research, stressing how the deviations from the so-called rational behavior, known as behavioral biases, contrast with the idea of rationality and efficiency claimed by the traditional economic theories. Furthermore, we present how three behavioral biases namely herding, over-reaction and loss aversion help in explaining and better understanding the behavior of market participants. Third, we present in Section 2.2.5 the Adaptive Markets Hypothesis (AMH), a relatively recent and interesting hypothesis that reconciles market efficiency with behavioral biases, by applying the principles of evolution and considering recent findings from other disciplines (e.g., cognitive neuroscience) on how market participants actually make decisions.

We conclude by presenting in Section 2.2.6 how the word surprise has been used in the context of economics and finance, particularly in financial markets. We include empirical observations and real world examples in this presentation.

2.2.1 Stock Markets

The most prominent and popular example of the financial markets is possibly stock markets. Stock markets play a very important role in societies by helping in improving

the capital allocation process and therefore contributing to economic growth. Broadly speaking, for companies, stock markets are a relatively cheap and good way of raising money, especially when compared to other classical ways such as bank lending (Frank, 2005, 2008; Graham, 1986). Similarly, for market participants, stock markets offer opportunities for investing in a wide range of companies as well as the possibility of participating in company earnings (e.g., dividends) (Bodie et al., 2004). For further works on the role of the financial markets, including the stock markets, see (Fischer and Merton, 1984; Greenwood and Smith, 1997; Wurgler, 2000).

The stock markets can be defined as a complex and dynamic socio-economic system in which a very large number of heterogeneous market participants like retail (individual) and institutional investors, investment banks, banks, and hedge funds, etc, negotiate a variety of different types of financial products such as stocks and derivatives (Blanchard and Johnson, 2012; Dopfer et al., 2004; Foster, 2005; Gao and Schmidt, 2005; Hull, 2005; Sornette, 2014). Participants are heterogeneous in the sense that they have, for instance, different beliefs, expectations, goals, budgets constraints, time horizons, and investment or trading strategies. Additionally, they can be either a human, artificial (e.g., fully automated), or hybrid (e.g., semi-automated) participant.

The negotiations among buyers and sellers are typically mediated by officially authorized institutions (e.g., brokers) so that the identities of the negotiating parts are unknown to each other. Buyers and sellers interact by submitting bids or asks, orders for buying or selling assets, respectively, in such a way that higher bids and lower asks are aligned at the top of the order book (Easley and O'Hara, 1995; Treleaven et al., 2013). A negotiation occurs when a bid matches an ask, i.e., a buyer accepts purchasing something for a price a seller is willing to receive. The underlying mechanism that regulates negotiations and price formation is known as continuous double auction mechanism (Biais et al., 2005; Madhavan, 2000). In a double-auction market, multiple buyers and sellers continuously interact in order to determine prices (Chang, 2012; Keynes, 1936). When multiple buyers compete against each other prices rise (Phelps, 2007), whereas when multiple sellers compete against each other prices fall, in a mechanics and rationale that is similar somehow to the law of supply and demand (Mas-Colell et al., 1995). As a strategic system, the output is the result of the aggregation of the individual decisions of all market participants interacting within this rich and diverse market “ecology” (Farmer, 2002).

In the uncertain environment of a stock market, **market participants need to constantly analyze and revise their beliefs, expectations, and investment strategies in the light of the massive amount of information available.** Such exogenous information include a wide variety of facts related not only to valuation and price formation mechanisms (e.g., earnings announcements) but also economic and financial indicators of different types (e.g., Gross Domestic Product (GDP) (Landefeld et al., 2008; Struzik, 2003), Unemployment Rate) (Warneryd, 2001). Besides, as we explain in Section 2.2.4, endogenous information by itself, i.e., the history of asset prices, dividends and trading volume, also plays an important role in general market dynamics. Another fundamental aspect of a stock market is that its participants need to interact in an environment that poses a set of different types of risks such as regulatory changes, terrorist attacks, global economy-wide shocks, insolvency rumors (Artzner et al., 1999; Fama and French, 1993; Friedman and Savage, 1948; Knight, 1921; Meder et al., 2013; Rumsfeld, 2002; Taleb, 2007, 2014). To cope with these risks, market participants use different tools (e.g., Value-at-Risk (VaR), presented in Section 4.2) to assess, subjectively (qualitatively) or objectively (quantitatively), the probability of a wide range of events which may have either a positive or negative influence on asset prices (Bloom, 2009; Kahneman, 2009; Lo and Mueller, 2010; Lo, 1999; Markowitz, 1952). Nevertheless, market participants interact based on the risk-return trade-off, meaning that while lower levels of risk are associated with low levels of potential returns, higher levels of risk are associated with higher levels of potential returns (Sharpe, 1994).

The behavior, evolution, and dynamics of the major financial markets, especially of the major stock market indexes (Lo, 2016b), has been extensively observed and investigated. Indeed, major stock market indexes, such as the S&P500 index², i.e., an index created by Standard & Poor's, based on the market capitalizations of the 500 largest companies which are publicly held on either the NYSE (The New York Stock Exchange)³ or NASDAQ⁴ stock exchanges, possibly the most important stock market index in the world, are constantly monitored since they are generally assumed to be important gauges for helping the understanding of the current and future state of the whole economic and financial system (e.g., (Aruoba et al., 2009; Beber et al., 2011; Stock and Watson, 2003)). To illustrate a typical stock market dynamics, we present

²<http://us.spindices.com/indices/equity/sp-500> [Accessed: July 2017]

³<https://www.nyse.com> [Accessed: July 2017]

⁴<http://www.nasdaq.com/> [Accessed: July 2017]

in Figures 2.5 and 2.6, respectively, the daily evolution and volatility (i.e., in this case the rate of change) of the S&P500 index from October 2005 to January 2015.



FIGURE 2.5: Daily evolution of the S&P500 index from October 2005 to January 2015. Figure generated with the Economic Research Federal Reserve Bank of St. Louis tool - <http://research.stlouisfed.org/> - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.




FIGURE 2.6: Daily volatility of the S&P500 from October 2005 to January 2015. Figure generated with the Economic Research Federal Reserve Bank of St. Louis tool - <http://research.stlouisfed.org/> - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.

The stock markets generate a feedback effect from market prices to the real economy, affecting a series of other variables (Apergis et al., 2015; Bond et al., 2011; Cochrane,

2005). In this kind of situation, however, it is very difficult (perhaps impossible) to establish a cause and effect relationship between such different types of variables, e.g., how market prices affect consumer confidence/expectations that, in turn, affect consumption that ultimately affects market prices (Chen, 2011). Nevertheless, there are historical and empirical evidence that when a stock market crash occurs, as happened in 1929 (e.g., (Bernanke, 1983)), 1987 (e.g., (Shiller, 1987)) and more recently in 2008 (e.g., (Farmer, 2012)), it might have significant and dramatic consequences for the real economy, including even the life of those who do not know what a stock market is (e.g., (Barro, 1990; Kindleberger et al., 2005; Lewis, 2011; Reinhart and Rogoff, 2009a)). For instance, we illustrate in Figure 2.7 the relationship between the S&P500 index, the U.S. Real Gross Domestic Product (GDP)⁵, and the U.S. Civilian Unemployment Rate⁶. As we can observe, there seems to exist an inverse relation between the first two (the S&P500 index and the GDP) and the third (unemployment rate). This kind of inverse relation between different variables are commonly found in economics and finance and is in line with micro (Frank, 2005; Mas-Colell et al., 1995) and macroeconomics theory (Blanchard and Johnson, 2012; Mankiw, 2014). In this particular case, for instance, during the recession period (shaded grey area), while the S&P500 index and the GDP fell by almost 40%, the unemployment rate rose by almost 100%.

⁵<https://fred.stlouisfed.org/series/GDPC1> [Accessed: July 2017]

⁶<https://fred.stlouisfed.org/series/UNRATE> [Accessed: July 2017]




FIGURE 2.7: Evolution of the S&P500 index, U.S. Real Gross Domestic Product (GDP), and U.S. Civilian Unemployment Rate, from October 2005 to January 2015. Figure generated with the Economic Research Federal Reserve Bank of St. Louis tool - <http://research.stlouisfed.org/> - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.

Stylized Facts

Empirical studies of the time series of stock returns have been documenting the presence of a series of non trivial statistical features or regularities across a wide range of different sorts of instruments, markets and time periods (Campbell et al., 1996; Pagan, 1996). **Such statistical properties are known as “stylized facts”** (Bouchaud et al., 2002; Chakraborti et al., 2011a; Challet et al., 2001b; Cont, 2001; Cutler et al., 1991; Schinckus, 2014) and their presence seems to be ubiquitous in all sorts of financial markets. The identification of these stylized facts is important because the ability of reproducing at least some the stylized facts is one of the criteria that has been used by researchers to claim that a given financial market model is effectively able to produce realistic results. Let us briefly present **the following five stylized facts** (for a complete and detailed list of stylized facts see (Cont, 2001)):

- **Heavy/fat tails:** in a normal distribution, considering the mean (μ) and the standard deviation (σ), about 68% of values are within $\mu \pm 1\sigma$; about 95% of values are within $\mu \pm 2\sigma$; and about 99.7% are within $\mu \pm 3\sigma$. This fact is known as the 3-sigma rule. Additionally, statistically speaking, a normal distribution has a

skewness (a measure of the asymmetry of the distribution) approximately equal to 0 and kurtosis (a measure of the “peakedness” of the distribution) approximately equal to 3. The human being heights, weights and IQ scores are typical examples of values that follow a normal distribution. Unlike the normal distribution, in a heavy/fat tail distribution returns are not normally distributed.

- **Aggregational Gaussianity:** the shape of the distribution is not the same at different time scales. As one increases the time scale over which returns are calculated, their distribution looks increasingly like a Normal distribution.
- **Gain/loss asymmetry:** returns display large drawdowns in values but not equally large upward movements.
- **Volatility clustering:** different measures of volatility display a positive auto-correlation over several days, i.e., high-volatility events tend to cluster in time. In other words, “large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes” (Mandelbrot, 1963, 1997). Volatility clustering resembles the concept of entropy used in a variety of areas such as information and communication theory.
- **Volume/volatility correlation:** volume is correlated with all measures of volatility. Volume is typically considered as a valuable indicator containing information regarding the quality and arrival of information which is not contained in prices.

2.2.2 Investment Strategies

The investment or trading strategies commonly used by market participants in the financial markets can be classified in the following three great categories (not necessarily mutually exclusive): fundamental analysis, technical analysis, and what is known as noise behavior or zero intelligence. It is important to bear in mind that our objective in this Section is simply to briefly present some mechanisms widely used by market participants in the context of a financial market and not to provide evidence in favor of or against one or another rationale. The debate, though interesting, is beyond the scope of this work.

Fundamental analysis

Fundamental analysis (Bodie et al., 2004; Hagstrom, 2005; Lynch, 1994) is concerned with the **estimation of the intrinsic (also referred to as fundamental or fair) price of an asset now or in a given point in future**. It begins with the estimation of the intrinsic value of a company through the use of different valuation models and techniques (e.g., discounted cash flow model). This estimation includes not only tangible assets but also intangible assets such as information regarding the management, the profile of its owner, her/his ambitions, goals, and plans with respect to the company. This work also involves the analysis of a wide range of factors like the understanding of the current macroeconomics scenario as well as the forecasting of possible scenarios. Additionally, it requires a close examination of financial information regarding the company such as cash flow, earnings, and investment plans. Finally, fundamental analysis is interested in identifying the internal and external risks which may affect the company. With the intrinsic value of a company calculated, the fair asset price is calculated by dividing the intrinsic value by the total number of shares of that company. The underlying assumption of fundamental analysis is that price always returns to its fair value at some point in time. Therefore, one interested in buying a given asset should look for assets whose current price is under the fair price, i.e., undervalued.

The work of fundamental analysis often results in a variety of **ratios** which are used to estimate whether a current asset price deviates from its fair price (e.g., (Campbell and Shiller, 1988a,b; Fama and French, 1988)). For instance, one ratio widely used is the P/E ratio, i.e., price-to-earnings ratio or simply “multiple”. The P/E ratio is calculated by dividing the price per share by the annual earnings per share. So, if a current asset price is \$70 and the earnings per share for the most recent 12 month period is \$7, then the asset has a P/E ratio of 10 ($70/7$). The P/E is generally used in the process of identifying undervalued (lower P/E ratios) or overvalued (higher P/E ratios) companies. However, decisions of buying or selling should not be based only on the P/E. Examples of other indicators are the EV/EBITDA ratio (i.e., Enterprise Value/Earnings Before Interest, Taxes, Depreciation and Amortization), and the DY ratio (i.e., annual dividends per share/price per share).

The process for estimating the intrinsic value of a company is essentially of subjective nature so that professionals (e.g., analysts) often diverge in estimating the intrinsic value of a same company. One of the problems with fundamental analysis relies on

the fact that fundamentals are not totally observable (Akerlof, 1970; Timmermann and Granger, 2004). Therefore, the success of an investment strategy based on fundamental analysis resides essentially in the estimation accuracy regarding the fair price.

Technical analysis

Technical analysis (also referred to as graphical analysis) (Edwards and Magee, 2010; Lo et al., 2000; Murphy, 1999) make use of a myriad of **statistical and mathematical indicators** derived from and build upon stocks data together with graphical patterns and tools to understand stock prices behavior, identify trends (Lempérière et al., 2014) and ultimately to discover and explore profit opportunities. The underlying assumption is that prices incorporate all relevant information and both past behavior and returns are rich in information concerning future behavior so that they can be used to some extent to predict or indicate future movements, i.e., history repeats itself. Finally, considering that all that matters is incorporated in asset prices, one does not need to be concerned with exogenous information, i.e., anything that happens outside the market in itself. In the strict sense of the definition, one may claim that questions such as those addressed by fundamental analysis can be “ignored”.

For instance, the MA (Moving Average) indicator is widely used to identify how prices evolve over a certain period of time. There are essentially two types of moving average the SMA (Simple Moving Average), in which all previous periods have the same weight, and the EMA (Exponential Moving Average), in which the most recent periods have a higher weight. MAs are often used in a variety of ways. For example, longer MA or EMA (e.g., 52, 200 periods) are generally used to identify long-term price references. Conversely, shorter MA or EMA (e.g., 5, 13 periods) are understood as reflecting short-term movements and/or opinions of market participants about a given asset. Examples of other technical indicators are the Bollinger Bands (two standard deviations away from a certain SMA), MACD (Moving Average Convergence/Divergence), OBV (On Balance Volume), and the RSI (Relative Strength Index). Some indicators such as “bottom and top”, “support and resistance”, and “top historical asset price” are assumed to be a sort of “psychological” barriers for both asset prices and market participants. Similarly, “ascending or descending bottoms and tops” are assumed to characterize “channels” or trends so that as long as prices oscillate inside the channel the trend defined by a given number of market participants is still valid.

Despite of its long history, there are those who claim that technical analysis is inconsistent with the Efficient Markets Hypothesis (EMH) (e.g., (Fama, 1965; Jen, 1970)), described in Section 2.2.3, and that technical analysis resembles more “alchemy” (Malkiel, 2016) than a truly scientific tool. Nevertheless, there are evidence from a great number of works with a significant diversity in terms of markets, historical periods, and methods, which suggest that the use of technical analysis is useful in finding profit opportunities (e.g., (Antoniou et al., 1997; Bessembinder and Chan, 1995; Blume et al., 1994; Brock et al., 1992; Chincaroli and Valls Pereira, 2015; Cochrane, 2008, 2011; Daniel and Titman, 1999; Edwards and Magee, 2010; Kwon and Kish, 2002; LeBaron, 1999; Lee et al., 2013; Lo et al., 2000; Menkhoff et al., 2012; Neely and Weller, 1999; Oberlechner, 2001)). Perhaps part of the success of technical analysis results from a phenomena known as **self-fulfilling prophecy** (Azariadis, 1981; Jordan; Merton, 1948). When a large number of market participants believe in a similar “prophecy”, they might behave in a way that their actions collectively contribute to turn the prophecy into reality. For example, if a significant number of market participants believe in the “prophecy” of “when price approximates a certain price region it will reverse”, i.e., as a result of the fact that some market participants might be using a similar MA (moving average) configuration, and when they are confronted with the scenario in which the price really approximates that price region, they might start buying the asset therefore raising the price and ultimately fulfilling the prophecy. The debate about technical analysis remains, and perhaps will be for a long time, active as those interested in obtaining concrete arguments need to cope with intrinsic difficulties such as the task of testing the EMH.

Noise or zero intelligence

Some market participants, known as “noise traders” (Black, 1986), **do not employ any “rational” mechanism, tend to form incorrect beliefs and expectations as well as base their actions on what they consider to be worth information but actually is simple noise.** We may also refer to such market participants as zero intelligence traders. The concept of zero intelligence was introduced by Gode and Sunder (1993) to describe a particular kind of market participant that has no strategy (behave at random), no intelligence, does not seek or maximize profits, does not have the ability to observe, remember, or learn. For a review of the zero intelligence methodology see (Ladley, 2012).

The idea of noise traders as well as the zero intelligence concept has been studied in several ways. Such studies include the general impact of noise traders in the financial markets (e.g., (Farmer et al., 2005; Kelly, 1997; Serrano-Padial, 2012)) as well as the comparison of the profitability of a random strategy with other strategies (e.g., (Biondo et al., 2013)). For instance, De Long and colleagues (DeLong, 1991; DeLong et al., 1990) demonstrated that noise traders can take disproportionate amount of risk that eventually lead them to earn higher expected return compared to the returns of “rational” investors (Palomino, 1996). Nevertheless, noise traders generally cause markets to be to some extent less efficient and perhaps their behavior can be used to understand the profit opportunities that tend to appear as well as the deviations in asset prices that occasionally occur (Hirshleifer and Luo, 2001; Shleifer and Summers, 1990).

2.2.3 Efficient Markets Hypothesis

Perhaps the **most iconic example of how traditional economic theories tend to consider market participants individually and of how the financial markets work** is illustrated by the Efficient Markets Hypothesis (EMH) (Fama, 1970, 1991) proposed by Fama (2011). The EMH relies on the assumption that market participants have stable and well-defined preferences (Simon, 1955). Additionally, it assumes that when participants are confronted with decisions that involve risk, they are able to correctly form their probabilistic assessments according to the laws of probability, calculating which of the alternative courses of action maximize their Expected Utility (EU) over n uncertain outcomes, defined as

$$EU = \sum_i^n p_i U_i \quad (2.2)$$

where p_i and U_i are, respectively, the probability and the utility of outcome i (Chiodo et al., 2003; Hirshleifer, 2001; Loewenstein et al., 2008; Rabin, 1998, 2002; Smith, 1991; Smith et al., 1988).

The EMH assumes that **market prices fully and instantaneously incorporate the information and expectations of all market participants**. There are essentially two types of information, endogenous and exogenous information. Endogenous information compromises, for instance, the history of asset prices, possibly dividends

and other variables related to asset prices such as the trading volume. Exogenous information correspond to any sort of information, facts, etc, that market participants commonly receive in the context of the financial markets such as earnings announcements, mergers and acquisitions. The underlying assumption is that market participants form expectations rationally and are able to make optimal decisions (Lo, 2005). Therefore, the majority of market participants are efficient in reflecting new information accurately (Black, 1986; Fama, 1965; Malkiel, 2003). Last but not least, the EMH assumes that market participants have no cost in acquiring and analyzing information (Tsang, 2009).

A market in which prices always “fully reflect” available information is called “efficient” (Fama, 1970). More formally, a market is efficient with respect to information set Ω_t if it is impossible to any market participant to make profits by trading on the basis of information set Ω_t (Jensen, 1978). There are three forms of market efficiency based on the set of variables contained in information set Ω_t , namely weak form, semi-strong-form, and strong form. In the weak form efficiency, Ω_t includes the history of asset prices, possibly dividends and other variables related to asset prices (e.g., trading volume). In the semi-strong form, Ω_t includes information contained in the weak form plus all information known to all market participants, i.e., publicly available information. Similarly, in the strong form efficiency, Ω_t includes information contained in the semi-strong form plus all information known to any market participant, both public and private information (Vega, 2006). Despite its apparent simplicity, the EMH is difficult to test in practice. Weak-form tests are concerned with how well historical asset prices and past returns can be used by market participants to predict future returns. Semi-strong-form tests are concerned with whether prices efficiently adjust to information that is obviously publicly available to all market participants at the same time, such as announcements of earnings, mergers and acquisitions. Strong-form tests are concerned with whether an individual or a group of market participants have privileged access to any information relevant for price formation, such as a forthcoming acquisition, so that such private information is not fully reflected in asset prices.

Additionally, according to the EMH, neither arbitrage opportunities exist nor an investment strategy would allow a market participant to earn above-average returns without accepting above-average risk. For example, Graham (1976) demonstrated doubts about the benefits of expending extensive active efforts in attempts to outperform the market. Furthermore, Barber et al. (2009) demonstrated how much individual investors lose by

trading. Similarly, Warren Buffet, perhaps one of the greatest investors of all time, at least in terms of absolute returns, claimed that for most investors, including both institutional and individual investors, the best way of investing is through an index fund that charges minimal fees so that investors following this strategy are certainly those that will beat the great majority of investment professionals (Buffet, 1996). Consistent with Buffet's ideas, Malkiel (2005) demonstrated that professional investment managers do not outperform their index benchmarks indexes (e.g., see (Foster and Warren, 2015, 2016) for other works). In an efficient market, however, "noise" behavior, presented in Section 2.2.2, and periods of market stress, might occasionally lead market participants to value assets in such a way that prices do not reflect the underlying fundamentals. Nevertheless, such significant deviations do not tend to last "too long" or become "too big" because market participants tend to correct them (Timmermann and Granger, 2004).

The idea of market efficiency has been always receiving critics (e.g., (Hellwig, 1980; Summers, 1986)). For instance, Grossman (1976) and Grossman and Stiglitz (1980) argue that perfectly informationally efficient markets are an impossibility. The basic argument to support this strong statement is essentially that if markets are perfectly efficient that would be no profit opportunities and therefore no incentives for market participating in gathering information as well as in spending efforts in the price discovery process/mechanism, i.e., the process of determining the price of an asset, as explained earlier in Section 2.2.2. Conversely, they claim that markets need to have sufficient profit opportunities (inefficiencies) to compensate information gathering and price discovery, two pillars of the gigantic financial services industry.

2.2.4 Behavioral Economics

Behavioral economics is **an alternative perspective to the understanding of human behavior in economic and financial contexts, including financial markets**. It has emerged, at least in part, in response to the difficulties faced by researchers in explaining some financial phenomena, anomalies, and empirical observations of asset price dynamics, by using only the assumptions, models, and theoretical background provided by the traditional economic theories, such as the EMH. Such evidence range from ordinary and relatively simple decisions (e.g., see (Ariely, 2010; Frank, 2008;

Thaler and Sunstein, 2008) for interesting examples) to complex aggregated behavior (e.g., the occurrence of bubbles (Akerlof and Shiller, 2009; Shiller, 2002a)).

Behavioral economics (e.g., (Barberis and Thaler, 2002; Camerer, 1995, 2014; Camerer et al., 2003; Kahneman and Tversky, 1979b; Martino et al., 2006; Mullainathan and Thaler, 2000; Shiller, 2003; Thaler, 2016)), also referred to in the literature as behavioral finance (Hirshleifer, 2015), **can be defined as the result of the combination of psychology and economics disciplines** that aims to understand human judgment and decision-making under uncertainty, specifically in order to identify the ways in which behavior differ from the standard models, as well as how this behavior matters in economic and financial contexts. Behavioral economics is particularly useful when humans are confronted with decisions that are complex and in which optimality is difficult to achieve.

Psychology has indeed been playing a significant role to the development of behavioral economics. For instance, the “The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel”, the so-called Nobel prize in Economics, was awarded in 2002 to Vernon L. Smith (an economist) and to Daniel Kahneman (a psychologist). Vernon L. Smith was awarded “for having established laboratory experiments as a tool in empirical economic analysis, especially in the study of alternative market mechanisms” (The Nobel Prize Foundation, 2002) (for a summary of Vernon L. Smith contributions see (Caginalp et al., 2003; Chapman, 2003)). Daniel Kahneman was awarded “for having integrated insights from psychological research into economic science, especially concerning human judgment and decision-making under uncertainty” (The Nobel Prize Foundation, 2002). Daniel Kahneman did a significant part of the work that lead him to win a Nobel prize in Economics with Amos Tversky, a fellow psychologist (for a summary of their contributions see (Shefrin and Statman, 2003)). Sadly, Amos Tversky died in 1996, otherwise he would probably have shared the prize with Daniel Kahneman and Vernon L. Smith.

Behavioral economics researchers have been indeed providing **extensive experimental evidence**, the majority resulted from laboratory experiments with human subjects, that humans have deviations from the so-called rational behavior (Barberis and Thaler, 2002; Gigerenzer and Gaissmaier, 2011; Mullainathan and Thaler, 2000; Summerfield and Tsetsos, 2015). Such deviations are known as **behavioral biases, and they are believed to be ubiquitous to humans, and several of them are clearly counterproductive from the economics and financial perspective**. Examples of

behavioral biases are herding (Chen, 2013; Eguiluz and Zimmermann, 2000; Prechter, 2001; Raafat et al., 2009), hyperbolic discounting (Laibson, 1997), loss aversion (Kahneman and Tversky, 1979b), miscalibration of probabilities (Lichtenstein et al., 1982), overconfidence (Barber et al., 2009; Barber and Odean, 2001; Benos, 1998; Gervais and Odean, 2001; Hirshleifer and Luo, 2001), overreaction (De Bondt and Thaler, 1985), psychological accounting (Tversky and Kahneman, 1981), and regret (Bell, 1983; Clarke et al., 1994). To illustrate how the behavioral biases affect human decision-making as well as how they help in better understanding the behavior of market participants, we explain later in this Section three of them, namely herding, overreaction and loss aversion.

Although the behavioral economics perspective might be seen, in principle, in conflict with the traditional economic perspective, the behavioral economics findings do not necessarily imply a total rejection of the assumptions (e.g., utility maximization, and efficiency) claimed by traditional economic theories (Sommer, 2013), such as the EMH. For example, as a concrete example of how these two different perspectives can be seen as alternative and complementary, interestingly, the so-called Nobel prize in Economics was awarded in 2013 to Eugene F. Fama, the proponent of the EMH (presented in Section 2.2.3), Lars Peter Hansen, and Robert J. Shiller, the latter being one of the first and more important researchers in the field of behavioral economics, “for their empirical analysis of asset prices” (The Nobel Prize Foundation, 2013).

Regardless of some critics (e.g., (Fama, 1998, 2014; Frankfurter and McGoun, 2000)), **some important conclusions can be drawn from the behavioral economics research.** Specifically, humans (and therefore market participants) do not always have enough time or the cognitive ability to process all the related information with accuracy, i.e., they have what is widely known as bounded rationality (Newell, 2005). Rather than maximizing the expected utility, being able to make optimal decisions, market participants maximize the perceived expected utility, based not on actual probabilities but on their beliefs about those probabilities (Oaksford and Chater, 2001). Additionally, they tend to exhibit behavioral biases that can lead them to economic ruin. Perhaps the most important conclusion is that market participants are not fully rational, as well as that the human behavior might be better understood by considering evidence from other disciplines, particularly from psychology, leading to the creation of novel, more behaviorally plausible and realistic scenarios. **The behavioral economics research has been helping** in better understanding the true nature of human judgment and

decision-making under uncertainty, particularly in economic and financial contexts, as well as in revealing the high level of sophistication and complexity of such system, specifically by shedding light on how a myriad of intertwined factors interact, affect and influence judgments and decisions.

Herding

Herding is an influential and well-documented feature of the human behavior. Herding can be defined as **a form of convergence or alignment of thoughts, behavior and actions of individuals in a group (herd)**. It generally takes places through local interaction and without centralized coordination (Bikhchandani et al., 1992). For a review of some theoretical frameworks for describing herding see (Raafat et al., 2009).

Herding behavior has been extensively documented in a number of diverse domains, particularly in **economics and finance** from different perspectives. The most relevant works in the context of this work include evidence on how herding may be used as a tool to explain the occurrence of bubbles (e.g., see the work of Shiller and colleagues (Akerlof and Shiller, 2009; Shiller, 2000, 2002b)), and an extensive and vast number of works that focus on investigating different aspects of herding in stock markets (Andersson et al., 2014; Chen, 2013; Eguiluz and Zimmermann, 2000; Hirshleifer and Hong Teoh, 2003; Phelps and Ng, 2014; Prechter, 2001; Yang et al., 2012). For a detailed review of the theory and evidence regarding herd behavior in the financial markets see (Hirshleifer and Hong Teoh, 2003).

For an economic and financial context, the basic rational of the herding effect can be defined as follows: a market participant thinks that other market participants (the herd, the market as a whole) possess (perhaps better) information that she/he does not possess regarding a given asset and, therefore, following them by imitating their actions seems the best action to take. So if prices are falling (rising) it should be better do sell (buy). In this illustration, for the sake of the understanding, we are only taking into account the voluntary actions of a market participant under normal circumstances and not actions that a market participant may be compelled to execute (e.g., an action resulting from a significant margin call). Interestingly, the herding effect can occur only in the presence of endogenous information, stressing again the complexity of economic and financial systems. We may also compare this basic

rationale with the **beauty contest** perspective presented by John Maynard Keynes (e.g., (Chambers and Dimson, 2013; Skidelsky, 2010)):

“ ... professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one’s judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees.” (Keynes, 1936).

Overreaction

Tversky and Kahneman (1974) reported that when people need to assess the probability of uncertain events, they rely on three heuristics (Gigerenzer and Gaissmaier, 2011) namely representativeness, availability, adjustment and anchoring to assess probabilities and to predict values. Representativeness is a rule of thumb by which people judge the probability or frequency of an event by considering how much the event resembles the available data (Kahneman and Tversky, 1972). Availability is a phenomenon by which people assess the frequency of a class or the probability of an event based on how easily instances or occurrences can be retrieve from memory. Adjustment and anchoring refers to a tendency of people to make estimates by starting from an initial value or familiar position, referred to as “anchor”, that is adjusted to yield the final value. Tversky and Kahneman also state that statistical principles are not learned from everyday experience because the relevant instances are not coded appropriately both by naive and sophisticated subjects. Therefore, the probability is actually subjective in the sense that different individuals are allowed to have different probabilities for the same event (Gilboa et al., 2008). Although the heuristics are economical and usually effective, they lead to severe systematic errors. For example, adjustments are typically insufficient,

and people think they see patterns in truly random sequences (Bar-Eli et al., 2006; Bar-Hillel and Wagenaar, 1991; Gilovich et al., 1985; Taleb, 2008; Wolford et al., 2004). Furthermore, it has been suggested that such biases contribute to yield overreaction.

Kahneman and Tversky (1979a) point out that in revising their beliefs, **people tend to overweight recent information and underweight prior information**, while Griffin and Tversky (1992) reported that people update their beliefs based on the “strength” and the “weight” of new evidence, where strength refers to aspects such as the salience and extremity, and weight refers to statistical informativeness such as the sample size. According to Griffin and Tversky, people pay too much attention to the strength of the evidence they are presented with and too little attention to its statistical weight.

Motivated by a variety of psychological evidence, especially by the idea of Griffin and Tversky (1992), Barberis et al. (1998) presented **a model of investor sentiment on how investors form expectations of future earnings**. They define overreaction as the fact that the average return, following a series of announcements of good news (e.g., higher than expected earnings), is lower than the average return following a series of bad news announcements. Subsequent news announcements are likely to contradict this optimism, leading to lower returns. The underlying idea is that after a series of announcements of good (bad) news, investors become overly optimistic (pessimistic) that future news announcements will also be good (bad) and, as a result, they overreact, sending the stock price to unduly high (low) levels. In other words, it leads to the prediction that stock prices overreact to consistent patterns of good or bad news.

The findings provided by Griffin and Tversky can be used, in conjunction with other theoretical background, to explain some highly unexpected market movements such as the October 19, 1987 stock market crash (also known as Black Monday), i.e., a drop of 22.6% in the Dow Jones Industrial Average (DJIA) (Shiller, 1987), and the May 6, 2010 stock market crash (what turned out to be known as Flash Crash), i.e., a drop of almost 9% in the course of about 36 minutes and a rebound with extraordinary velocity (U.S. Commodity Futures Trading Commission and U.S. Securities & Exchange Commission, 2010). **The overreaction of market participants might be caused**, among other things, due to high-strength, low-weight news event, in spite of neither changes in the underlying fundamentals nor the occurrence of a highly unexpected event (Cutler et al., 1989).

Finally, De Bondt and Thaler (1985, 1987) investigated whether market participants of the U.S. stock market tend to “overreact”, i.e., they overweight recent information and underweight base rate data, to unexpected and dramatic news events, regardless of whether the events are positive or negative. Their study provided experimental and survey evidence indicating that market participants overreaction leads prices to temporarily depart from their underlying fundamental value (rationality), with prices initially biased by either excessive optimism or pessimism in response to the arrival of new information. Specifically, they tested whether the subsequent price reversal of assets that have experienced either extreme gains or losses are pronounced. Consistent with the overreaction hypothesis, portfolios of prior losers assets outperformed prior winners assets, meaning that prior losers would be more attractive investments than prior winners.

Loss Aversion

The work of Kahneman and Tversky (1979b), known as **prospect theory** (Tversky and Kahneman, 1992), is considered one of the seminal works in the field of behavioral economics. Kahneman and Tversky identified, among other findings, what turned out to be known as loss aversion, i.e., the notion that losses are weighted more heavily than gains. Furthermore, they found that people are concerned with the changes in their wealth rather than with its final state so that they tend to be more sensitive to decreases in their wealth than to increases (Barberis et al., 2001; Schmidt and Zank, 2005; Zhang and Semmler, 2009). It leads to a tendency of people to exhibit “risk aversion” behavior in choices that involve gains and “risk seeking” behavior in choices that involve losses (Benartzi and Thaler, 1995; Studer et al., 2015; Zhang et al., 2014). Some effects, especially the certainty effect (Tversky and Kahneman, 1981), i.e., the fact that people tend to underweight outcomes that are merely probable in comparison with outcomes that are obtained with certainty, contributes to loss aversion.

For instance, consider the following **example** (extracted from (Lo, 2004) with minor adjustments). Suppose one needs to choose between two “investment opportunities” A and B: A yields a sure profit of \$240000; B yields \$1 million with 25% probability and \$0 with 75% probability. Investment A has an expected value of \$240000 (i.e., 240000×1.0), whereas investment B has an expected value of \$250000 (i.e., $1 \text{ million} \times 0.25$), which represents a higher return of investing in A, though, one will receive either \$1

million or \$0. When confronted with this kind of choice, there is no right or wrong choice since it is a matter of personal preference (Beckmann et al., 2008; Chui et al., 2010). Nevertheless, most subjects prefer investment A over B, a kind of behavior characterized as “**risk aversion**”. Likewise, suppose one needs to choose between two “investment opportunities” C and D: C yields a sure loss of \$750000; D yields a loss of \$1 million with 75% probability and \$0 with 25% probability. Although this kind of scenario may seem, in principle, not plausible, many economic and financial situations might obligate market participants to face choosing between “the lesser of two evils”. When confronted with this kind of choice, most subjects prefer investment D over C, a kind of behavior characterized as “**risk seeking**”.

The identification of this kind of behavior made it possible a better understanding of how market participants deal with gains and losses. When a market participant buys an asset that she/he believes will have certain expected return possibly within certain time horizon, two possible scenarios she/he will face is one in which the asset appreciates and another in which the asset depreciates. Analysis of trading records have provided evidence that market participants realize their profitable positions more readily than their unprofitable ones (Odean, 1998), and thus exhibited what is known as the disposition effect (Shefrin and Statman, 1985), i.e., a tendency to hold losing investments for too long and, on the other hand, to sell winning investments too soon. Although the prospect theory was proposed in late 1970’s, the research concerning the application of the concepts of loss aversion to the financial markets is still very active (e.g., see (Barberis et al., 2001; Grune and Semmler, 2008; Schmidt and Traub, 2002; Zhang and Semmler, 2009)).

2.2.5 Adaptive Markets Hypothesis

The Adaptive Markets Hypothesis (AMH) (Lo, 2004, 2005, 2012) is a relatively recent and still under development hypothesis that tries to **bring together the EMH with behavioral biases by applying the evolutionary principles of innovation, adaptation, competition, and natural selection to economic interactions**. Additionally, AMH is based on recent research in the cognitive neuroscience of decision-making as well as is heavily influenced by recent advances in the discipline of evolutionary psychology (e.g., (Cosmides and Tooby, 2013)). AMH represents a shift of thinking in financial economics from the physics to biological sciences (Farmer, 2002;

Farmer and Lo, 1999; Lo and Mueller, 2010), offering a more complete explanation of the behavior of the financial markets and their participants (Brennan and Lo, 2011, 2012).

The AMH consists of six assumptions. The first assumption, the only that is shared with the EMH, is that market participants are selfish, i.e., they act in their own self-interest. Second, market participants are neither perfectly rational nor completely irrational but are intelligent, forward-looking, competitive investors who frequently make mistakes (Hayek, 1945); third, they are capable of learning from their mistakes and adapt their expectations and behavior to new economic realities and market conditions though. Fourth, competition drives adaptation and innovation and; fifth, natural selection shapes market ecology (Blume and Easley, 1992). Adaptation results from the interactions among market participants, which are governed by natural selection, i.e., survival of the fittest. Therefore, market participants constantly need to innovate in order to discover new profit opportunities. Additionally, current market condition is a product of this selection process. Sixth, the sum of all the previous assumptions is that evolution, i.e., selfish market participants, innovation, adaptation, competition, and natural selection, determines market dynamics.

Lo (2005) points out some **concrete implications of the AHM**. The **first implication** is that the relation between risk and reward is unlikely to be stable over time. As market conditions change due to modifications in regulation and the introduction of new taxes, for example, the relation between risk and reward is affected. Additionally, natural selection determines who participates in market interactions. Participants who experienced significant losses are more likely to definitely or temporarily exit the market.

The **second implication** is that arbitrage opportunities do occasionally exist. In the context of the AMH, market efficiency is not a binary state but rather a continuum. Market efficiency is considered highly context-dependent and dynamic, i.e., it depends not only of information but also on market conditions, the number and nature of market participants. For instance, on the one hand, when a small number of market participants are competing for a given asset, price formation is likely to be less efficient and perhaps price deviates from the “fair” value. On the other hand, price formation is likely to be more efficient when a large number of market participants are competing and therefore price tends to be more “realistic”. The degree of market inefficiency

determines the effort investors will expend to gather and trade on information (Verrecchia, 1982). These profit opportunities must actually occur, acting as incentives for those interested in gathering information as well as are essential to the price discovery aspect of the financial markets, as explained in Section 2.2.3. Although as opportunities are discovered, they tend to be explored and disappear, new opportunities tend to continually occur considering that some market participants die (bankruptcy) or become extinct, exiting the market, others are born, i.e., enter the market, and the fittest-richest ones may change their expectations, goals, and behavior. Additionally, phenomena such as cycles, trends, panics, bubbles, and crashes can and do co-exist resulting in complex market dynamics (Kindleberger et al., 2005; Kirman, 1993).

The **third implication** is that investment strategies performance may vary depending on the market conditions in which it is applied to, i.e., they can perform well under certain conditions and poorly under others. For example, during normal market conditions, market participants are heterogeneous, with diverse goals, expectations, and information, resulting somehow in a kind of wisdom of crowds in terms of price formation. On the other hand, during periods of market dislocation, when market participants share same goals (e.g., capital preservation), expectations, and information, the madness of mobs can generate panic and crashes, in a similar way that euphoria can generate bubbles (Lo, 2010; Raafat et al., 2009).

Finally, the **fourth implication** is that, while other aspects, such as profit maximization, utility maximization, are certainly relevant for market participants, “survival” is the only objective that matters in the long run. Innovation is in this context the key to survival. As the relation between risk and reward is unlikely to be stable over time, market participants who innovate, evolve their capabilities, and adapt their goals, expectations, and behavior to changing market conditions are less likely to become extinct.

The AMH can be viewed as a new version of the EMH, derived from evolutionary principles (Lo, 2005), although of course the AMH is not yet as well developed as the EMH. Despite this, some recent works (Getmansky et al., 2015; Neely et al., 2009) have been providing evidence in favor of the AMH, demonstrating how the AMH can explain not only departures from the EMH but also how market may frequently shift from the wisdom of crowds to the madness of mobs and back again.

We consider that the Adaptive Markets Hypothesis (AMH) is currently, to the best of our knowledge, the most appropriate hypothesis to understand the financial markets, both the behavior of market participants individually and the financial markets globally, and therefore we support the AMH and are significantly influenced by it in our work.

2.2.6 Surprise in economics and finance: importance, evidence and computation

While in Section 2.1.3 we addressed surprise from the cognitive science perspective, we review in this section how the word surprise and related terms such as less (more) than expected have been frequently used in economics and finance.

As we have presented earlier, forecasting is one of the essential tasks in the context of economics and finance. Market participants, both professionals (e.g., financial advisors, economists) and individual market participants (investors), need to deal with the difficult task of assessing the occurrence of a wide range of different kinds of future events (Lo and Mueller, 2010). For instance, market participants use different sets of tools to come up with forecasts and estimations for a variety of economic and financial indicators (Banerjee and Marcellino, 2006; Marcellino, 2006) with respect to the behavior of individual companies (e.g., revenue, earnings per share) as well as the behavior of the general economy (e.g., Gross Domestic Product (GDP) and Unemployment Rate). This kind of assessment, though difficult in practice, is important because a series of strategies, like those presented in Section 2.2.2, depend somewhat on the probability of the occurrence of such events as well as on the behavior of economic and financial indicators, which may have a positive or negative influence on asset prices (Taleb, 2007, 2008, 2014).

Market consensus

A common idea used in this context is the idea of “market consensus”, which is traditionally used as an attempt to create a **general market expectation or belief about something**. The so-called market consensus is generated, for example, by economic and financial services companies which periodically survey groups of economists and professionals on a set of indicators in order to compute ex ante estimates. Then, when data is released, it is usually compared to different and relevant market consensus in

order to understand if it is surprising, “lower (higher) than expected”, “beat (miss) the expectation”, or “better (worse) than expected”.

Interestingly, what seems to be relevant for asset behavior and prices are not the released numbers in itself, but actually whether numbers are **surprising** in some way. Generally speaking, a better (worse) than expected data about a relevant indicator may signal that a given company or the economy is behaving better (worse) than the expected. Some indicators such as the Citigroup Economic Surprise Index (CESI) try to gauge that (see (Scotti, 2016) for more details). Considering the rationale used by the fundamental analysis perspective, presented in Section 2.2.2, the basic interpretation is that when the actual data is higher (lower) or better (worse) than the expected, market participants should, according to Efficient Markets Hypothesis (EMH), presented in Section 2.2.3, react and behave in response to this surprising information accordingly (Foster et al., 1984).

Earnings surprise

Some of the most prominent examples come from the research related to earnings surprise (e.g., see (Bartov et al., 2002; Bernard and Thomas, 1989; Brown, 2001, 2003; Chordia et al., 2009; Chordia and Shivakumar, 2006; Copeland et al., 2004; Kasznik and McNichols, 2002; Skinner and Sloan, 2002; Soffer et al., 2000)). The basic rationale for the earnings surprise is the following: a better (worse) than expected earnings per share should be taken into account by, for example, fundamental analysts which will revise and probably increase (decrease) the estimation for the fair price of the company accordingly. Similarly, market participants will adjust their expectations, which will then reflect in the concrete actions of buying or selling which in turn affect asset prices.

The adjustment often results in both price-momentum drift and price drift after earnings announcements (Brown, 2001, 2003; Kothari et al., 2006; Sadka, 2006). The post-earnings-announcement drift (PEAD) (Ball and Brown, 1968) is an anomaly that conflicts with the assumptions claimed by the market efficiency theories such as the EMH (Chordia and Shivakumar, 2006). To briefly present how corporate profits affect asset prices, we show in Figure 2.8 the evolution of the *S&P500* index in comparison with corporate business profits before tax.




FIGURE 2.8: *S&P500* index, and corporate business profits before tax, from July 2005 to July 2015, quarterly adjusted, generated with the Economic Research Federal Reserve Bank of St. Louis site - <http://research.stlouisfed.org/> - and then adapted by us. The shaded area indicates a period of recession in the U.S. economy.

Let us illustrate with a **real world example** the impact of earnings surprise on asset prices. As of 23-07-2015, with the market closed, the e-commerce company Amazon.com, Inc.⁷ announced financial results for its second quarter (April, May and June) ended June 30, 2015⁸. Shortly after the announcement, several media companies⁹ start using terms such as surprise and better-than-expected to refer to key financial variables released such as revenue and earnings per share. Indeed, Amazon released better-than-expected results at least when compared to the expectations of small groups of professionals, the so-called market consensus. The reported earnings per share of \$0.19 for the second quarter was significantly better-than-expected when compared to, for example, the consensus earnings per share estimate of -\$0.15, computed by Zacks Investment Research, based on the forecast of 16 analysts¹⁰ as well as the consensus earnings per

⁷<http://www.amazon.com/> [Accessed: July 2017]

⁸<http://phx.corporate-ir.net/phoenix.zhtml?c=97664&p=irol-irhome> [Accessed: July 2017]

⁹<http://www.nytimes.com/2015/07/24/technology/amazon-earnings-q2.html> [Accessed: July 2017]

¹⁰<http://www.zacks.com/stock/research/AMZN/earnings-announcements> [Accessed: July 2017]

share estimate of $-\$0.14$ ¹¹, compiled by Thomson Reuters I/B/E/S Estimates¹² (the number of analysts was not publicly available). On the next trading day (24-07-2015) Amazon shares surged almost 10% despite of adverse market conditions, i.e., S&P500 index and the NASDAQ Composite fell by more than 1% on this day. Based on this remarkable appreciation, not considering, for the sake of this illustration, the impact of complex market mechanisms (e.g., margin calls, options, etc), we can claim that the results were overall considered by the majority of market participants as a surprise.

Computing earnings surprise

In this context, there are several different methods for computing earnings surprise (e.g., (Kaestner, 2006a,b)). The earnings surprise is calculated by performing some kind of comparison between the actual earnings per share with the consensus estimate. The basic method for computing the Unexpected Earnings (UE_q) is calculated by

$$UE_q = EPS_q / EST_q \quad (2.3)$$

where EPS_q is the actual earnings, EST_q is the consensus estimate (generally the mean or the median of the forecasts) in the time preceding the announcement, and q is the time unit, e.g., quarter. The Unexpected Earnings UE_q can also be calculated by a slightly different method:

$$UE_{q_2} = EPS_q - EST_q \quad (2.4)$$

Also, there are two different and more sophisticated scaling methods for computing the unexpected earnings. The Scaled Unexpected Earnings ($SCUE_q$) is calculated by

$$SCUE_q = UE_{q_2} / abs(EPS_q) \quad (2.5)$$

where abs is the absolute value of the reported Earnings Per Share (EPS_q).

¹¹<http://www.reuters.com/article/2015/07/24/us-amazon-com-results-idUSKCN0PX2HX20150724> [Accessed: July 2017]

¹²<http://financial.thomsonreuters.com/en/products/data-analytics/company-data/ibes-estimates.html> [Accessed: July 2017]

Similarly, the Standardized Unexpected Earnings (SUE_q) is calculated by

$$SUE_q = UE_{q2} / \sigma EST_q \quad (2.6)$$

where σEST_q is the standard deviation of the consensus estimate. The same rationale used to compute earnings surprise can be applied to compute the surprise “felt” by market participants with respect to other micro and macro economic and financial indicators.

2.3 Agent-based Financial Markets

Agent-based Computational Economics (ACE) (Chen, 2012; Tesfatsion, 2002), also referred to in the literature as Agent-based Computational Finance (LeBaron, 2000, 2006), can be thought of as a branch of a wider area known as agent-based modeling (Axelrod, 1997; Bonabeau, 2002; Getchell, 2008; Macal and North, 2005, 2006). ACE is a relatively new field of research that consists in an innovative methodology that has been applied to the investigation of a wide range of economic and financial subjects such as macroeconomics (e.g., (Assenza et al., 2015; Gatti et al., 2010; Lengnick, 2013)), bubbles, crashes, and meltdown modeling (e.g., (Brock and Hommes, 1998; Buchanan, 2009; Giardina and Bouchaud, 2003; Hommes and Iori, 2015; Karino and Kawagoe, 2009; Klimek et al., 2015)). The application of the agent-based methodology to the financial markets has been one of the most prolific areas.

Agent-based financial markets can be seen as a subfield of Agent-based Computational Economics (ACE), that consists in a fertile field to use the agent-based methodology. It offers the opportunity to depart from the conventional thinking provided by traditional economic theories such as the assumption of stable (static) and well-defined preferences, the rationality of agents as well as the believe in an efficient market, assumptions presented in Section 2.2.3. For instance, researchers have the possibility to design, implement, and realize a variety of different experiments with artificial agents that have bounded rationality, behavioral biases, and cognitive judgment and reasoning features and limitations. Besides, alternative hypotheses, such as the Adaptive Market Hypothesis (AMH), presented in Section 2.2.5, can be to some extent

tested as well as interesting, though complicated, phenomena such as emergence (e.g., (Harper and Lewis, 2012)), bubbles, cycles, panics, and trends, can be investigated.

From the agent-based financial market perspective, a financial market (e.g., stock market) can be seen as a socio-economic complex system (e.g., (May et al., 2008; Sterman, 1994)), modeled from bottom up, and simulated as a multi-agent system (Weiss, 1999; Wooldridge, 2002) composed of a large number of heterogeneous artificial agents (Arthur, 1993; Davidsson, 2000; Drogoul et al., 2003; Russell and Norvig, 2009; Wooldridge and Jennings, 1995). Specifically, researchers can create diverse populations (i.e., different “market ecologies” (Farmer, 2002)) of heterogeneous artificial agents (market participants) in terms of strategies (e.g., different weights to fundamental, technical, and noise components), budgets, goals, risk aversion, etc, as well as to consider artificial agents with human-like features such as different memory systems (e.g., different memory sizes), and memory processes (e.g., different levels of forgetfulness), and so forth. Additionally, they can in practice play with such mechanisms and parameters, adjusting and modifying them according to a specific purpose, resulting in an enormous number of possibilities for investigation in terms of rich, novel, and diverse scenarios. More importantly, it offers the possibility of making controlled experiments and simulations, isolating any given parameter, and carrying out as much simulations as necessary in order to observe, and systematically validate how a given parameter may affect general market dynamics, price formation, price discovery, bid-ask spreads, volatility, etc (Cristelli, 2011; Fagiolo et al., 2006, 2007; Janssen and Ostrom, 2006; Windrum et al., 2007). Agent-based financial markets have been considered as a more natural and behaviorally plausible way for describing and simulating the financial markets in comparison with other approaches (e.g., statistical, mathematical, traditional) (Batten, 2000; Buchanan, 2009; Farmer and Foley, 2009; Farmer et al., 2012; Gualdi et al., 2015; Markose et al., 2007).

Agent-based financial markets can play an important role in future economic and financial modeling. It make it possible to address questions that otherwise would be difficult (some of them probably impossible) to answer via other approaches. More specifically, agent-based financial markets have a vast potential to provide real and significant contributions to society (e.g., informing policy makers) by giving novel and realistic insights into how market participants behave individually as well as how a given economic and financial system behave globally. A series of agent-based financial markets works have been providing experimental evidence toward establishing a

foundation that may ultimately turn the field into a truly alternative, and also complementary, paradigm to traditional ones (Bouchaud, 2008; Chakraborti et al., 2011a,b; Farmer and Foley, 2009; Gatti et al., 2010).

2.3.1 Questions in designing artificial agents in the context of Agent-based Financial Markets

In designing artificial agents in the context of agent-based financial markets, the first guideline that should be taken into account is that **artificial agents must be heterogeneous** (Kirman, 2006; Lo, 2004, 2010). Additionally, some key **design questions** can be drawn from the literature (Chen, 2008; LeBaron, 2001, 2006; Wan et al., 2002).

The first design question, and perhaps the most important, is how to represent the preferences of the artificial agent. Preferences include, for example, defining what will be the goal and time horizon, how many assets the artificial agent will be allowed to operate at the same time, which methods will be used to create beliefs and expectations, and ultimately which investment strategies will be employed. Preferences also consist of how much risk the artificial agent will be able to take, and which methods will be used to assess the probabilities regarding the occurrence of potentially good or bad events. Finally, it may include, depending on the objective of the research, not only behavioral biases such as overreaction and loss aversion (presented in Section 2.2.4) but also memory biases such as mood-dependent and mood-congruent memory (presented in Section 2.1.1). Preferences may be static or, in a more realistic representation, dynamic, changing due to modifications in general market conditions.

The second design question is how artificial agents deal with information. For market participants, information is one of the most important issues to be addresses, as presented in Section 2.2.3 and 2.2.6, for example. Therefore, artificial agents should be capable of receiving and processing information in a similar way. Artificial agents with access only to endogenous information resemble market participants who base their decisions solely on technical analysis, as presented in Section 2.2.2. Finally, it is important to decide how agents process and turn the exogenous incoming information into decisions.

The third design question refers to how artificial agents learn from mistakes and evolve. As described in Section 2.2.5, adaptation is not only the better way of achieving

a consistent level of profits but also significantly contributes to achieving the only objective that really matters that is survival. In the dynamic environment of the stock markets, changes can occur very fast demanding rapid reactions of artificial agents who frequently need to adjust and modify their beliefs, expectations, and initial investment strategies, as presented in Section 2.2.2. For example, artificial agents who employ fundamental analysis may need to periodically recalculate the intrinsic/fair value of a given company as a result of new incoming information about changes in underlying factors with respect to the company, which may increase or decrease the fair price. Additionally, artificial agents need to be able to learn from their mistakes, and adapt their preferences accordingly.

2.3.2 Agent-based Financial Markets platforms

In the context of agent-based financial markets, one of the earliest proposals dates from early 1990's with the development of the Santa Fe Artificial Stock Market (SFI) (Arthur, 1994; Arthur et al., 1997; Lebaron, 2002). Following the influential SFI other approaches such as the eAuctionHouse (Wurman et al., 1998) and the eMediator (Sandholm, 2000) were proposed. More recently, we would like to stress the following three proposals. For a comprehensive survey of ABM platforms see (Nikolai and Madey, 2009).

The Penn-Lehman Automated Trading Project (PLAT) (Kearns and Ortiz, 2003) is a broad investigation of algorithms and strategies for automated trading in the financial markets by using the Penn Exchange Simulator (PXS), which is implemented in C and runs on Unix and Linux platforms.

The Agent Exchange (AgEx) (Castro, 2009; Castro and Sichman, 2009)¹³ is a financial market simulation tool for software agents. It was written in Java and based on the JADE platform. AgEx makes it possible the development and communication between autonomous agents who manage portfolios, as well as consists in a framework for the development of new agents.

The Java Auction Simulation API (JASA) (Phelps, 2007)¹⁴ is an extension of the Java Agent-Based Modelling (JABM) toolkit (Phelps, 2012)¹⁵. The JABM toolkit is a Java framework for building multi purpose agent-based simulation models using

¹³<http://agex.sourceforge.net/> [Accessed: July 2017]

¹⁴<https://sourceforge.net/projects/jasa/> [Accessed: July 2017]

¹⁵<https://sourceforge.net/projects/jabm/> [Accessed: July 2017]

a discrete-event simulation framework. JASA make available to researchers in the field of Agent-based Computational Economics (ACE) an easy to use API devoted to the design, implementation, and deployment of high-performance experiments, agents and multi-agent-based simulations. The JASA platform allows the creation of a wide range of different scenarios in economics and finance with a great diversity of built-in functionalities in terms of auction protocols, types of agents, investment strategies, and reports, for example. As an open source software, both JABM and JASA facilitate the introduction of extensions without great effort.

We consider the JABM and JASA platforms as powerful and useful tools by providing researchers with an environment in which they are able to concentrate their efforts on the research questions being addressed rather than on secondary questions. As we will explain later, our cognitive model was designed, implemented, and integrated into JABM and JASA. All the experiments we carried out with agent and multi-agent-based simulations ran with/on JABM and JASA.

2.3.3 Agent-based Models for studying financial markets

A great number of different agent-based models of financial markets have been proposed (e.g., (Hoffmann et al., 2006; Hommes, 2001; Lettau, 1997; Lu et al., 2015; Maslov, 2000; Palmer et al., 1994)). For a comprehensive survey of a variety of agent-based models see (Alfi et al., 2009a,b; Boer-Sorban, 2008; Chakraborti et al., 2011a,b; Lovric, 2011a; Samanidou et al., 2007). Regarding the complexity of the models, for the sake of simplicity, we might divide them essentially into two categories.

On the one hand, **simple models** like the El Farol Bar problem (Arthur, 1994) and the Minority Game (Challet and Zhang, 1997) are good starting point for modeling the financial markets. By capturing some of the features of the financial markets, they offer a set of opportunities of understanding key questions regarding its dynamics (e.g., price formation, emergence).

On the other hand, more **sophisticated models** like the Santa Fe Artificial Stock Market (Arthur et al., 1997; Lebaron, 2002), Caldarelli et al. (1997), Lux and Marchesi (1999, 2000), Chiarella and Iori (2002), and Giardina and Bouchaud (2003), tend to have mechanisms that are more similar to the mechanisms of the financial markets (e.g., order book, auction mechanisms). The majority of the proposed agent-based models,

in spite of their level of complexity or sophistication, have been able to replicate some of the stylized facts of the financial markets (e.g., (Cont, 2007)), especially the heavy/fat tails and volatility clustering facts, presented in Section 2.2.1. As a result, such models have been more or less considered as valid representations of the financial markets. For a comparative analysis of all previous agent-based models, except the model of Chiarella and Iori (2002), in terms of seven categories, “Aim of the model”, “Agents and strategies”, “Number of agents”, “Price formation”, “Origin of Stylized Facts”, “Discussion realism/tractability”, and “Self-Organized Criticality”, see (Cristelli, 2011).

We now detail the following three agent-based models: the El Farol Bar problem introduced by Arthur (1994) as well as the Minority Game proposed by Challet and Zhang (1997), the Lux and Marchesi model (Lux and Marchesi, 1999, 2000), and the order-driven market model proposed by Chiarella and Iori (2002).

El Farol Bar problem

The El Farol Bar problem was introduced in Arthur (1994). In this problem, there is a bar with a fixed capacity and a certain number of clients need to periodically and independently choose between two actions, either go the bar or stay at home. However, if the number of clients who goes to the bar is above or equals its capacity, the bar becomes too crowded and those who attended did not have fun. For example, in his seminal paper (Arthur, 1994), Arthur proposes a bar capacity of 60 and 100 clients. Clients generally make use of a strategy that provides them with a prediction for the next attendance, indicating the best action to take. The only information available is the historical attendance and there is no communication between clients.

Formally, the El Farol Bar problem can be defined as follows. Let C be the bar capacity; n be the number of clients; $PatronAgent_k$ be the representation of each client, where $k = 1$ to n ; $rounds$ be an entity to represent time; $Action = (a_1 = stayAtHome) \oplus (a_2 = goToTheBar)$ be the possible mutually exclusive and exhaustive actions each $PatronAgent_k$ needs to select at each $round_i$; Att_i be the attendance at $round_i$, i.e., sum of all $PatronAgent_k$ who selected the action a_2 ; $Attendance$ be an entity with size equals $rounds$ that stores the historical attendance of each Att_i ; and $Bartender$ be an entity that represents the bar, and is responsible for computing Att_i and storing it in $Attendance$.

We present in **Algorithm 1** the general steps of the El Farol bar problem.

```

C ← 60
n ← 100
Attendance ← ∅
Actions ← {a1 = stayAtHome, a2 = goToTheBar}

foreach roundi, i = 1 to rounds do
  Att ← 0
  foreach PatronAgentk, k = 1 to n do
    | PatronAgentk selects aj ∈ Actions
  end
  Att ← Bartender sums all PatronAgentk who selected a2
  Bartender broadcasts Att to all PatronAgentk

  foreach PatronAgentk, k = 1 to n do
    | PatronAgentk checks whether Att < C
  end
  Attendance[i] ← Att
end

```

Algorithm 1: El Farol Bar problem: basic original algorithm.

We present in Figure 2.9 a typical attendance scenario for the El Farol bar over time for a same run for different time horizons, considering $C = 60$ and $n = 100$.

The context presented by the El Farol has the particular feature of being under a self-denying prophecy regime. A self-denying prophecy (also known as self-defeating prophecy) is the complementary opposite of a self-fulfilling prophecy, presented in Section 2.2.2. When a certain number of market participants believe in a similar “prophecy”, they might behave in a way that their actions collectively contribute to deny the prophecy. More specifically, in the context of the El Farol, for example, if for any given round_i all PatronAgent_k believe in the “prophecy” that the best action to take is goToTheBar, all PatronAgent_k will attend, the bar will become too crowded ($C \geq 60$), and then their actions deny the “prophecy”. In such a paradoxical scenario, agents are somewhat forced to be heterogeneous in the sense that they have to employ, for

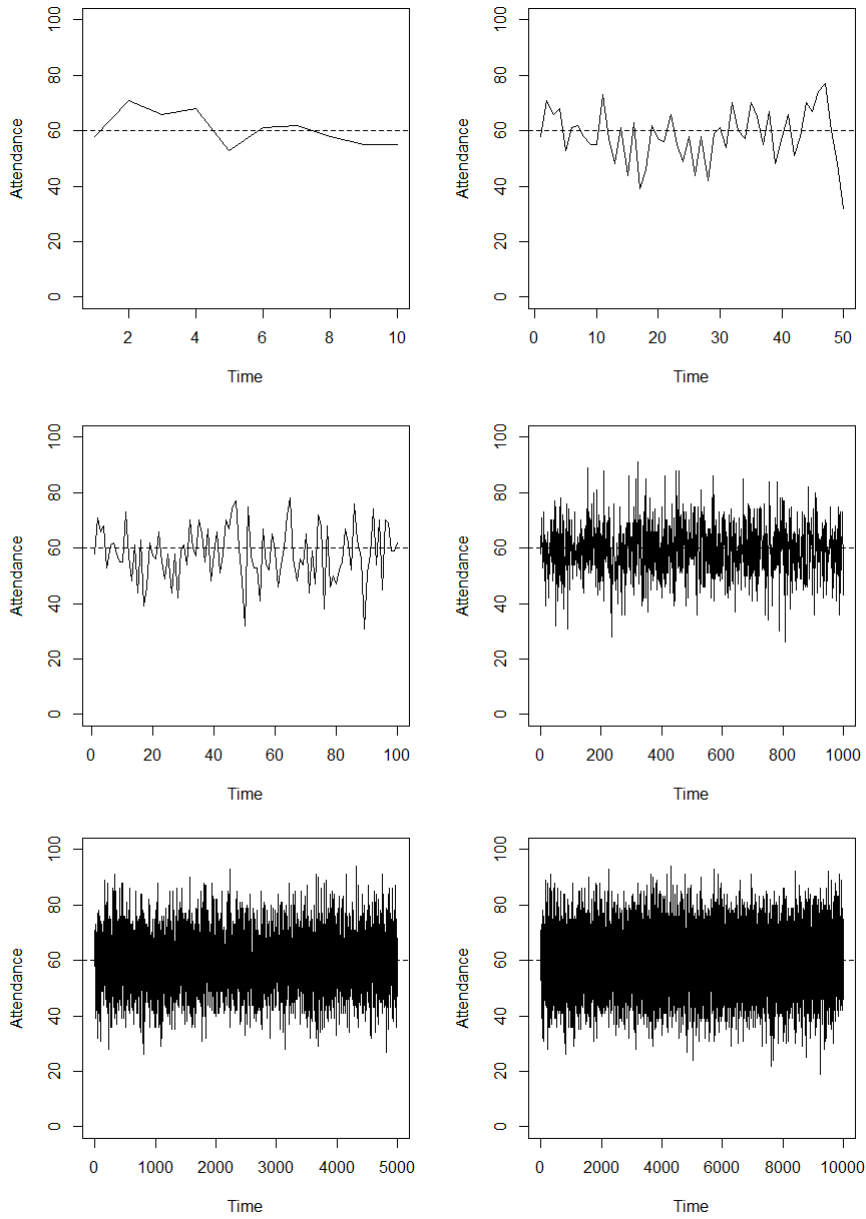


FIGURE 2.9: El Farol bar problem: an example of the dynamics of attendance over time for a same run.

instance, different strategies for creating predictions about the next attendance. Nevertheless, as a strategic environment, the result heavily depends on the choice made by other agents.

The El Farol Bar problem is naturally related to the **Minority Game (MG)** idea introduced by Challet and Zhang (1997). A MG is the most drastic simplification of

the El Farol Bar problem. It is a simple adaptive multi-agent model of the financial markets that can be used to describe the possible interactions of market participants in the financial markets (Challet et al., 2001a, 2000). A MG is an abstraction of situations in which a set of agents is faced with a binary choice. A basic MG is defined as follows. Consider a population of N agents competing in a repeated game, where N is an odd integer. At each round, each agent has to choose between either two actions, typically, +1 or -1 (such numbers can be seen for instance as a “bid” or “sell” action). The minority choices win the game at that round and all the winning agents are rewarded. MGs are believed to be one of the simple adaptive multi-agent model of the financial markets, capturing some key features of a generic market mechanism (e.g., price formation) as well as of the competition between adaptive heterogeneous agents (e.g., inductive reasoning), (e.g., (Brandouy, 2005; Yeung and Zhang, 2008)).

Lux and Marchesi

Different agent-based models somewhat rely on the three components, fundamentalist, chartist (technical), and noise, see (Bak et al., 1997; Boswijk et al., 2007; Chan et al., 1999; Feroni and Agliari, 2008; Takahashi and Terano, 2003). Such components of the agent decision-making process are essentially representations of the investment strategies presented in Section 2.2.2.

Lux and Marchesi (1999, 2000) propose a multi-agent model of financial markets in which the pool of agents is divided into two groups, namely fundamentalists and chartists. Fundamentalists believe in fair/fundamental price (p_f) derived from fundamental analysis and therefore only buy (sell) if and only if the current price (p_t) is below (above) the fundamental price (p_f). Chartists are noise traders whose behavior depends on a sort of herding mechanism and historical prices. They are divided into two subgroups, optimists or trend followers and pessimists or contrarians (Lux, 1998). Although the total number of agents is fixed, the ratio between fundamentalists and chartists may vary during the simulation. The overall market dynamics basically depends on three elements: chartists switching between optimistic and pessimistic opinion, switching between chartist and fundamentalist strategy, and the price formation process.

Order-driven driven market model proposed by Chiarella and Iori

Chiarella and Iori (2002) introduce an order-driven market model with heterogeneous agents trading via a central order matching mechanism in a double auction. Their model, like other models mentioned earlier, essentially rely on the assumption that demands of agents depend on three components: a fundamentalist, a chartist (technical), and a noise-induced component. Chiarella and Iori assume that agents know the fundamental value (p^f) of the asset (which is constant) as well as the history of prices and, for the sake of the simplicity, that agents can only trade one stock at a time. At the moment of deciding on whether to submit an order or not, which will prevail in the interval $\{t, t + \tau\}$ during which the order will be active, each agent makes an expectation about the spot return ($r_{t,t+\tau}^i$), according to

$$r_{t,t+\tau}^i = g_1^i \frac{(p^f - p_t)}{p_t} + g_2^i r_{L_i} + n_i \epsilon_t \quad (2.7)$$

where g_1^i is the weight of the fundamentalist component, g_2^i is the weight of the chartist component, and n_i is the weight of the noise-induced component. Additionally, $g_1^i > 0$, and $g_2^i > 0$ in the case of trend followers and $g_2^i < 0$ in the case of contrarians. Specifically, $g_1^i \sim |\mathcal{N}(0, \sigma_1)|$, $g_2^i \sim \mathcal{N}(0, \sigma_2)$, $n_i \sim \mathcal{N}(0, n_0)$, and $\epsilon \sim \mathcal{N}(0, 1)$. The quantity r_{L_i} is the moving average value of the return over L_i interval, a number between $(1, L_{max})$ uniformly and independently generated. A positive (negative) expected spot return $r_{t,t+\tau}^i$ indicates a long (short) position. Then, the expected future price ($p_{t+\tau}^i$) at time $t + \tau$, is calculated by each agent as

$$p_{t+\tau}^i = p_t e^{r_{t,t+\tau}^i}. \quad (2.8)$$

If an agent expects a price increase (decrease) she/he decides to buy (sell) one unit. They assume that an agent is willing to submit a bid (ask) order to buy (sell) at price b_t^i (a_t^i) lower (higher) than her/his expected future price ($p_{t+\tau}^i$) according to

$$b_t^i = p_{t+\tau}^i (1 - k^i) \quad (2.9)$$

$$a_t^i = p_{t+\tau}^i (1 + k^i) \quad (2.10)$$

where the k_i is uniformly distributed in the interval $(0, k_{max})$, with $k_{max} \leq 1$. Finally, if $b_t^i(a_t^i)$ is smaller (larger) than the current quoted ask (bid), the agent submits a limit order at $b_t^i(a_t^i)$, whereas if $b_t^i(a_t^i)$ is larger (smaller) than the current quoted ask (bid) the agent submits a market order at $b_t^i(a_t^i)$ and a transaction at price $b_t^i(a_t^i)$ occurs. A market order is a buy or sell order to be executed immediately at current market price. A limit order is an order to buy (sell) at no more (at no less) than a specific price (Treleaven et al., 2013).

The model proposed by Chiarella and Iori is able to reproduce many of the complex phenomena observed in real stock markets (some of them presented in Section 2.2.1). Besides, they demonstrated that with their model one is able to shed light on, among other things, how heterogeneity, i.e., different types of agents in terms of fundamentalist, chartist, and noise weights, yields significant different results. For instance, they stressed the role of fundamentalists in stabilizing the market, whereas demonstrated that chartists have the opposite effect. Their results indicate that all three types of agents are to some extent necessary to generate realistic price dynamics.

2.4 Related Work

First of all, it is important to bear in mind that, as we have been stressing, the interest in the understanding of the financial markets is not new. On the contrary, a wide variety of works have been proposed. A significant part of them, perhaps due to monetary incentives, focus on maximizing gains by means of conceiving optimal investment or trading strategies. Although we acknowledge that gains are naturally important in the context of the financial markets, our goal in this work is to explore which contributions the application of the cognitive science perspective might bring to the understanding of the financial markets. Our work involves the realization of agent and multi-agent-based simulations and backtesting with our implementation of the artificial cognitive agents as well as case studies with the proposed cognitive model. Throughout the thesis we present different sets of works which may resemble our work on some aspects. However, in this section we present a subset of the works in the context of the financial markets. This subset focus on Agent-based Financial Markets, backtesting and case studies whose goal is devoted to model and simulate market participants by means of artificial agents who resemble a human market participant in terms of cognition,

emotions, etc. To the best of our knowledge, few works with such features can be found in the literature.

Kodia and colleagues (Kodia and Said, 2009; Kodia et al., 2010a, 2009, 2010b) have been working on a novel way of modeling and simulating the stock market dynamics by essentially considering the cognitive behavior of market participants. A market participant is characterized as an agent with an experience degree that varies from a novice to an expert represented, respectively, by the numbers 0 and 1. Additionally, the behavior of each agent is defined in terms of three pairs of behavioral attitudes namely pessimism/optimism, speculation/caution, and mimetism/leadership. Each pair describes a behavioral attitude that can be either checked or reversed represented, respectively, by the numbers +1 and -1. For example, a optimist speculator mimetic participant and a pessimist, cautious, and leader participant can be represented as having the behavior (+1, +1, +1) and (-1, -1, -1), respectively. In other words, an optimistic agent has a different degree of confidence in a given outcome, a cautious agent examine a large number of messages compared to a speculator, an imitator reproduces the reaction of its surroundings agents, and a leader holds a dominant market position and takes the initiative of buying or selling assets. Broadly speaking, the agent decision-making process is based on cognitive and rational paradigm, biased by behavioral attitudes. Agents receive various kinds of both qualitative and quantitative stimuli which are weighted according to the agent behavioral profile. For example, agents can interact with other agents by giving and receiving advice or opinions of its neighbors. Based on the MadKit platform (Gutknecht and Ferber, 2001), Kodia and colleagues developed a multi-agent based simulation tool named SiSMar (Simulation Stock Market), that is used to understand the influence of psychological character of market participants and its neighborhood on its decision-making, as well as their impact on the market in terms of price fluctuation. They claim that the contribution of their work is to consider the stock market as a social organization of autonomous actors with dependent heterogeneous beliefs and different behavioral attitudes.

Another work that uses a similar approach is described by Pereira et al. (2009). Participants were modeled as cognitive agents which use a sort of BDI architecture. Agents have a mental state consisting of beliefs defined by a belief base, knowledge, and desires, whose behavior is governed by a deliberative process consisting of a set of desire-generation rules. Knowledge refer to information that can be accessed directly by the agent, whereas, on the other hand, beliefs refer to information that cannot be directly

observed namely the future states of affairs. Additionally, agents have access to a set of technical indicators which constitute the only sources of their beliefs. Beliefs, desires, knowledge, and goals are represented by means of fuzzy propositions, i.e., propositions whose degree of truth can assume any value between 0 and 1. The market structure is realistic, but organized in a very simplified way. For example, a single asset is assumed to exist. A fundamental idea which underlie the work of Pereira and colleagues is that under the pressure of competition-driven selection profitable agents proliferate, whereas unprofitable agents become extinct. Therefore, agents evolve by means of an evolutionary algorithm. Unlike other approaches, agents use the same rules but with different parameters for each agent. Additionally, the performance of each agent is compared to the performance of its peers so that 30% of the individuals agents having the lowest performance are eliminated from the simulation.

Lovric (2011b) explores agent-based financial artificial markets as a tool for studying behavioral biases of individual investors. The proposed conceptual model of individual investor brings together different cognitive elements and behavioral phenomena such as risk attitude, time preference, strategies, goals, motivation, emotions, heuristics and biases, personality, which are studied by using agent-based simulations. The aim of model is descriptive in the sense that their goal is be able to capture and describe observed behavioral phenomena. He claims to be confirmed the added value of agent-based financial market models in studying the topics of behavioral finance.

The three previous works share with us the same assumption, i.e., considering market participants as cognitive agents, and therefore modeling and simulating the economic and financial system or problem accordingly. However, both our goal and approach are significantly different. They do not employ neither knowledge related to research on emotions, decision-making, and memories nor rely on any detailed psychological or cognitive emotion theory to support their work. Nevertheless, we agree with them in the sense that this kind of approach consists indeed in a novel and interesting approach for understanding the financial markets.

Chapter 3

TribeCA: a Conceptual Cognitive Model and Implementation

We start this chapter by presenting in Section 3.1 our conceptual cognitive model, named TribeCA (Trading and investing with behavioral-emotional Cognitive Agents), the agent architecture and its main modules, namely Sensors, Reasoning, Goals, Memory, and Actions.

Then, in Section 3.2 we instantiate the TribeCA, describing how it might be used in economics and finance, by presenting **a cognitive modeling approach to the El Farol**. The El Farol bar problem, presented in Section 2.3.3, is a relatively simple but powerful model that can be used as a good starting point for modeling and understanding the financial markets both individually and globally. We then introduce a concept of rationality based on the belief threshold concept (**BT**), and show the cognitive *PatronAgent* algorithm. We also describe in Section 3.2.1 three concepts we used to measure efficiency in our experiments, namely the Cumulative Efficiency (**CuE**), Cumulative Happiness (**CuH**), and Happiness Distribution (**HaD**) as well as briefly present an example of how *CuE*, *CuH*, and *HaD* are calculated.

We present in Appendix A the main classes of both JABM and our proposed cognitive model. To test our conceptual model as well as to create the foundations for the realization of the agent-based experimental evaluation, presented in Chapter 5, we developed an implementation of our conceptual cognitive model and integrated it into

the JASA and JABM tools, presented in Section 2.3.2. As a result, we are able to carry out experiments in economics and finance with artificial cognitive agents.

3.1 TribeCA: a Conceptual Cognitive Model

The TribeCA (Trading and investing with behavioral-emotional Cognitive Agents) model we propose has roots in both the cognitive emotion theories, particularly the Belief-Desire Theory of Emotions (BDTE) (presented in Section 2.1.2), the surprise process proposed by Meyer and colleagues (presented in Section 2.1.3), and the artificial surprise model proposed by Macedo and colleagues (presented in Section 2.1.4).

We show in Figure 3.1 the architecture we adopted for the artificial cognitive agent. As in many other agent architectures, the architecture includes the following modules: Sensors, Reasoning, Goals, Memory and Actions. We detail the main components of our conceptual model.

Agent is an artificial entity that has the capability to perceive and affect the world. An agent has a set of sensors, set of actions, goals, a reasoning mechanism, and memory systems and processes.

Sensors are the perceptual system that allows the agent to receive inputs from the environment.

Actions are the output from the agent to the environment.

Goals can be defined as some action that the agent wants or has to do, while desires are future states of the environment that the agent would like to happen. Goals relate to the possible plan of action by taking into account current situation and past experience.

Reasoning is in the core of the architecture. It receives information both from the internal (from memory) and external world (through the sensors) and outputs an action that has been selected for execution. The first of the steps is concerned with receiving the input through the sensors. Inputs (new beliefs) are then processed by the BBC (Belief-Belief Comparator) and BDC (Belief-Desire Comparator) which, by the memory processes of encoding and retrieving, compare the inputs with the existing beliefs and desires stored in Memory. BBC and BDC empower the agent with the

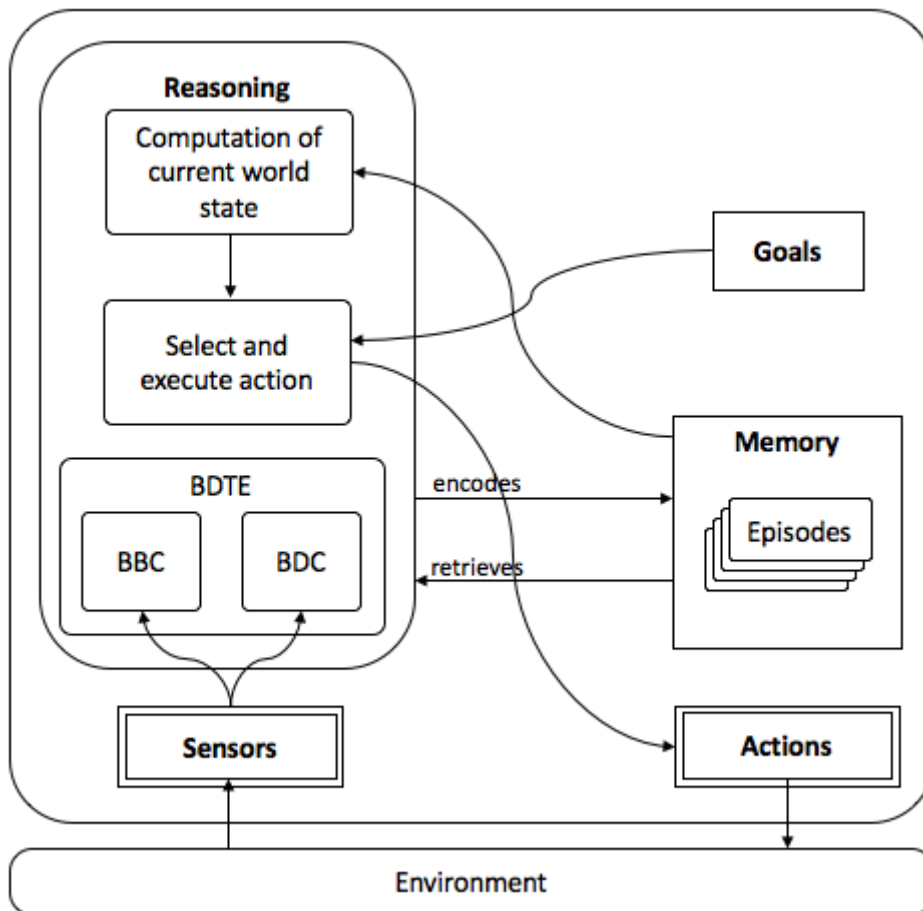


FIGURE 3.1: Architecture of the proposed artificial cognitive agent.

ability to compute emotions. In our architecture, the agent has the capacity of computing the following five emotions: hope, fear, surprise, happiness, and unhappiness emotions. Hope, fear, happiness, and unhappiness are computed based on the BDTE, specifically by applying the quantitative formulation presented in Table 2.2. Similarly, the architecture implements the surprise process defined by Meyer and colleagues, and its computation is carried out by the artificial surprise model proposed by Macedo and colleagues, specifically by Equation 2.1. The second step is the computation of the current world state. This is carried out by generating expectations which are based on the knowledge (episodes) stored in the Memory module. Then, in the third step, based on the expectations generated in the previous step, the agent selects and executes the action that is considered to be the best to achieve its goal(s).

Memory of the agent stores information about the world. Inspired in the human

episodic memory, presented in Section 2.1.1, the memory of the agent includes a set of episodes that may encode and store a variety of information. Specifically in this work, the belief, the set of possible outcomes regarding the belief, what really happened, and the set of emotions “felt” by the agent. We compute emotions in two different moments. Before receiving the input, the agent might feel either hope or fear based on whether the agent has a desire in favor of or against to a given future outcome, respectively. When the agent receives the input it might feel surprise, and then happiness or unhappiness concerning the input. Therefore, the set of emotions for each episode consist of three different emotions, namely hope, fear, surprise, and either happiness or unhappiness with its respective values.

3.2 A cognitive modeling approach to the El Farol

To illustrate how the TribeCA can be used in economics and finance, we present in this section a possible cognitive modeling approach to the El Farol bar problem based on the TribeCA. Let us start by briefly reviewing some key concepts of the Belief-Desire Theory of Emotions (BDTE), one of the three pillars of the TribeCA, presented earlier.

A proposition p is represented as a tuple $\langle S, B, D \rangle$ where S is the mental language expressing the proposition p , B and D are quantities representing, respectively, the agent’s degree of belief and desire regarding proposition p . The strength of a belief in a proposition p at time t , is defined as $b(p, t)$, where $b(p, t) \in \mathbb{R}$ and $0.0 \leq b(p, t) \leq 1.0$, where 1.0 denotes certainty that p , 0.5 maximal uncertainty, and 0.0 certainty that not p . Similarly, the strength of a desire about a proposition p at time t , is defined as $d(p, t)$, where $d(p, t) \in \mathbb{Z}$ and $-100 \leq d(p, t) \leq +100$, where positive values denote desire in favor of p , negative values denote desire against p , and 0 denotes indifference. A new belief is the belief or fact that agents receive basically through its sensors.

The Belief-Belief Comparator (BBC) and the Belief-Desire Comparator (BDC) are responsible for comparing each newly acquired belief to all pre-existing beliefs and desires, respectively, looking for a match or a mismatch. BDTE defines emotions as products or signals produced by the BBC and BDC. For example, suppose an agent has $b(p, t) = 0.9$ and $d(p, t) = +80$. When the agent receives a new belief at time t regarding proposition p , the BBC yields a belief-confirmation signal, since what the

agent believed at t_{-1} as most likely really happened. Similarly, the BDC yields a desire-fulfillment signal, since the agent has a desire in favor of p .

In addition to the capacity of computing emotions through the Belief-Belief Comparator (BBC) and the Belief-Desire Comparator (BDC), our agents are able to deal with previous knowledge with respect to whether the **current strategy succeeded or failed** and use such knowledge to calculate its current belief in the strategy. A strategy succeeds when it indicates the correct action, i.e., if the strategy predicted that the bar will be crowded (or not) and it turned out to be crowded (or not) the action indicated by the strategy was the correct (wrong) one.

According to the BDTE, the El Farol bar problem scenario can be modelled as follows. The belief in the fact that “**My strategy works**” is defined as $b(p,t)$, the belief in proposition p at time t , formally, where 1.0 denotes certainty that the strategy really works, 0.5 denotes maximal uncertainty that is the agent does not know whether the strategy works or not, and 0.0 denotes certainty that the strategy does not work. The $b(p,t)$ is calculated by a Bayesian process considering the memory size (MS) of the agent. In this work, we opted for using the BayesNet algorithm available in the data mining software Weka¹. The MS is used to store the whole experience, consists of both positive and negative instances, of the agent in using the current strategy in a given number of last rounds, limiting how many rounds the agent is able to “remember”. This limitation is carried out by means of a simple sliding window mechanism, forcing the agent to “forget” the oldest instances. Therefore, the belief in the correctness of the strategy will change, either increasing or decreasing, because of a change in the memory of the agent.

Practically, on the one hand, a $b(p,t) \geq 0.5$ means that the agent has some degree of belief in the fact that its current strategy works and so it makes sense to a “rational” agent to maintain using the current strategy. On the other hand, a $b(p,t) < 0.5$ means that the agent has some degree of belief in the fact that its current strategy does not work and so it makes sense changing to a new strategy. Therefore, we rely on the concepts of the BDTE to derive a concept of rationality. We define a belief threshold (BT) by which the agent must change its current strategy. While a BT higher than 0.5 can be seen as a behavioral bias regarding the level of correctness of a strategy, a BT lower than 0.5 can be seen as a behavioral bias or some psychological resistance

¹<http://www.cs.waikato.ac.nz/ml/weka/>

with respect to changing its current strategy. A BT lower than 0.5 can be thought of as a conceptualization of the loss realization that leads to a general disposition to sell winners too early and hold losers too long. An agent is considered to act rationally when the belief threshold (BT) of 0.5 is fully respected, whereas an agent is considered to act not so rationally when he/she does not fully respect the rational BT of 0.5. Finally, if $b(p,t) == 0.5$ the agent maintains using the current strategy, and when an agent starts using a strategy its initial $b(p,t) = 0.5$. Similarly, according to the BDTE, we assume all agents have a truly desire in favor of p , meaning the strategy to work, i.e., $d(p,t) = +100$.

For all agents, the basic algorithm is as follows:

- (1) At the start of the simulation, the agent randomly selects a strategy and removes it from the set of available strategies (which will be detailed later in this Chapter);
- (2) for each round, the agent stores the positive or negative experience of using the strategy in its memory and updates its $b(p,t)$ by increasing or decreasing it via a Bayesian process (i.e., BayesNet algorithm available in the data mining software Weka);
- (3) computes the happiness according to the BDTE, meaning that happiness is in this work defined as the desire of the agent, i.e., +100 if the strategy provides the correct action and;
- (4) if and only if $b(p,t) < BT$, the agent changes its current strategy by randomly selecting one of the remaining strategies in the set. When the agent has tested all the available strategies, all strategies are added back to the set and the algorithm backs to step (1). In this situation, the agent has no prior memory of using a given strategy of the set.

We detail in **Algorithm 2** the Cognitive *PatronAgent* rationale.

```

Actions  $\leftarrow$  {a1 = stayAtHome, a2 = goToTheBar}
b(p, t)  $\leftarrow$   $\alpha$ ; BT  $\leftarrow$   $\beta$ ; MS  $\leftarrow$   $\gamma$ ; Memory  $\leftarrow$   $\emptyset$ ;
Strategies = {Strategy1, ..., Strategyn};
StrategiesTemporary  $\leftarrow$  Strategies
CurrentStrategy  $\leftarrow$  strategyi  $\in$  StrategiesTemporary, i randomly
Removes CurrentStrategy from StrategiesTemporary
foreach roundi, i = 1 to rounds do
    | ForecastedAttendance  $\leftarrow$  forecast given by CurrentStrategy computation
    | if ForecastedAttendance < C then
    |   | SelectedAction  $\leftarrow$  a2
    | else
    |   | SelectedAction  $\leftarrow$  a1
    | end
    | Att  $\leftarrow$  real attendance from Bartender
    | if (Att < C) == (SelectedAction = a2) then
    |   | isStrategyCorrect  $\leftarrow$  TRUE
    | else
    |   | isStrategyCorrect  $\leftarrow$  FALSE
    | end
    | if isStrategyCorrect then
    |   | Increase b(p, t) via Bayesian process (BayesNet algorithm)
    | else
    |   | Decrease b(p, t) via Bayesian process (BayesNet algorithm)
    | end
    | if i > MS then
    |   | Removes Memory[1]
    | end
    | Memory[i]  $\leftarrow$  isStrategyCorrect
    | if b(p, t) < BT then
    |   | if StrategiesTemporary !=  $\emptyset$  then
    |     | Removes CurrentStrategy from StrategiesTemporary
    |   | end
    |   | if StrategiesTemporary ==  $\emptyset$  then
    |     | StrategiesTemporary  $\leftarrow$  Strategies
    |     | Removes CurrentStrategy from StrategiesTemporary
    |   | end
    |   | CurrentStrategy  $\leftarrow$  strategyi  $\in$  StrategiesTemporary, i randomly selected
    |   | Removes CurrentStrategy from StrategiesTemporary
    | end
end
end

```

Algorithm 2: Cognitive *PatronAgent* algorithm.

3.2.1 Cumulative Efficiency, Cumulative Happiness, and Happiness Distribution

In addition to the concept of rationality, efficiency is also a key issue in our work. Although we are aware of the difficulty both on defining what efficiency means as well as on how to measure it, issues presented in Section 2.2.3, we rely on the concepts of the BDTE, particularly on the *BT*, and on three basic concepts proposed by Baccan and Macedo (2013a,b), namely the **Cumulative Efficiency (*CuE*)**, the **Cumulative Happiness (*CuH*)**, and **how happiness is distributed among agents (*HaD*)**.

The first is the Cumulative Efficiency (*CuE*) that is the cumulative sum of all attendances below the bar capacity. The idea is to quantify how efficient the resource, that is in this case the bar, has been used by the artificial agents. The *CuE'* at round i is calculated as

$$CuE'_i = \sum_{k=1}^i E_k, \quad i = 1, 2, \dots, n \quad (3.1)$$

where $attendance_i$ is attendance at round i , $E_k = attendance_k$ if $attendance_k < 60$, $E_k = 0$ otherwise, and n is the number of rounds.

The second is the Cumulative Happiness (*CuH*) that is the cumulative sum of all happiness “felt” by the artificial agents. An artificial agent “feel” happiness when its strategy works, i.e., when it indicates the correct action. The *CuH'* at round i is calculated as

$$CuH'_i = \sum_{k=1}^i \left(\sum_{j=1}^{na} HappinessAgent_{jk} \right), \quad i = 1, 2, \dots, n \quad (3.2)$$

where $HappinessAgent_{jk}$ is the happiness “felt” by agent j at round i , $HappinessAgent_{jk} = 100$ if strategy succeeds, $HappinessAgent_{jk} = 0$ otherwise, na is the number of artificial agents, and n is the number of rounds.

The third is **how happiness is distributed among artificial agents (*HaD*)** that is the sum of all rounds in which each artificial agent “felt” happiness. The *HaD* for each agent j is calculated as

$$HaD_j = \sum_{i=1}^n HappinessAgent_{ji}, \quad j = 1, 2, \dots, na \quad (3.3)$$

where $HappinessAgent_{ji}$ is the happiness “felt” by agent j at round i , na is the number of artificial agents, and n is the number of rounds.

Whereas CuH tries to measure happiness from a global perspective, HaD tries to measure happiness from an individual perspective. From the economics perspective, happiness can be thought of as wealth and, therefore, CuH and HaD can be considered measures that try to quantify individual wealth and how fair the wealth is distributed, respectively.

***CuE*, and *CuH* optimal values**

To illustrate how $b(p,t)$, CuE , CuH , and HaD work, consider the following example. Assume that there are the following two groups of artificial agents: G1 consists of 59 agents using a fixed strategy that indicates the action go to the bar; and G2 consists of 41 agents using a fixed strategy that indicates the action stay at home. In this example, the attendance is 59, below the bar capacity, and therefore the right action to take would be go to the bar. Therefore, for each round, while all artificial agents of G1 would “feel” happiness, all artificial agents of G2 would “feel” unhappiness. Practically, in the first round, $CuE_1 = 59$ and $CuH_1 = 5900$ ($59 \cdot 100$), in the second round, $CuE_2 = 118$ ($59 + 59$) and $CuH_2 = 11800$ ($5900 + 5900$) and so forth.

We would like to stress that this is the scenario that provides **optimal results in terms of CuE and CuH** and that such optimal values are used by us as references for calculating and plotting the results of CuE and CuH . The optimal CuH to CuE ratio is 100.

Formally, the optimal $CuE_{optimal}$ is calculated as

$$CuE_{optimal_i} = \sum_{k=1}^i 59, \quad i = 1, 2, \dots, n \quad (3.4)$$

Similarly, the $CuH_{optimal}$ is calculated as

$$CuH_{optimal_i} = \sum_{k=1}^i 5900, \quad i = 1, 2, \dots, n \quad (3.5)$$

Then, the **Cumulative Efficiency** (CuE) is calculated as

$$CuE = CuE' / CuE_{optimal} \quad (3.6)$$

Similarly, the **Cumulative Happiness** (CuH) is calculated as

$$CuH = CuH' / CuH_{optimal} \quad (3.7)$$

From the perspective of the use of the resource (i.e., CuE), we can claim that artificial agents have been using the bar in a quite good and efficient way, as well as agents are quite happy globally (i.e., CuH). However, in this case, HaD values would exhibit a high standard deviation, meaning that the distribution of happiness among artificial agents is not indeed “fair”. Finally, as the number of rounds increases, $b(p,t)$ of G1 (G2) agents would be close to 1.0 (0.0), meaning that they “firmly believe” its strategy works (does not work). The better the strategy, the higher the $b(p,t)$.

Chapter 4

Traditional Modeling versus Cognitive Modeling of Surprise in Economics and Finance

The word surprise and terms somewhat related to surprise such as “higher (lower) than expected” have been extensively used in the context of economics and finance, as presented in Section 2.2.6. In this chapter we present two different perspectives on modeling surprise in economics and finance, namely the traditional (economic and financial) perspective and the cognitive perspective.

We present two case studies that show how the cognitive science perspective of surprise presented in Chapter 2 may be applied to model, simulate, and thus better understand some key economic and financial concepts such as market consensus and risk management. We demonstrate how the cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997) and the artificial surprise model proposed by Macedo et al. (2004), two core components of the TribeCA, can be applied to compute artificial surprise.

Case Study I, presented in Section 4.1, focuses on modeling surprise with respect to the idea of market consensus. Market consensus, as explained in Section 2.2.6, is a concept widely used in economics and finance, particularly in the task of forecasting, that tries to measure the general market expectation or belief about a future relevant event. For this case study, we selected the Unemployment Rate indicator, compute

the surprise by using the economic and financial approach and the cognitive science approach, and compare the results.

Case Study II, presented in Section 4.2, brings together the task of risk management with the cognitive science perspective on surprise. Specifically, we carried out a case study with an artificial cognitive agent using a popular risk management tool known as VaR (Value-at-Risk) historical. Our goal is to better understand and then inform market participants (human agents) about the risks as well as about the intensity of surprise they may feel and be able to cope with it in using VaR (Value-at-Risk) historical under two different regimes, namely a Calm period and a Crash period.

We consider these two case studies as a modest contribution, though a good starting point, towards truly understanding the surprise “felt” by market participants in economics and finance.

4.1 Case Study I: Surprise in Market Consensus

As presented in Section 2.2.6, market consensus is a widely used concept in economics and finance. The so-called market consensus tries to capture the general expectation of market participants with respect to some indicators. Such indicators in this context may include key economic and financial indicators such as the CPI (Consumer Price Index), GDP (Gross Domestic Product), and Unemployment Rate as well as company indicators such as earnings, and earnings per share.

We carried out a case study to **compare the computation of surprise from the perspective of economics and finance to the perspective of cognitive science in respect to “market consensus” regarding forecasts**. To this end, we obtained the data related to the forecasts free of charge from the Wall Street Journal (WSJ) Economic Forecasting Survey¹. The WSJ monthly surveys a group of nearly 50 economists on more than 10 major economic indicators with respect to the US economy such as GDP, CPI, and Unemployment Rate. The number of economists who actually participate may vary from month to month. For this case study we selected the Unemployment Rate, a well known relevant economic indicator, as briefly presented in Section 2.2.1.

¹<http://projects.wsj.com/econforecast/> [Accessed: July 2017]

Similarly, we obtained the real/released data regarding the Unemployment Rate from the Economic Research Federal Reserve Bank of St. Louis², also free of charge.

We present in Table 4.1 the data relevant for this case study. Each month economists (forecasters) are asked to give their estimation on what might be the Unemployment Rate for the next June and December. Therefore, we compiled the WSJ monthly survey data from June 1, 2003 to June 1, 2015 taken by semesters in a total of 25 observations, and compute the mean and standard deviation of the forecasts for each observation. We then compare the forecasts in the month preceding the real data (May and November) with the real data released in the beginning of the subsequent month (June and December).

Economic and financial approach and cognitive science approach

We employed the following two different approaches in this case study: an economic and financial approach and a cognitive science approach.

The **economic and financial approach** is straightforward. We applied the earnings surprise rationale, described in Section 2.2.6, to compute the unexpectedness or surprise “felt” by market participants with respect to the Unemployment Rate (UR). First, we computed the Unexpected Unemployment Rate (UUR) as

$$UUR_m = UR_m - EST_m \quad (4.1)$$

where m refers to the monthly periodicity, UR_m is the real/released Unemployment Rate, and EST_m is the consensus estimate for month m . For this case study we assume the mean of the forecasts as the consensus estimate in the month preceding the announcement.

Second, we computed the Scaled Unexpected Unemployment Rate ($SCUUR$)

$$SCUUR_m = \frac{UUR_m}{abs(UR_m)} \quad (4.2)$$

²<http://research.stlouisfed.org/> [Accessed: July 2017]

TABLE 4.1: Unemployment Rate: date, number of forecasts, consensus estimate (mean), standard deviation of the forecasts, and real Unemployment Rate date, in a total of 25 observations.

	Unemployment rate			
	Forecasts			Real data
Date	Number	Consensus (EST_m)	Standard deviation (σEST_m)	UR_m
05-2003	53	5.9691	0.1625	6.30
11-2003	52	6.0019	0.0843	5.70
05-2004	54	5.5600	0.0852	5.60
11-2004	53	5.3574	0.1039	5.40
05-2005	55	5.1768	0.0738	5.00
11-2005	55	5.0196	0.0942	4.90
05-2006	55	4.6772	0.0709	4.60
11-2006	55	4.5375	0.1054	4.40
05-2007	59	4.5634	0.0912	4.60
11-2007	52	4.7694	0.0778	5.00
05-2008	51	5.2227	0.1149	5.60
11-2008	53	6.7815	0.1705	7.30
05-2009	51	9.2481	0.2330	9.50
11-2009	51	10.2654	0.2649	9.90
05-2010	56	9.6737	0.1553	9.40
11-2010	47	9.5583	0.1998	9.30
05-2011	54	8.7873	0.1441	9.10
11-2011	51	9.0096	0.1241	8.50
05-2012	49	8.0960	0.1399	8.20
11-2012	48	7.7816	0.1269	7.90
05-2013	48	7.5122	0.1053	7.50
11-2013	43	7.1500	0.1439	6.70
05-2014	46	6.3277	0.1077	6.10
11-2014	42	5.7295	0.0978	5.60
05-2015	53	5.3389	0.1156	5.30

where UUR_m is given by Equation 4.1, abs is the absolute value of the real/released UR_m , though in this case study there will be only positive values for the underlying indicator, since there is not a negative Unemployment Rate. Nevertheless, we keep the abs function in order to maintain it in line with the original formula.

Third, we computed the Standardized Unexpected Unemployment Rate ($SUURE$)

$$SUURE_m = \frac{UUR_m}{\sigma EST_m} \quad (4.3)$$

where UR_m is given by Equation 4.1, and σEST_m is the standard deviation of the forecasts.

The **cognitive science approach** we consider is based on the artificial surprise model proposed by Macedo and colleagues, presented in Section 2.1.4. Let (Ω, A, P) be a probability space where Ω is the sample space (i.e., the set of possible outcomes of the event), $A = \{A_1, A_2, \dots, A_m\}$, is a σ -field of subsets of Ω (also called the event space, i.e., all the possible events), and P a probability measure which assigns a real number $P(F)$ to every member F of the σ -field A . Let $E = \{E_1, E_2, \dots, E_n\}$, $E_i \in A$, be a set of mutually exclusive events in that probability space with probabilities $P(E_i) \geq 0$, such that $\sum_{i=1}^n P(E_i) = 1$. Let E_h be the highest expected event from E . **The intensity of surprise about an event** E_g , defined as $S(E_g)$, is calculated as

$$S(E_g) = \log_2(1 + P(E_h) - P(E_g)) \quad (4.4)$$

where E_h is the event with the highest probability in the set. In each set of mutually exclusive events, there is always at least one event whose occurrence is unsurprising, namely E_h .

Unlike the straightforward economic and financial perspective, the cognitive science perspective requires the creation of the outcomes and the estimation of its subjective probabilities. For instance, in this case study we rely on the assumption that the EST_m follows a Normal distribution. Therefore, based on the 3-sigma rule, presented in Section 2.2.1, we are interested in computing the artificial surprise when the UR_m violates the ± 2 standard deviation of the mean forecast, regardless of a violation higher than -2 or higher than $+2$ standard deviation, i.e., $\mu \pm 2\sigma$. Therefore, we assume that about 95% of the UR_m values are within $EST_m \pm 2\sigma EST_m$. In other words, when the artificial agent is confronted with the real data (UR_m) one of the following two possible outcomes occurs: UR_m is (1) either within $EST_m \pm 2\sigma EST_m$ or (2) higher than $EST_m \pm 2\sigma EST_m$. Additionally, the belief of the artificial agent in the first outcome and in the second outcome is, respectively, 0.95 and 0.05. Regarding the artificial surprise, the artificial agent will “feel” no surprise since what the he/she considered as more likely (outcome 1) happened or, conversely, surprise with a high intensity since what he/she considered as less likely (outcome 2) happened. More precisely, the surprise

about outcome 2 would be 0.9259994, i.e., $S(E_g) = \log_2(1 + 0.95 - 0.05)$. Figure 4.11 presents the intensity of surprise (0.9259994) “felt” by the artificial agent over time.

4.1.1 Results and discussion

Regarding the economic and financial approach, Figure 4.1 shows the dynamics of the forecasts regarding Unemployment Rate and the number of forecasts per month (right axis) versus the real Unemployment Rate over time (left axis). Figure 4.2 shows some statistical indicators regarding forecasts, namely the mean, mode, median, highest and lowest forecast for each month. Similarly, Figures 4.3 and 4.4 show, respectively, the real data concerning the Unemployment Rate versus the mean forecast as well as the real data concerning the Unemployment Rate to mean forecast ratio. Figures 4.5, 4.6, and 4.7 show, respectively, Equations 4.1, 4.2, and 4.3. Likewise, Figure 4.12 shows the comparison between Equations 4.1, 4.2 and 4.3.




FIGURE 4.1: Unemployment rate (left) versus real data versus number of forecasts (right).

Regarding the cognitive science approach, Figure 4.9 shows the mean forecast along with its ± 2 standard deviation versus the real Unemployment Rate data. Figure 4.10 shows when the real Unemployment Rate data is higher or lower than the standard deviation of the mean forecast. Figure 4.11 shows the surprise intensity according




FIGURE 4.2: Unemployment rate forecasts: some statistical indicators.




FIGURE 4.3: Unemployment rate: real data versus mean forecast.




FIGURE 4.4: Unemployment rate: real data to mean forecast ratio.




FIGURE 4.5: Unexpected Unemployment Rate (Equation 4.1).




FIGURE 4.6: Scaled Unexpected Unemployment Rate (Equation 4.2).




FIGURE 4.7: Standardized Unexpected Unemployment Rate (Equation 4.3).




FIGURE 4.8: Comparison of the dynamics of Equations 4.1, 4.2 and 4.3.

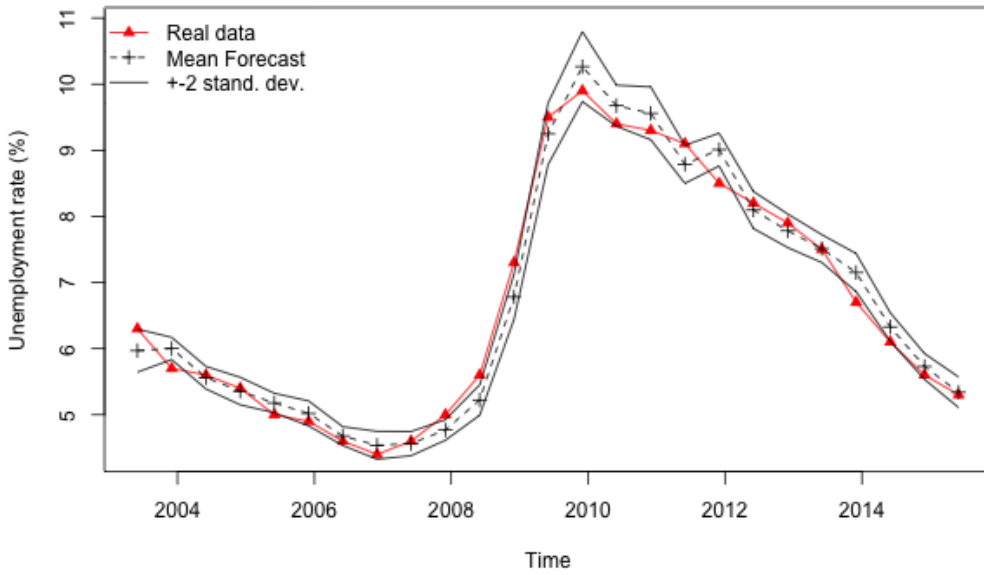


FIGURE 4.9: Unemployment rate: mean forecast, standard deviation, and real data.

to artificial surprise model proposed by Macedo and colleagues. Last but not least, Figure 4.12 shows the comparison between Equations 4.1, 4.2 and 4.3 with the surprise intensity according to the artificial surprise model proposed by Macedo and colleagues (Equation 4.4)

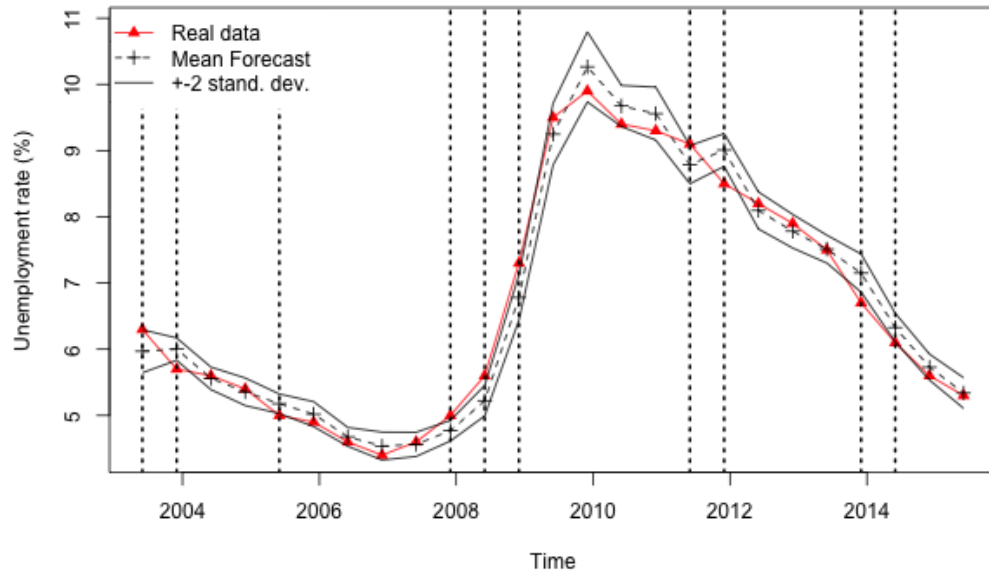


FIGURE 4.10: Real data is higher either than +2 standard deviation or -2 standard deviation of the mean forecast, and therefore there is a violation of the 0.954 confidence level.

First of all, it is important to bear in mind that the success of this case study does not rely on the use of a particular economic or financial indicator. Our focus is on comparing the computation of surprise with respect to market consensus from two different perspectives, namely the cognitive science perspective with the economics and finance perspective. Instead of using the Unemployment Rate, a well known important economic indicator, as briefly presented in Section 2.2.1, we could indeed have selected other indicators (e.g., GDP) without any impact on our results. Therefore, our goal in this case study is not to provide evidence neither in favor of nor against to the use of any particular economic or financial indicator.

The cognitive science perspective, presented throughout Section 2.1, essentially relies on the subjective probabilities of the artificial agents with respect to the possible outcomes to compute the artificial surprise (Equation 4.4). In this case, the higher the




FIGURE 4.11: Surprise intensity according to the artificial surprise model proposed by Macedo and colleagues (Equation 4.4).




FIGURE 4.12: Comparison of the dynamics of Equations 4.1, 4.2 and 4.3 with the surprise intensity (vertical dashed line) according to the artificial surprise model proposed by Macedo and colleagues (Equation 4.4).

difference between the subjective probabilities/beliefs, the higher may be the surprise. Additionally, assuming that surprise is a neutral emotion, the artificial surprise model proposed by Macedo and colleagues, presented in Section 2.1.4, computes the surprise intensity “felt” by artificial agents, not qualifying whether the intensity is positive or negative for the artificial agent.

The economic and financial perspective, in its turn, can be thought of as a pure mathematical approach. For example, Equation 4.3 uses the standard deviation to try to capture whether forecasters diverge or converge in their forecasts as well as to represent whether the surprise is positive or negative, however, it requires an understanding of the related indicator. For example, a lower than expected Unemployment Rate is a good signal, whereas a lower than expected GDP is a bad signal.

For the creation of the outcomes and the estimation of its subjective probabilities, we rely on the assumption that the EST_m follows a Normal distribution. We consider that relying on this assumption does not have any impact on our results, especially taking into account that we used this rational as a step of the process of creating the so-called market consensus. Despite that, one should indeed be careful in assuming the Normal distribution pattern in other contexts, since there is substantial evidence that financial asset returns are not normally distributed (Lo and Mueller, 2010). Therefore, other methods would ideally include mechanisms such as self-reports so that forecasters should be able to express their confidence in their forecasts. With this information one may be able to weight the forecasts more precisely. Another approach may include the analysis of different financial instruments regarding a given asset (e.g., derivative instruments such as options) to try to infer the beliefs of market participants by observing their actions.

Considering the cognitive science perspective, what we are able to say is that when data regarding an indicator is released, it may indeed elicit surprise but only to a given set of persons, the so-called market consensus. Surprise is a subjective process that essentially belongs to the individual. What surprises me may not surprise you, because our memories, beliefs, and expectations are certainly different. Therefore, we believe that the term surprise should be more carefully used in the context of economics and finance. Instead of carrying meaningful and relevant information, terms such as “beat (miss) market expectation”, “lower (higher) than expected”, seem to carry just “noise”, may influence the behavior of market participants in a detrimental way and, in the end, does not seem to make much sense, at least from this perspective.

Polls versus Prediction Markets: surprise and cognitive surprise and UK referendum on the UK membership of the EU

Let us now apply the same rationale to other kinds of market consensus, namely opinion polls and prediction markets, by illustrating the case of the recent United Kingdom (UK) referendum on the United Kingdom's membership of the European Union³. As of June 23, 2016 the UK citizens were consulted on whether the UK should remain a member of the European Union or leave the European Union. More precisely, we compare the ex-ante expectations (prior to the referendum) obtained through two different mechanisms, namely opinion polls and prediction markets, to the referendum result as well observe the behavior of the financial markets when the official result was released.

On the one hand, different opinion poll trackers^{4,5,6} were pointing out small advantages (roughly of 2%) for one choice against the other, with the remain choice winning in the majority of the opinion polls. Furthermore, the analysis of the poll of polls make it possible to claim that at that moment there was a draw, i.e., 50% of the opinions in favor of the remain choice whereas 50% of the opinions in favor of the leave choice. Statistically and mathematically speaking, considering aspects traditionally used in polls such as margin of error, it might be claimed that either the remain choice or the leave choice could win.

On the other hand, prediction markets were forecasting with great likelihood the victory of the remain choice. Prediction markets (Berg et al., 2008; Horn et al., 2014; Snowberg et al., 2012), also known as decision markets, such as the Iowa Electronic Market⁷, are a special case of behavioral markets designed specifically to aggregate information and to forecast uncertain future events such as political contests, sporting events, and economic and financial outcomes and macroeconomic variables. Prediction markets operate on the same principle as the pari-mutuel horse racing. Broadly speaking, considering a claim, i.e., a statement that something will happen by a specific future date, bettors will bet (generally real money) for or against this claim. For instance, considering data from the prediction market BetFair⁸, on election day a bet for the remain choice would

³<https://www.eureferendum.gov.uk/> [Accessed: July 2017]

⁴<http://www.economist.com/blogs/graphicdetail/2016/06/britain-s-eu-referendum> [Accessed: July 2017]

⁵<http://www.bloomberg.com/graphics/2016-brex-it-watch/> [Accessed: July 2017]

⁶<http://graphics.wsj.com/brexit-whos-voting-what/> [Accessed: July 2017]

⁷<http://tippie.uiowa.edu/iem/> [Accessed: July 2017]

⁸<https://betting.betfair.com/politics/brexit/> [Accessed: July 2017]

pay 1.13 whereas a bet against it (for the leave choice) would pay 8.00. In other words, the bettor would receive 1.13 (8.00) times the money invested in the remain (leave) choice in case of victory of the remain (leave) choice. These ratios therefore represent the aggregated market's expectations (or market's beliefs) with respect to the claim. In this example, on election day, the quantitative subjective beliefs of market participants (or in this case the market expectation) regarding the victory of the remain choice was 88%, i.e., $1 - (1.13/9.13)$, whereas the victory of the leave choice was 12%, i.e., $1 - (8.00/9.13)$. Predictions markets quickly incorporate new information, and are highly effective in predicting different types of events (e.g., (Ray, 2006)).

From the artificial surprise model of Macedo and colleagues, the UK referendum event might be modelled as follows: there are two mutually exclusive and exhaustive outcomes, namely (E_1) as the remain choice and (E_2) as the leave choice. From the opinion polls perspective, specifically the poll of polls scenario described earlier, the surprise about either event E_1 or E_2 would be 0, i.e., $S(E_i) = \log_2(1 + 0.50 - 0.50)$, no surprise at all, regardless of the outcome. However, from the prediction markets perspective, the surprise about event E_2 would be 0.8155, i.e., $S(E_2) = \log_2(1 + 0.88 - 0.12)$. Interestingly, from the cognitive science perspective, when there is maximal uncertainty, i.e., outcomes with roughly the same quantitative subjective belief, the surprise would be "felt" with very low intensity, roughly 0.0.

As of June 24, 2016, 51.9% of the UK citizens decided that the UK should leave the European Union, the event that turned out to be known as Brexit. On that day, the Brexit sent the financial markets into turmoil. The British pound lost more than 8% in relation to the United States of America dollar, plunging to the lowest value since 1985. The Stoxx 600⁹, an index that represents the 600 large, mid and small capitalization companies across 18 countries of Europe, generally considered the equivalent to the S&P500, lose 7% in the biggest drop since the financial crisis resulted from the Crash of 2008. The S&P 500 lost 3.59%. New York crude oil retreats 6.8%. The VIX (also known as "the fear index") (Fleming et al., 1995; Whaley, 1993, 2000, 2009) soared more than 50%. Market participants flew to quality assets as they became more risk averse, e.g., Gold price surges as much as 8.1% on twice the average daily volume. Both the opinion polls and the prediction markets were wrong in predicting that the remain choice would win. Based on the behavior of the financial markets immediately after the results, we may claim that the Brexit scenario has not been priced in by

⁹<https://www.stoxx.com/index-details?symbol=SXXP> [Accessed: July 2017]

market participants. On the contrary, the Brexit seems to have been received as a huge surprise by market participants, particularly to those that were relying on the prediction markets.

In the context of forecasting there are several problems such as overoptimism biases (Ackert et al., 2008; Lejarraga et al., 2016; Trueman, 1994). One of these problems seems to reside in the situation when everyone agrees about something (Makarov and Rytchkov, 2012). On the one hand, when opinions excessively converge, due to either uncertainty, herd, or the madness of mobs effect, anything that contradicts the market belief may surprise market participants as well as have an exaggerated effect on markets. On the other hand, when opinions diverge, this diversity creates a healthy and stable market ecology resulting in a kind of equilibrium.

4.2 Case Study II: Surprise in Risk Management

As presented in Section 2.2.1, market participants agents in the context of the financial markets need to deal with a series of different kinds of risks. In this context, the **Value-at-Risk (VaR) tool is one of the most popular financial risk measures**, used by financial institutions all over the world (Berkowitz and O'Brien, 2002; Gouriéroux et al., 2000; Kawata and Kijima, 2007). The objective of the VaR tool is to measure the probability for a significant loss in a portfolio of financial assets (Alexander and Sarabia, 2012). Generally speaking, it is assumed that for a given time horizon t and a confidence level p , VaR is the loss in market value over the time horizon t that is exceeded with probability $1 - p$ (Christoffersen and Pelletier, 2004; Duffie and Pan, 1997). For example, suppose a period of one-day ($t = 1$) and a confidence level p of 95%, the VaR would be 0.05 or 5% the critical value. There are several different methods for calculating VaR. For instance, let us briefly present the following two methods to calculate VaR, namely the statistical approach and the historical approach (Jorion, 2006).

The VaR statistical assumes that the historical returns respect the Efficient Market Hypothesis (EMH), presented in Section 2.2.3. The EMH in turn assumes that the series of historical financial returns are Gaussian, with the average value μ of zero and constant variance of σ^2 , i.e., returns $\sim \mathcal{N}(0, \sigma^2)$. Based on the EMH assumptions as

well as on the Gaussian characteristics, it is possible to compute the VaR statistical for a confidence level p of 99% and 95% that are -2σ and -3σ , respectively.

The VaR historical is an alternative way to calculate VaR by ranking the historical returns from the smallest to the highest. Suppose that the series of T returns are r_1, \dots, r_t , we define that this series of returns are ranked if $r_1 \leq r_2 \leq \dots \leq r_T$. In this case, the VaR historical is the return on the position integer $((1 - p) T)$. More precisely, Algorithm 3 describes how the VaR historical is computed as a function of the series of returns (window size) and the confidence level (p). For instance, we present in Figure 4.13 the following scenario: consider the S&P500 index, presented in Section 2.2.1, a begin as of 01-03-2004, four window sizes generally used in the VaR historical computation, namely the 50 (from begin to 10-05-2004), 125 (from begin to 26-08-2004), 250 (from begin to 24-02-2005), and 500 (from begin to 22-02-2006) values, and two confidence levels (p) equal 0.99 and 0.95. Just for a time reference, a financial market year typically consists of 252 trading days.

```
probs ← seq(0, 1, 0.01)
windowSize ← returnPeriodicity[di : dj]
quantiles ← quantile(windowSize, probs)
varP95estimation ← quantiles[6]
varP99estimation ← quantiles[2]
```

Algorithm 3: $VaR_{historical}$: Value-at-Risk (VaR) historical function.

We consider the VaR historical as the most appropriate method for this case study for several reasons. First, because the VaR historical is a conceptually simple method, widely used by practitioners (see for example (David Cabedo and Moya, 2003; Halbleib and Pohlmeier, 2012; Hendricks, 1996)). Second, because it is an easy to understand measure that computes the estimation based on historical data and a confidence level (we will explain later in Section 4.2.1 the particular importance of the confidence level for this case study). Last but not least, because, unlike the VaR statistical, VaR historical is free of the assumption about the Gaussian distribution of the series of returns (Jorion, 2006) and, therefore, relies on more realistic assumptions, as previously presented in Section 2.2.1.

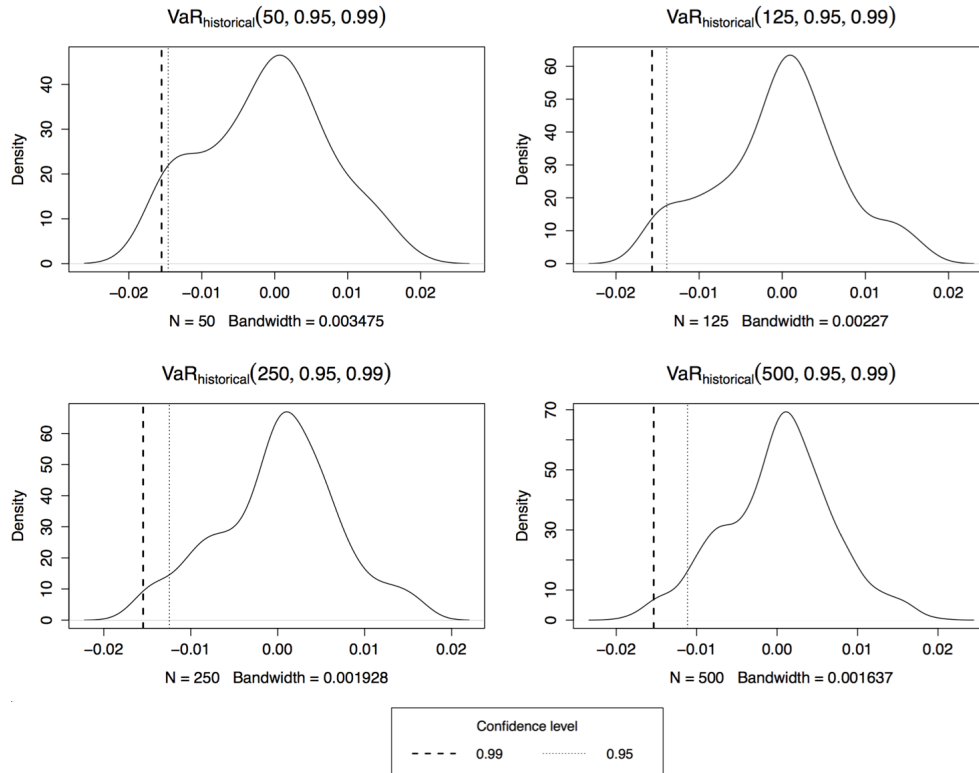


FIGURE 4.13: Example of a $VaR_{historical}$ computation considering daily returns, a begin as of 01-03-2004, four window sizes namely 50 (from begin to 10-05-2004), 125 (from begin to 26-08-2004), 250 (from begin to 24-02-2005), and 500 (from begin to 22-02-2006) values, and two confidence levels (p) equal 0.99 and 0.95.

4.2.1 Cognitive surprise of using Value-at-Risk historical

The goal of this case study is to **compute the cognitive surprise “felt” by an artificial agent relying on Value-at-Risk (VaR) historical under two different market scenarios, namely a Calm period and a Crash period.** The use of an artificial cognitive agent make it possible a better understanding of the risks market participants (human agents) face as well as informing market participants about the intensity of surprise they may be able to cope with in using VaR historical.

We applied a **method** similar to the work of Halbleib and Pohlmeier (2012) to divide the S&P500 index, presented in Section 2.2.1, from 26-11-1990 to 01-07-2009 into two parts, a **Calm period** from 26-11-1990 to 31-08-2008 (total of 4478 days), and a **Crash period** from 01-09-2008 to 01-07-2009 (total of 210 days). The main motivation for this division is the occurrence of the events that lead to the financial crisis of 2008,

particularly the bankruptcy of the investment bank Lehman Brothers in the middle of September, 2008. The data was obtained free of charge from Yahoo Finance¹⁰.

Figures 4.14 and 4.15 show the S&P500 index daily close and daily return of the whole period, respectively. Figures 4.16 and 4.17 present the daily close and the daily returns of the S&P500 of the Calm period, respectively. Similarly, Figures 4.18 and 4.19 present the daily close and the daily returns of the S&P500 of the Crash period, respectively.




FIGURE 4.14: Whole period: S&P500, daily close.

Rolling window size approach and alpha approach

To compute the VaR historical, we employed the following two different approaches: a Rolling window size approach and an Alpha approach.

The **Rolling window size approach** consists in specifying a rolling window size ($window_z$) and a confidence level (p), which combination is used by the VaR historical to compute the estimation by considering the most recent returns. The windows size set contains the 50, 125, 250, and 500 values. The confidence level set contains the 0.95 and 0.99 values. The combination of the window size values with the confidence

¹⁰<http://finance.yahoo.com/> [Accessed: July 2017]




FIGURE 4.15: Whole period: S&P500, daily return.




FIGURE 4.16: Calm period: S&P500, daily close.




FIGURE 4.17: Calm period: S&P500, daily return



FIGURE 4.18: Crash period: S&P500, daily close.




FIGURE 4.19: Crash period: S&P500, daily return.

level values result in eight different treatments. **Algorithm 4** describes the steps of a single simulation of the Rolling window size approach. Similarly, Figure 4.20 presents a schematic representation of the Rolling window size approach.

```

violationsP95  $\leftarrow \emptyset$ 
violationsP99  $\leftarrow \emptyset$ 
i  $\leftarrow$  according to an uniform distribution function, select i within the period (for
the Crash period we fixed i equivalent to the day 01-09-2008)
j  $\leftarrow 0$ 
foreach dk, k = i + 1 to i + 210 do
    dbeginTrain  $\leftarrow i - 1 - window_z$ 
    dendTrain  $\leftarrow i - 1$ 
    windowSize  $\leftarrow dailyReturn[d_{beginTrain} : d_{endTrain}]$ 
    violationsP95[j]  $\leftarrow dailyReturn[d_k] < VaR_{historical}(windowSize, p = 0.95)$ 
    violationsP99[j]  $\leftarrow dailyReturn[d_k] < VaR_{historical}(windowSize, p = 0.99)$ 
    i  $\leftarrow i + 1$ 
    j  $\leftarrow j + 1$ 
end

```

Algorithm 4: Rolling window size approach: algorithm for a single simulation.

The **Alpha approach** consists in specifying a decay function so that recent daily returns gain more weight as opposed to old daily returns. Decay functions are generally

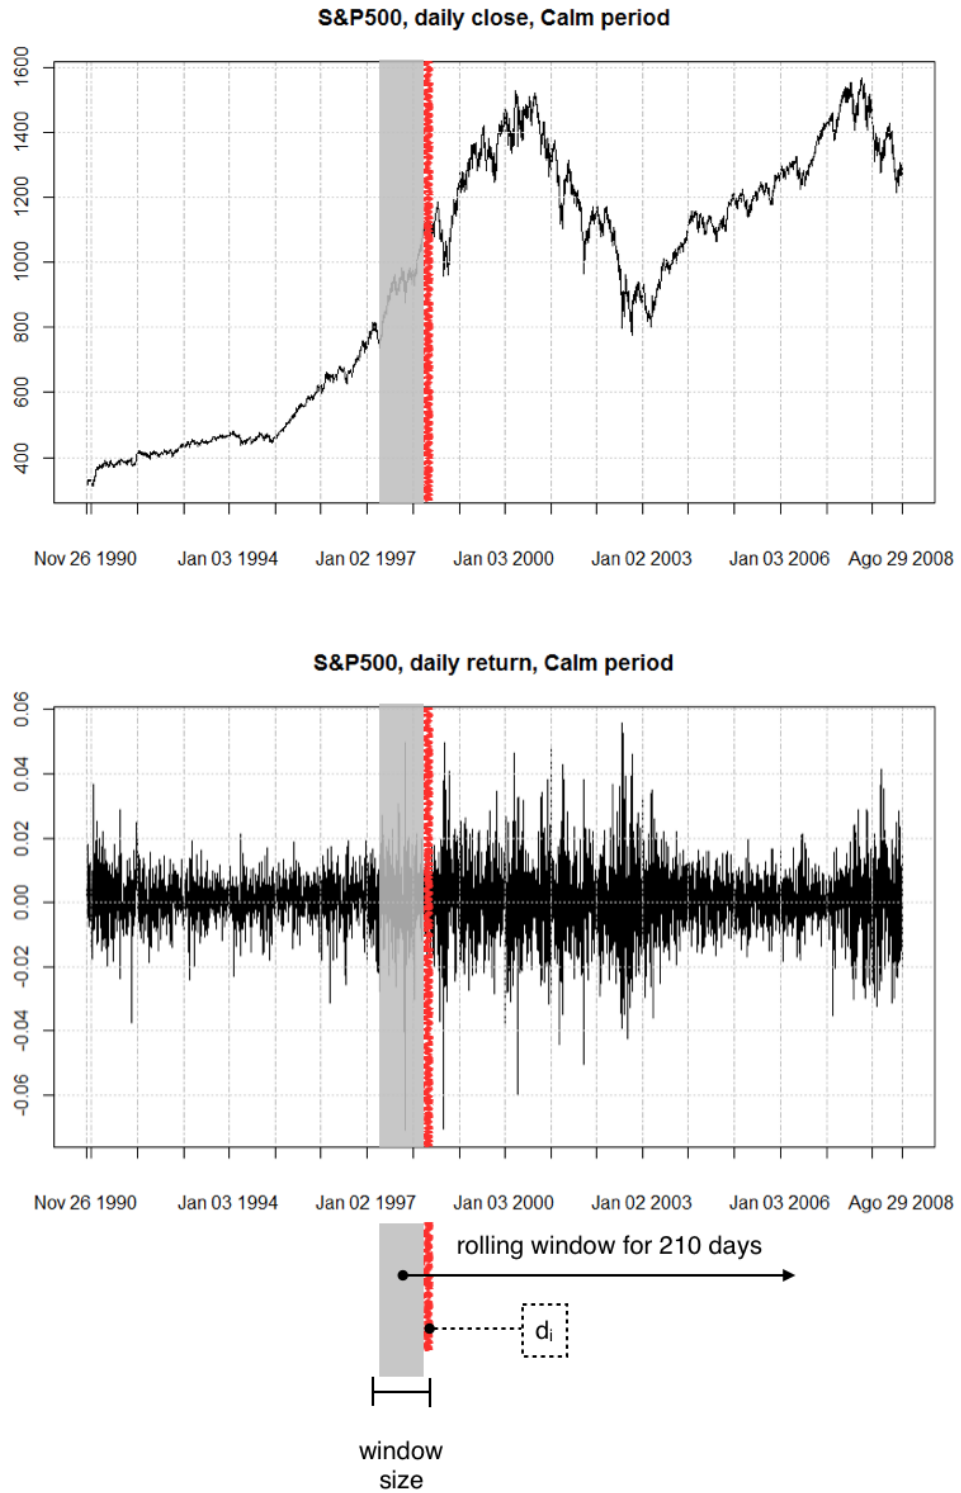


FIGURE 4.20: Rolling window size approach: a schematic representation for the Calm period for a given simulation.

used with the objective of emulating to a certain extent the human memory process of “forgetting” as well as to contemplate some findings on how humans use past experience in decision-making, as presented in Section 2.2.4, which indicate that in revising their beliefs, people tend to overweight recent information and underweight prior information (e.g., (Griffin and Tversky, 1992)). The alpha set contains the 0.995, 0.99, 0.97, 0.94 values. A higher (lower) alpha implies a smaller (higher) level of “forgetfulness”. Unlike the Rolling window size approach, we opted to use just a confidence level of 0.99. The reason is that we observed in our initial experiments that the combination of a confidence level of 0.95 with the previous alpha values caused the agent to “forget” too many returns, generating in the end a quite low and poor VaR historical estimation. Therefore, in the alpha approach we have four treatments. **Algorithm 5** describes the steps of a single simulation of the Alpha approach.

```

violationsP99  $\leftarrow \emptyset$ 
i  $\leftarrow$  according to an uniform distribution function, select i within the period (for
the Crash period we fixed the i equivalent to the day 01-09-2008)
foreach dk, k = i + 1 to i + 210 do
    dbeginTrain  $\leftarrow$  i - 1 - windowz
    dendTrain  $\leftarrow$  i - 1
    windowSize  $\leftarrow$  dailyReturn[dbeginTrain : dendTrain]
    alphaWeights  $\leftarrow$  generate alpha weights as a function of windowSize
    adjWindowSize  $\leftarrow$  windowSize * alphaWeights
    violationsP99[j]  $\leftarrow$  dailyReturn[dk] < VaRhistorical(adjWindowSize, p =
0.99)
    i  $\leftarrow$  i + 1
    j  $\leftarrow$  j + 1
end

```

Algorithm 5: Alpha approach: algorithm for a single simulation.

Artificial surprise rationale

Let us now describe how we address the question of modeling surprise in the context of the **artificial surprise**, presented in Section 2.1.4. We essentially applied the concepts, ideas, and method presented by Bacca et al. (2014b), described in the first case study. This work can be thought of as a continuation and expansion of our initial work. We assume, for the sake of the experiment and simplicity, the confidence levels *p* (0.99,

and 0.95) as the subjective belief of the artificial agent in the accurateness of the VaR historical estimation. By making this assumption and considering a higher subjective belief, we are empowering the artificial agent with a “firm believe” in the accurateness of the VaR historical estimation. So, suppose an event E_g as VaR historical estimation that can assume two mutually exclusive events, meaning that it can be either correct (E_1), i.e., daily return is not lower than estimation, or incorrect (E_2), i.e., daily return is lower than estimation.

In other words, the artificial agent will “feel” no surprise since what the artificial agent considered as more likely, i.e., a correct VaR estimation (E_1), happened or, conversely, surprise with a high intensity since what he/she considered as less likely, i.e., an incorrect VaR estimation (E_2), happened. More precisely, for the confidence level of 0.99 the surprise about event E_2 would be 0.9855004, i.e., $S(E_2) = \log_2(1+0.99-0.01)$. Similarly, for the confidence level of 0.95 the surprise about event E_2 would be 0.9259994, i.e., $S(E_g) = \log_2(1+0.95-0.05)$. For each day d_k we stored the violations of the VaR historical estimations, i.e., if the daily return of $d_k \leq \text{VaR}_p$, so that later we are able to compute the cognitive surprise, $surprise_k$.

In the end we have a sequence $\{surprise_1, \dots, surprise_{210}\}$. Afterwards, for each simulation, we compute the cumulative sum of the surprise. It means that we generate a sequence of 210 elements as a result of the partial sums $surprise_1$, $surprise_1 + surprise_2$, $surprise_1 + surprise_2 + surprise_3$, and so forth. In the case of the Calm period, we added the cumulative sum of the surprise of each treatment and then average it by the number of simulations. The cumulative sum and the normalization make it easier the observation of surprise over time as well as the comparison of surprise between the Calm and the Crash period. The average cumulative sum of the surprise for a given treatment is presented in the next figures for the Calm period.

This case study takes into account the resulting **twelve different treatments**, eight treatments from the Rolling window size approach and four treatments from the Alpha approach. We ran 10^4 independent simulations for each one of the twelve treatments for the Calm period and one simulation for each treatment for the Crash period, since the begin day ($d_i = 01-09-2008$) is fixed for this period.

4.2.2 Results and discussion

We present in Figures 4.21 and 4.22 the density and the box plot of daily returns of the Calm period and of the Crash period, respectively. Additionally, Table 4.2 presents some statistical indicators regarding the Calm and Crash period. Figure 4.23 and 4.24 show the behavior of all the Rolling window size treatments and the Alpha treatments, respectively, for the Calm period. Similarly, Figure 4.25 and 4.26 show the behavior of the Rolling window size treatment and Alpha treatment, respectively, for the Crash period.




FIGURE 4.21: Calm and Crash period: S&P500, density of daily returns.

TABLE 4.2: Statistical data about the Calm and Crash period.

	Calm period	Date	Crash period	Date
Min	-0.0711275	26-11-1990	-0.0946951	02-09-2008
1st Qu.	-0.0046386		-0.0189830	
Median	0.0004558		-0.0002665	
Mean	0.0003125		-0.0015464	
3rd Qu.	0.0055558		0.0149512	
Max	0.0557443	29-08-2008	0.1095720	01-07-2009
Sd	0.0100892		0.03077633	
Kurtosis	6.693523		4.448898	
Skewness	-0.1080146		-0.01451907	




FIGURE 4.22: Calm and Crash period: S&P500, box plot of daily returns.




FIGURE 4.23: Rolling window size approach: treatments for the Calm period.




FIGURE 4.24: Alpha approach: treatments for the Calm period.




FIGURE 4.25: Rolling window size approach: treatments for the Crash period with VaR_{95} and VaR_{99} .




FIGURE 4.26: Alpha approach: treatments for the Crash period.

Figure 4.27 and 4.28 show the comparison of different Rolling window treatments with VaR₉₅ and with VaR₉₉, respectively, for the Calm period versus the Crash period. Similarly, Figure 4.29 shows the comparison of different Alpha treatments for the Calm period with the Crash period.

First of all, it is important to bear in mind that the financial markets are by their very nature complex and dynamic systems. Such complexity significantly increases when we take into account the sophisticated and complex human cognitive and decision-making mechanism, as presented in Section 2.1.

As a result, as others (e.g., (Taleb, 2014)), we consider that the task of risk management in finance and economics is indeed quite difficult. Therefore, we do not have in this case study the goal of providing evidence neither in favor of nor against to a particular risk management tool or system. Instead, our goal is to compute, in a systematic and clear method, the cognitive surprise “felt” by an artificial agent relying on the VaR historical tool during two different periods, namely a Calm period and a Crash period.

Let us then begin the discussion by analyzing **some characteristics of the Calm period in comparison with the Crash period**. We can see in Figure 4.14 and




FIGURE 4.27: Comparing different Rolling window size treatments with VaR_{95} for the Calm and Crash period.




FIGURE 4.28: Comparing different Rolling window size treatments with VaR_{99} for the Calm and Crash period.




FIGURE 4.29: Comparing different Alpha treatments for the Calm and Crash period.

Figure 4.15 that daily returns seem, as expected and in line with the existing literature, to reproduce some statistical regularities that are often found in a large set of different assets and markets known as stylized facts (presented in Section 2.2.1). Specifically, we can observe in Table 4.2 that returns do not follow a Gaussian distribution (returns are not normally distributed), seem to exhibit what is known as fat-tails, as well as to reproduce the volatility clustering fact, i.e., high-volatility events tend to cluster in time. We can see in Figure 4.22 and in Figure 4.21 that the comparison between the daily returns of the Calm period with the daily returns of the Crash period allows us to claim that the daily returns for the Crash period seem to exhibit fat tail distributions. We can also see in Figure 4.18 that during the Crash period the SP&500 index depreciated almost 50% in value (from 1300 points to less than 700 points) in a short period of time (from 31-08-2008 to the beginning of March, 2009).

We now turn our attention to the analysis of the cognitive surprise. For the **Calm period**, we can observe in Figure 4.23 that for the Rolling window size treatments, the memory size of 50 is the one that yields a higher surprise, regardless of the confidence level p . Regarding the Alpha treatments, we can see in Figure 4.24 that the higher the α , the lower the surprise. For the **Crash period**, we can observe in Figure

4.25 the cumulative surprise of the Rolling window size treatments with the confidence level of 0.95 and 0.99. We can see that the lower the window size, the lower the surprise. Regarding the Alpha treatments, we can see in Figure 4.26 that, similarly to the behavior for the Calm period, the higher the *alpha*, the lower the surprise. It may be explained by the fact that the Crash period is indeed a period in which the volatility is high. Therefore, a VaR historical based on a larger window size takes more time to adapt itself to this new and changing environment in comparison with a VaR historical based on a smaller window size.

In **comparing the Calm period with the Crash period**, we can observe in Figure 4.27 and in Figure 4.28 that for the Rolling window treatments the cumulative surprise is, as expected, higher during the Crash period in comparison with the Calm period, regardless of the confidence level. Similarly, we can observe in Figure 4.29 that for the Alpha treatments the cumulative surprise is higher during the Crash period in comparison with the treatment using the same configuration for the Calm period.

Interestingly, if we assume that *alpha* somewhat emulates the human memory process of “forgetting” and considering that a lower *alpha* implies a higher level of “forgetfulness”, we may argue that market participants should be careful in forgetting the past, at least in the context of the financial markets, since the cognitive surprise “felt” by an artificial agent under this treatment is significantly higher than the Rolling window size treatments.

Generally speaking, each day the artificial agent “felt” surprise represents a failed VaR historical estimation. Therefore, the higher the surprise, the wrong a given VaR historical treatment is. As a result, the analysis of the results indicates that it may be quite difficult for a human agent (market participant) not to feel surprise when relying on VaR historical, regardless of the resulting VaR configuration in terms of window size, *alpha*, confidence level, etc. Indeed, there are several issues with models like the VaR historical that take into account historical data for computing their estimation. These models somewhat assume the past is a good indicator of what may happen in the future. In other words, they somehow consider that history repeats itself through time. However, this inductive reasoning often underestimate the probability of rare events and extreme events (Hertwig et al., 2004) and, consequently, underestimates the level of risk.

Additionally, an essential flaw of this kind of rationale is not to truly acknowledge that the “absence of evidence is not evidence of absence” and the existence of “unknown unknowns” (Rumsfeld, 2002). As an example of this kind of rationale, consider, for instance, the so-called turkey paradox (Taleb, 2008). There is a butcher and a turkey. Every day for some period, let us say 100 days, the butcher feeds the turkey. As time goes by, the turkey increases its belief in the fact that in the next day it will receive food from the butcher. However, at a given day, for the “shock” and “surprise” of the turkey, instead of being feed by the butcher, the butcher kills the turkey. The same inductive analogy may be applied to the black swan scenario (Taleb, 2007). The reasoning that all swans someone have seen are white therefore all swans are white is essentially wrong, since the mere existence of one black swan contradicts the core premise of the reasoning, refuting it completely. Some market participants tend to rely on a similar reasoning by investing or trading a variety of financial instruments that yields small but regular gains, until eventually the day a single operation blows all the gains, possibly resulting in significant losses (Barberis, 2013; Taleb, 2014).

Chapter 5

Agent-based Experimental Evaluation with the El Farol

In this chapter we present the agent-based experimental evaluation we carried out with the implementation of the TribeCA, our conceptual cognitive model, presented throughout Chapter 3, with the El Farol bar problem. As presented in Section 2.3.3, the El Farol bar problem is an agent-based model for studying the financial markets. Despite of being apparently a simple model, the El Farol is a good starting point towards the understanding of the behavior of both individual market participants and the financial markets as a whole.

We start Section 5.1 by describing the common features of the agent-based experiments we carried out. We then detail in each of the following sections the three experiments.

Experiment I, presented in Section 5.2, presents our initial experiment with twelve different treatments in terms of number of strategies available, the ability of changing the current strategy or not, and different memory size (MS), to observe the behavior of the artificial cognitive agents we modelled both individually and globally.

Experiment II, presented in Section 5.3, focuses on the impact of rationality (i.e., the belief threshold (BT)) on market efficiency (CuE , CuH , and HaD). In seven different treatments we gradually increase the BT by a factor of 0.05 from 0.20 to 0.50.

Finally, **Experiment III**, presented in Section 5.4, investigates the impact of the memory size (MS) of the cognitive agent on market efficiency (CuE , CuH , and HaD).

In thirty-four different treatments we gradually increase the MS from 10 to 2500 (all rounds of the experiment).

5.1 Agent-based Experiments with the El Farol

All experiments and their respective treatments share the following common initial configuration: initial $b(p,t) = 0.5$, $d(p, t) = +100$, the number of artificial agents is 100 and the capacity is of 60 (both consistent with the seminal paper on the El Farol (Arthur, 1994)), and the number of rounds is 2500 (a number we consider to be sufficient to understand the general dynamics). Additionally, we present in Table 5.1 the set of strategies available.

TABLE 5.1: Set of available strategies for all agent-based experiments and treatments.

Strategy	Acronym	Description
Noise Trader Strategy	<i>NTS</i>	Generates a forecast, uniformly chosen between 0 and 100.
Simple Moving Average Strategy	<i>SMAS</i>	Generates a forecast by computing the simple moving average with a given window size, uniformly chosen between 2 and 100.
Exponential Moving Average Strategy	<i>EMAS</i>	Generates a forecast by computing the exponential moving average with a given window size. In this strategy, alpha, uniformly chosen between 0.01 and 0.99, allows recent values referring to the attendance to gain more weight as opposed to old values.
Opposite Strategy	<i>OPS</i>	Generates a forecast that is equals the complementary attendance of the last round (i.e., forecast = 100 - attendance of the last round).
Same Strategy	<i>SAS</i>	Generates a forecast that is equals the attendance of the last round.
Lagged Strategy	<i>LAS</i>	Generates a forecast that is equals the attendance of a given past round, uniformly chosen between 1 and 5.

As presented in Section 2.2.2, the *NTS* captures the essence of the noise or zero intelligence idea, whereas the *SMAS*, *EMAS*, *OPS*, *SAS*, and *LAS* are essentially strategies based on technical aspects. We used the Mersenne Twister algorithm (Matsumoto

and Nishimura, 1998), a high-quality pseudo-random number generation, to draw all random variables for the *NTS*, *SMAS*, *EMAS*, and *LAS*.

For each treatment we have run 100 independent simulations. We added up the data from all simulations and divided them by 100, resulting in one time series for each one of the variables to be observed (e.g., *CuE*, *CuH* and *HaD*). Additionally, *CuE*, *CuH* were also divided by the optimal *CuE*, *CuH*, respectively, so that the figures throughout this chapter indicate the results of the treatments as a percentage of the optimal, i.e., *CuE*, % of optimal and *CuH*, % of optimal. **Algorithm 6** summarizes the *Cognitive PatronAgent* initial configuration.

```

C ← 60
n ← 100
rounds ← 2500
Actions ← {a1 = stayAtHome, a2 = goToTheBar}
Strategies = {NTS, SMAS, EMAS, OPS, SAS, LAS}
b(p, t) ← 0.5;
    
```

Algorithm 6: Cognitive *PatronAgent* initial configuration, common for all experiments and treatments.

5.2 Experiment I: Revisiting the El Farol

We start by presenting our initial experiment to explore how the artificial cognitive agents we modelled behave in the context of the El Farol problem. The experiment consists of twelve treatments separated into three groups. In the first group (G1), from Treatment 1 (T1) to Treatment 4 (T4), artificial agents have all strategies available (i.e., *NTS*, *SMAS*, *EMAS*, *OPS*, *SAS*, *LAS*), and do not have the ability to change current strategy. In the second group (G2), from Treatment 5 (T5) to Treatment 8 (T8), artificial agents have all strategies available as well as the ability to change current strategy according to $b(p, t)$ and *BT*. In the third group (G3), from Treatment 9 (T9) to Treatment 12 (T12), artificial agents can only use the *NTS* and, like the artificial agents of the first group, do not have the ability to change current strategy. For each group we vary the memory size (*MS*) from 60, 100, 500, to 2500. We summarize in Table 5.2 the features of the experiment we conducted.

Besides the *CuE*, *CuH*, and *HaD*, for this experiment we use two other metrics namely the Global Belief in Strategy (*GBS*) and the Global Surprise (*GSu*).

The **Global Belief in Strategy (*GBS*)** at round i is calculated as

$$GBS_i = \sum_{j=1}^{100} BeliefAgent_{ji}, \quad i = 1, 2, \dots, n \quad (5.1)$$

where $BeliefAgent_{ji}$ is the $b(p, t)$ (i.e., the belief of the agent in the correctness of the current strategy) of the agent j at round i , j is the number of artificial agents, and n is the number of rounds.

The **Global Surprise (*GSu*)** at round i is calculated as

$$GSu_i = \sum_{j=1}^{100} Surprise_{ji}, \quad i = 1, 2, \dots, n \quad (5.2)$$

where $Surprise_{ji}$ is the cognitive surprise “felt” by the artificial agent j at round i , computed by the artificial surprise model proposed by Macedo and Cardoso (presented in Section 2.1.4), j is the number of artificial agents, and n is the number of rounds.

TABLE 5.2: Experiment with different strategies available (Strategies), the ability of changing current strategy or not (Fixed), and memory size (MS). For all treatments, belief threshold (BT) = 0.5.

Treatment	Strategies	Fixed	MS
1	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	Yes	60
2	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	Yes	100
3	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	Yes	500
4	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	Yes	2500
5	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	No	60
6	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	No	100
7	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	No	500
8	<i>NTS, SMAS, EMAS, OPS, SAS, LAS</i>	No	2500
9	<i>NTS</i>	Yes	60
10	<i>NTS</i>	Yes	100
11	<i>NTS</i>	Yes	500
12	<i>NTS</i>	Yes	2500

Results

For each treatment within a group we first present the results in terms of the GBS (Global Belief in Strategy), CuE (Cumulative Efficiency), CuH (Cumulative Happiness), and Cumulative GSu (Global Surprise). For each of these four variables we show the results of each of the 100 simulations (dashed gray line) and the mean (solid black line). Additionally, we compare the mean of all treatments within the same group with respect to GBS , CuE , CuH , and Cumulative GSu .

For the **first group** (G1), Treatment 1 (T1), Treatment 2 (T2), Treatment (T3), and Treatment (T4), we show in Figures 5.1, 5.2, 5.3, and 5.4, respectively, the GBS , CuE , CuH , and Cumulative GSu . Figure 5.5 shows the comparison of the mean of treatments with respect to these four previous variables.

For the **second group** (G2), Treatment 5 (T5), Treatment 6 (T6), Treatment 7 (T7), and Treatment 8 (T8), we show in Figures 5.6, 5.7, 5.8, and 5.9, respectively, the GBS , CuE , CuH , and Cumulative GSu . Figure 5.10 shows the comparison of the mean of treatments with respect to these four previous variables.




FIGURE 5.1: *GBS* (Global Belief in Strategy) for treatments of the first group.

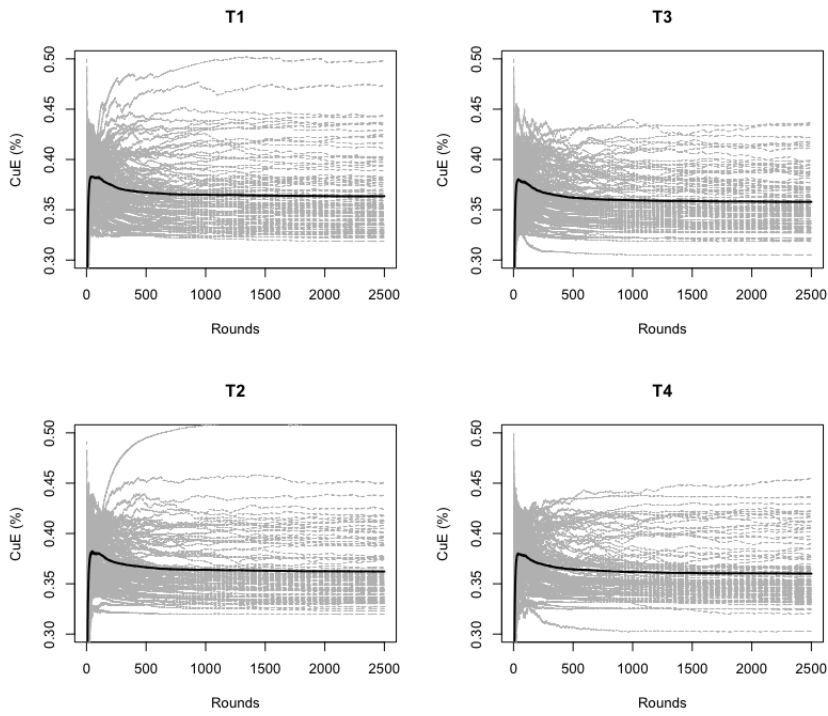


FIGURE 5.2: *CuE* (Cumulative Efficiency) for treatments of the first group.




FIGURE 5.3: CuH (Cumulative Happiness) for treatments of the first group.

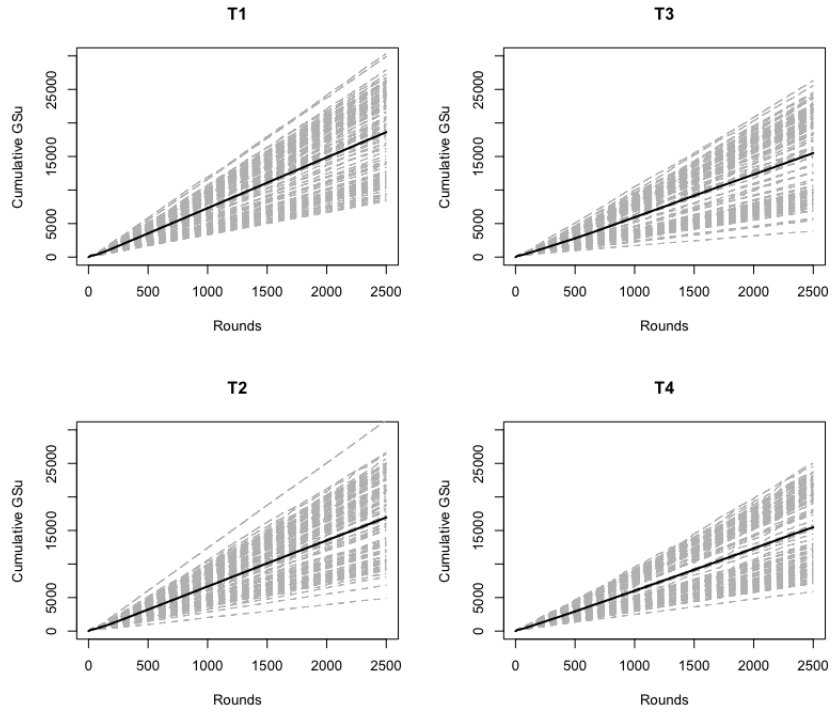


FIGURE 5.4: Cumulative GSu (Global Surprise) for treatments of the first group.

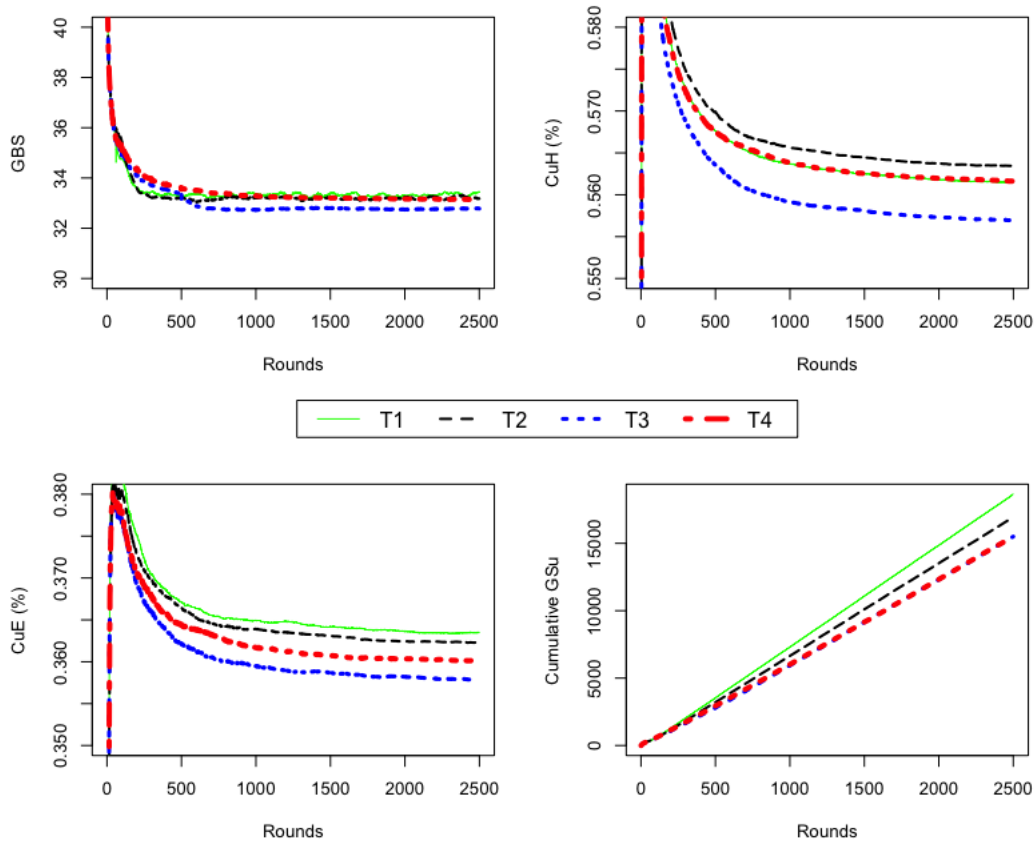


FIGURE 5.5: Comparing the mean of treatments T1($MS=60$), T2($MS=100$), T3($MS=500$), and T4($MS=2500$) of the first group, i.e., all strategies, changing of strategy not allowed, in terms of GBS , CuH , Cumulative GSu , and CuE (clockwise, from top left).

For the **third group** (G3), Treatment 9 (T9), Treatment 10 (T10), Treatment 11 (T11), and Treatment 12 (T12), we show in Figures 5.11, 5.12, 5.13, and 5.14, respectively, the GBS , CuE , CuH , and Cumulative GSu . Figure 5.15 shows the comparison of the mean of treatments with respect to these four previous variables.

We also show in Figures 5.16, 5.17, 5.18, and 5.19, the comparison of the mean of GBS , CuE , CuH , and Cumulative GSu variables, respectively, between different groups, namely G1 (T1, T2, T3, T4), G2 (T5, T6, T7, T8), and G3 (T9, T10, T11, T12).

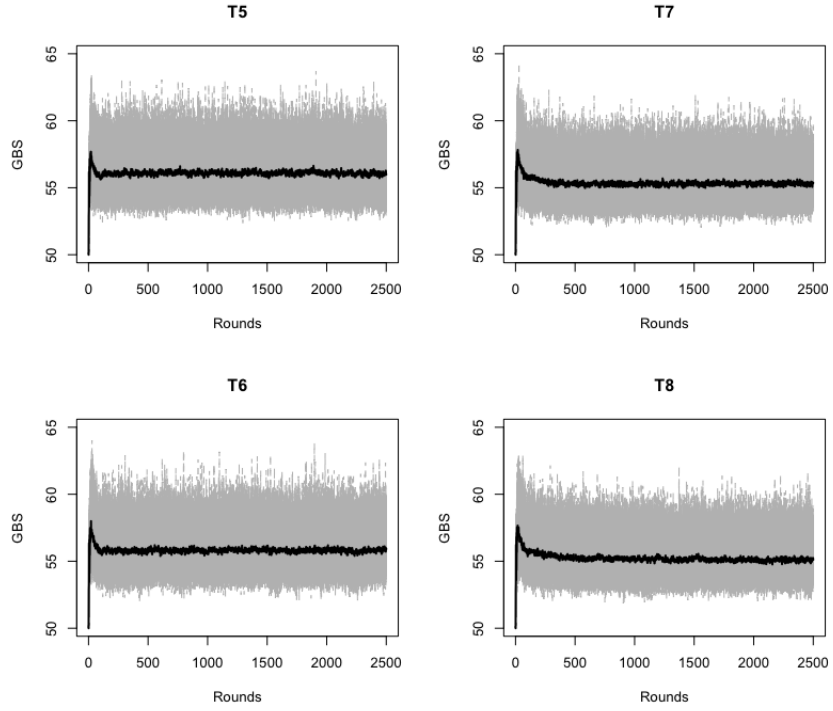


FIGURE 5.6: *GBS* (Global Belief in Strategy) for treatments of the second group.

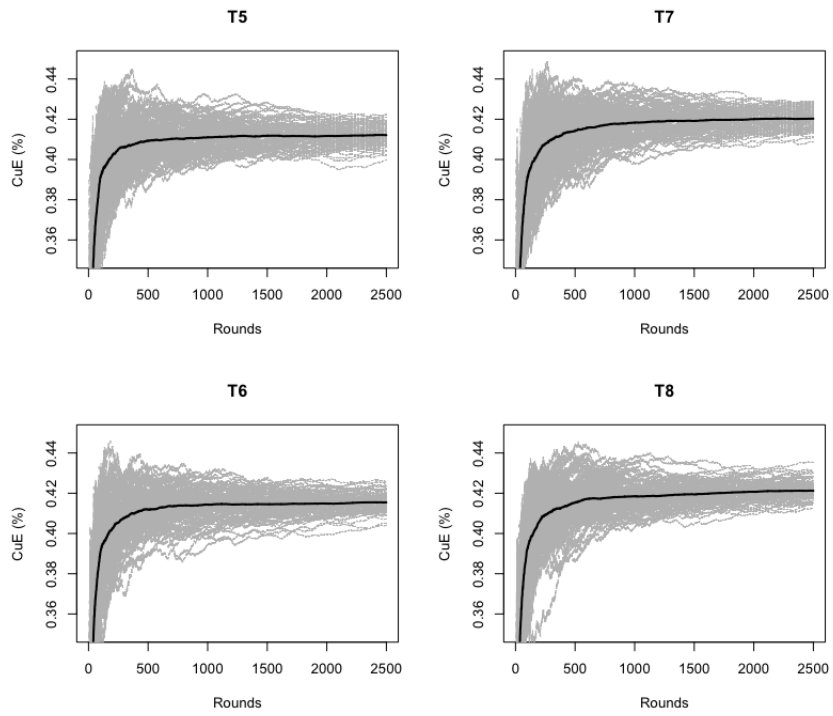


FIGURE 5.7: *CuE* (Cumulative Efficiency) for treatments of the second group.




FIGURE 5.8: CuH (Cumulative Happiness) for treatments of the second group.

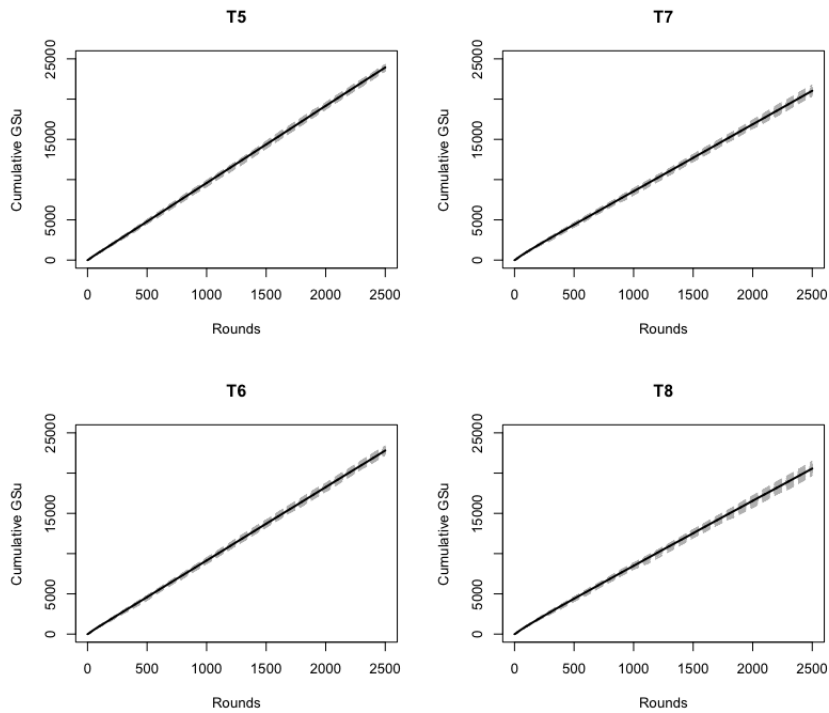


FIGURE 5.9: Cumulative GSu (Global Surprise) for treatments of the second group.

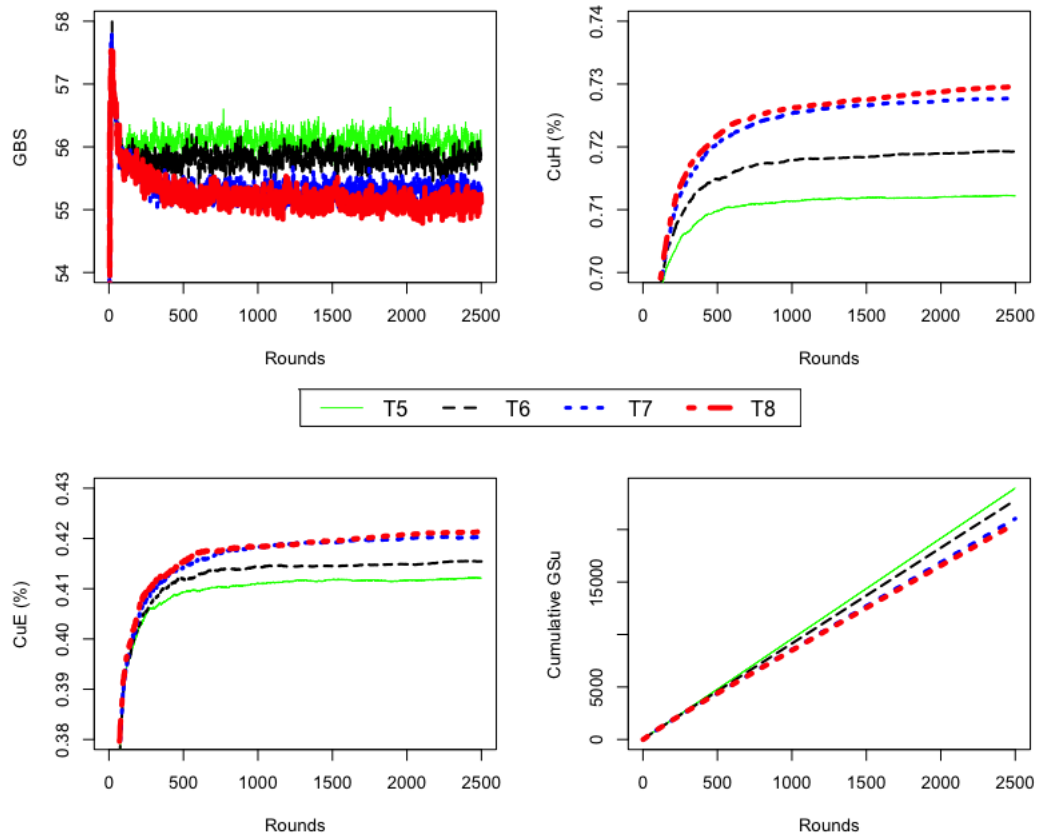


FIGURE 5.10: Comparing the mean of treatments T5($MS=60$), T6($MS=100$), T7($MS=500$), and T8($MS=2500$) of the second group, i.e., all strategies, changing of strategy allowed, in terms of GBS , CuH , Cumulative GSu , and CuE (clockwise, from top left).

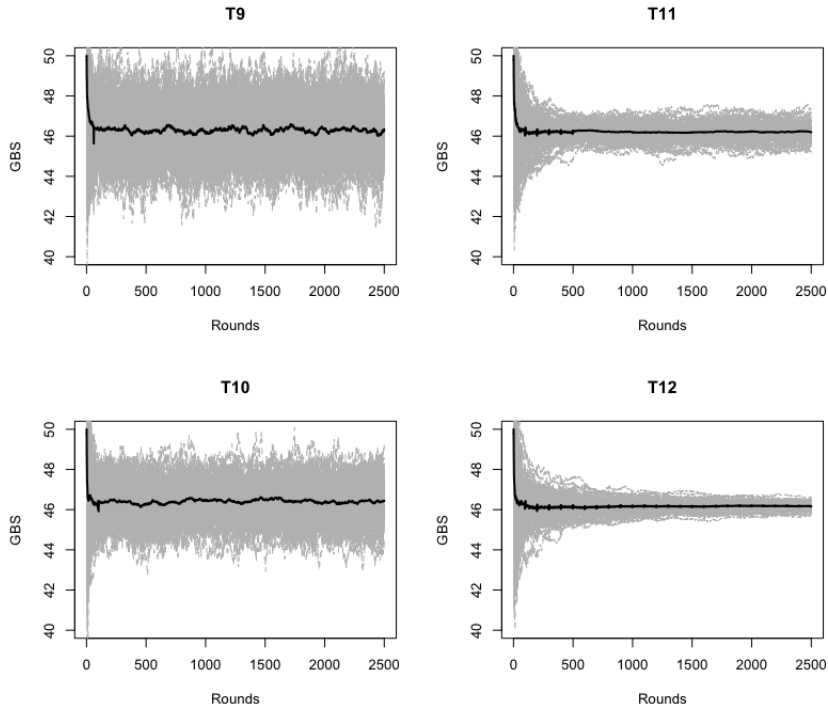


FIGURE 5.11: *GBS* (Global Belief in Strategy) for treatments of the third group.




FIGURE 5.12: *CuE* (Cumulative Efficiency) for treatments of the third group.




FIGURE 5.13: CuH (Cumulative Happiness) for treatments of the third group.

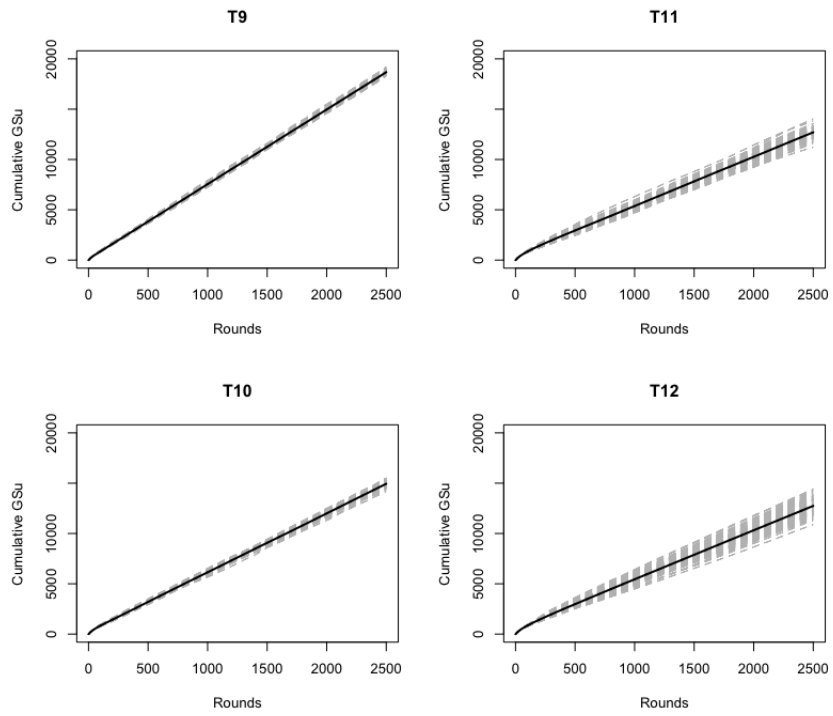


FIGURE 5.14: Cumulative GSu (Global Surprise) for treatments of the third group.

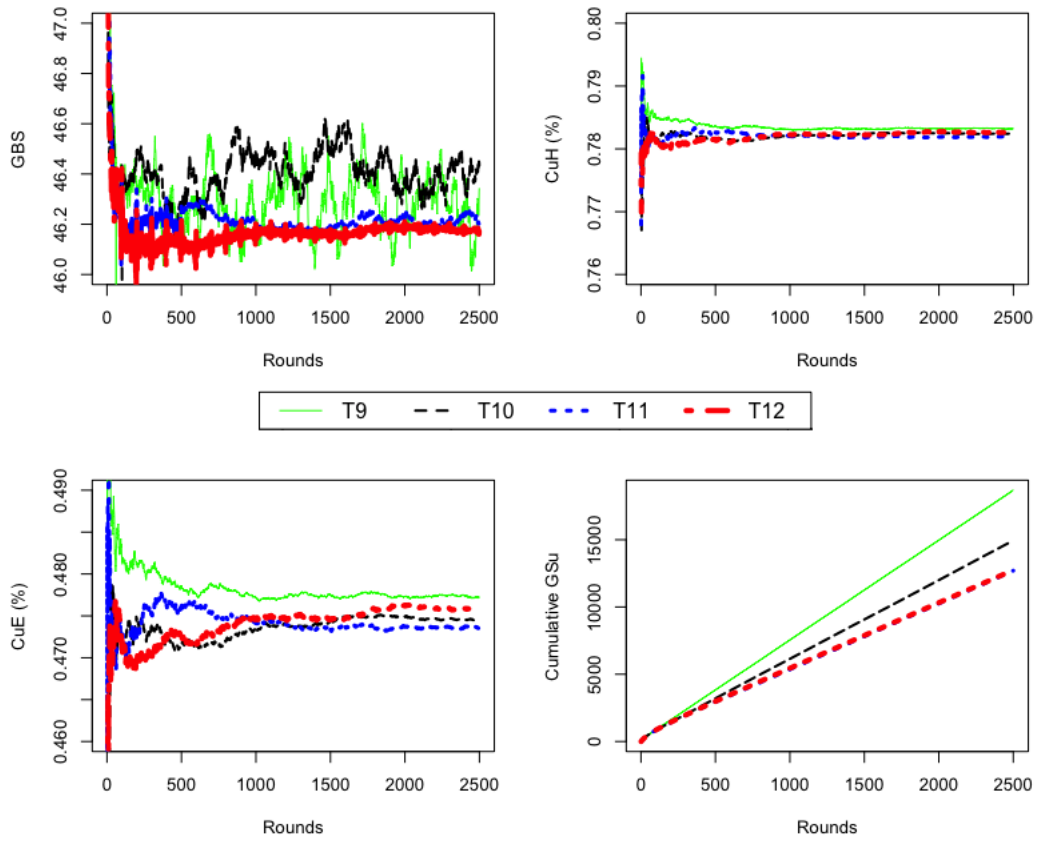


FIGURE 5.15: Comparing the mean of treatments T9($MS=60$), T10($MS=100$), T11($MS=500$), and T12($MS=2500$) of the third group, i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed, in terms of *GBS*, *CuH*, Cumulative *GSu*, and *CuE* (clockwise, from top left).




FIGURE 5.16: Comparing *GBS* (Global Belief Strategy) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed).




FIGURE 5.17: Comparing CuE (Cumulative Efficiency) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed).




FIGURE 5.18: Comparing CuH (Cumulative Happiness) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., *NTS* (Noise Trader Strategy), changing of strategy not allowed).




FIGURE 5.19: Comparing Cumulative GSu (Global Surprise) between groups: G1 (i.e., all strategies, changing of strategy not allowed), G2 (i.e., all strategies, changing of strategy allowed), and G3 (i.e., NTS (Noise Trader Strategy), changing of strategy not allowed).

Discussion

In this first exploratory experiment, we aim to observe the behavior of the cognitive artificial agents we modelled in the context of the El Farol problem. More specifically we are interested in observing how different ecologies in terms of different memory sizes, different strategies, and two different conditions on whether agents are able to change strategies might affect CuE (efficiency), CuH (happiness), and GSu (surprise) as well as the general market dynamics. Our goal is not to provide evidence with the goal of claiming a cause and effect relation between one or more variables. Instead, we simply aim to explore and therefore obtain some initial insights on the behavior of the artificial cognitive agents we modelled that may lead us to carry out other experiment(s) with other specific research goals.

The experiment is divided into three groups, with each group having specific features with respect to the set of strategies available, the ability to change the current strategy or not, and the memory size. For all treatments, the belief threshold (BT) is equals to 0.5. In the first group (G1), artificial agents select one of the six strategies available at the beginning of the simulation and stick to it for all rounds. In the second group (G2), artificial agents also select one of the six strategies but, in contrast to G1, artificial agents change the current strategy when the belief in the correctness of the strategy is lower than the belief threshold, i.e., $b(p,t) < BT$. Finally, in the third group (G3), artificial agents select the NTS (Noise Trader Strategy) and stick to it for all rounds.

Let us start by discussing the results from different treatments within a same group. For G1 treatments, we can observe in Figure 5.5 that memory size does not seem to significantly impact GBS , CuH , CuE , and GSu . Nevertheless, the results of T2($MS=100$) are better in comparison with T3($MS=500$) in terms of CuE and CuH , suggesting that the benefits of using a larger memory size might have a positive but limited impact on this ecology. For G2 treatments, like the results of G1, we can observe in Figure 5.10 that memory size does not seem to significantly impact GBS , CuH , CuE , and GSu . However, unlike the results found in G1, a larger memory size that is T7($MS=500$) seems to yield slightly better results in comparison with a smaller memory size that is T6($MS=100$) in terms of CuE and CuH . Additionally, as we increase the memory size, the results are slightly better in terms of CuE and CuH whereas GSu decreases. For G3 treatments, we can observe in Figure 5.15 that a smaller memory size that

is $T9(MS=60)$ seems to yield slightly better results in terms of CuE whereas GSu is higher.

Let us now compare treatments from different groups. Regarding GBS , we can observe in Figure 5.16 that the results of G2 treatments are higher than the G3 treatments that are, in their turn, higher than the G1 treatments. It is a result that we foresee since for G2 treatments artificial cognitive agents have the ability to change current strategy when there is a violation of the BT , meaning that they are often changing between strategies and therefore resetting their belief in the correctness of the strategy ($b(p,t)$) to 0.5. Regarding CuE and CuH , we can observe in 5.17 and in Figure 5.18, respectively, that the results of G3 treatments are higher than of G2 that are higher than G1. For all G3 treatments, artificial agents use the NTS (Noise Trader Strategy) for all rounds. Interestingly, the results of G3 are better in comparison with G2 and G1, regardless of the fact that globally, artificial agents in G3 do not believe in the correctness of the strategy they are currently using ($b(p,t)$), as stressed by the lower GBS , below the $b(p,t)=0.5$. Similarly, we can observe in Figure 5.19 that the G2 results are the highest in terms of GSu .

The results of G3 are indeed better in comparison with G1 and G2 in terms of efficiency (CuE and CuH). However, based on the comparison of the features of each group with the behavior of humans, we consider that it would be difficult for humans to use a totally random strategy in similar contexts, especially considering aspects such as the desire of predictability. As a result, we may claim that G2 treatments have more resemblances with a real world scenario and, therefore, this is the configuration that should be further investigated, as we describe in the next section. Last but not least, we consider that a drawback of this experiment is to consider rationality as a binary concept by using a fixed value for BT , in this case 0.5, and forcing artificial cognitive agents to change the current strategy if and only if $b(p,t) < BT$.

5.3 Experiment II: The impact of rationality on market efficiency

As presented throughout Section 2.2, there is a debate about the rationality of economic agents and how this rationality affects individual and global decision-making in

economic and financial contexts. Therefore, there is also a debate on whether markets, including the financial markets, are efficient or not.

We conducted an experiment with seven treatments to explore how the artificial cognitive agents we modelled behave in the context of the El Farol problem, with special focus on **how different belief thresholds (BT) impact market efficiency (CuE , CuH and HaD)**. We gradually increase the BT by a factor of 0.05 from a biased belief threshold, represented in Treatment 1 (T1) with $BT=0.20$, to a strictly “pure” rational belief threshold, represented in Treatment 7 (T7) with $BT=0.50$. The belief threshold value is the unique variable that we changed between each of the seven treatments. We summarize in Table 5.3 the features of the experiment we conducted.

TABLE 5.3: Experiment with different belief thresholds (BT). For all treatments, memory size (MS) = 100.

Treatment	Belief Threshold (BT)	Color in Figures
1	0.20	green
2	0.25	yellow
3	0.30	black
4	0.35	gray
5	0.40	blue
6	0.45	pink
7	0.50	red

Results

We start by presenting the results in terms of the attendance. Unlike the other variables, for the attendance only, we opt for presenting the results of just some runs (randomly selected), instead of normalizing the data and then presenting it. Considering the variations of the dynamics over time (which can be observed in the next figures) the normalized attendance would not preserve the interesting dynamics we can observe by looking at individual runs. Figures 5.20, 5.21, 5.22, 5.23, 5.24, 5.25, and 5.26 present the result for Treatments 1, 2, 3, 4, 5, 6, and 7, respectively.

Regarding market efficiency, we first show in Figure 5.27 and 5.28 show the results of the CuE (Cumulative Efficiency) over time and in general (box plot), respectively. Similarly, Figure 5.29 and 5.30 show the results of the CuH (Cumulative Happiness) over time and in general (box plot), respectively. Finally, Figure 5.31 and 5.32 show the results of the HaD (distribution of happiness among agents) in terms of density

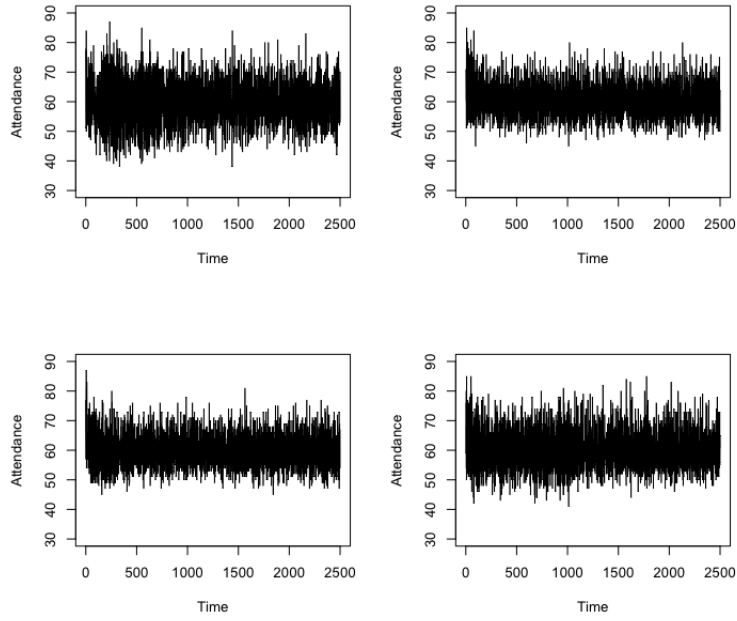


FIGURE 5.20: Attendance for four runs of Treatment 1: $BT=0.20$.

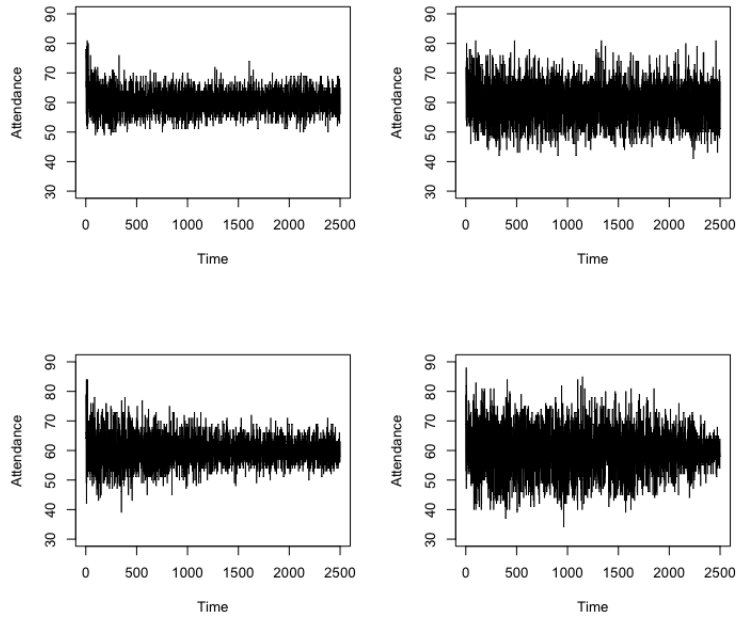


FIGURE 5.21: Attendance for four runs of Treatment 2: $BT=0.25$.

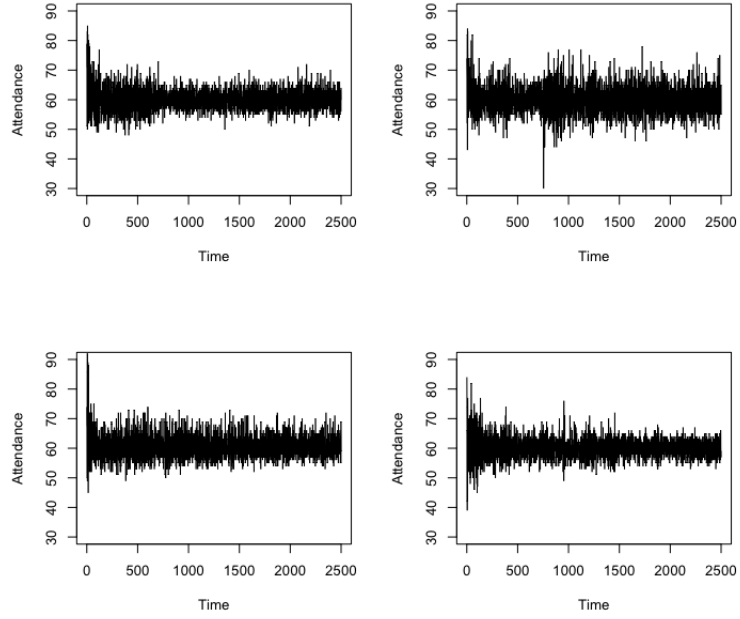


FIGURE 5.22: Attendance for four runs of Treatment 3: $BT=0.30$.

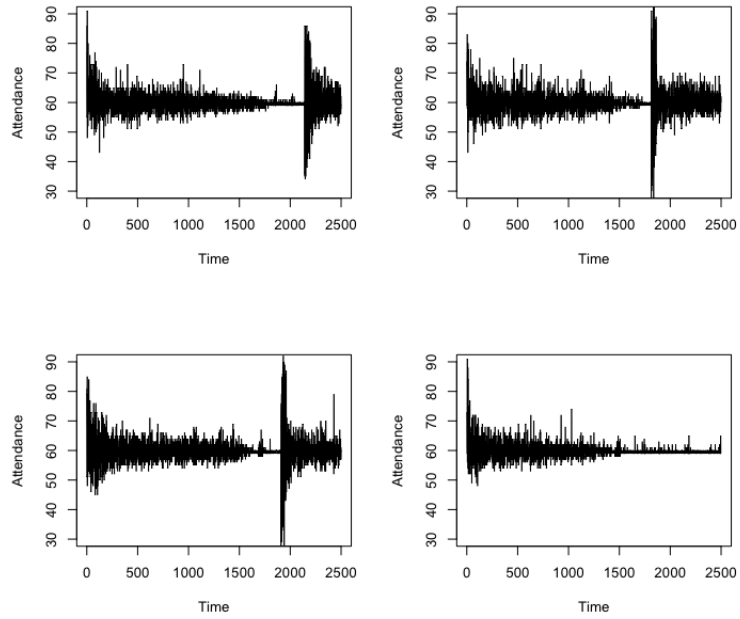


FIGURE 5.23: Attendance for four runs of Treatment 4: $BT=0.35$.




FIGURE 5.24: Attendance for four runs of Treatment 5: $BT=0.40$.




FIGURE 5.25: Attendance for four runs of Treatment 6: $BT=0.45$.

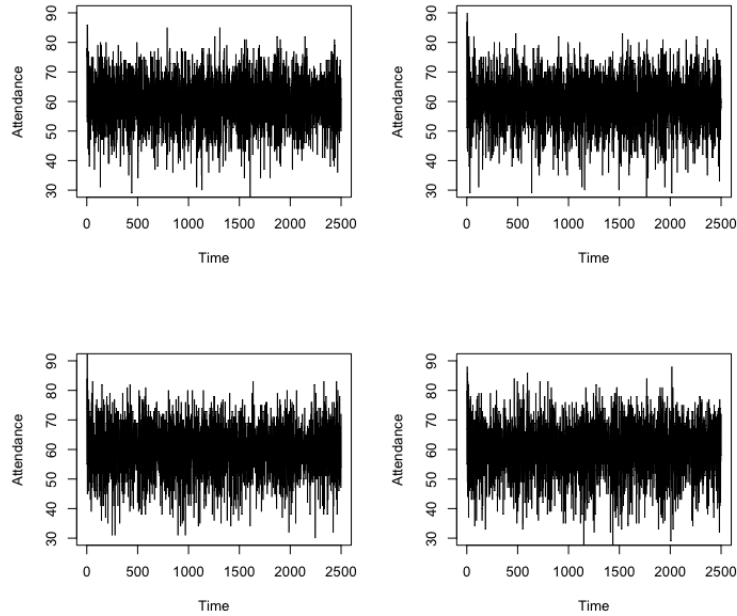


FIGURE 5.26: Attendance for four runs of Treatment 7: $BT=0.50$.

and in general (box plot), respectively. We would like to stress two opposing dynamics in comparing T1 (totally biased treatment) with T7 (pure rational treatment). We can see that T1 is the treatment that yields (A) the highest results in terms of CuE ; (B) good results in terms of CuH ; as well as (C) good results in terms of HaD . Conversely, we can see that T7 is the treatment that yields the lowest values in terms of CuE , CuH , and HaD .

Regarding the dynamics in terms of the strategies, Figure 5.33 presents the frequency of strategies for each treatment, from Treatment 1 ($BT = 0.20$) to Treatment 7 ($BT = 0.50$). Additionally, Figure 5.34 and 5.35 show the number of changes of strategies per round over time and in general (box plot), respectively. As previously presented, the artificial agent changes from current strategy when the belief in the correctness of the current strategy ($b(p, t)$) is lower than the belief threshold (BT).




FIGURE 5.27: *CuE* (Cumulative Efficiency) over time, zoomed in.




FIGURE 5.28: *CuE* (Cumulative Efficiency), box plot.




FIGURE 5.29: *CuH* (Cumulative Happiness) over time, zoomed in.




FIGURE 5.30: *CuH* (Cumulative Happiness), box plot.




FIGURE 5.31: *HaD* (distribution of happiness among agents): density.




FIGURE 5.32: *HaD* (distribution of happiness among agents), box plot.

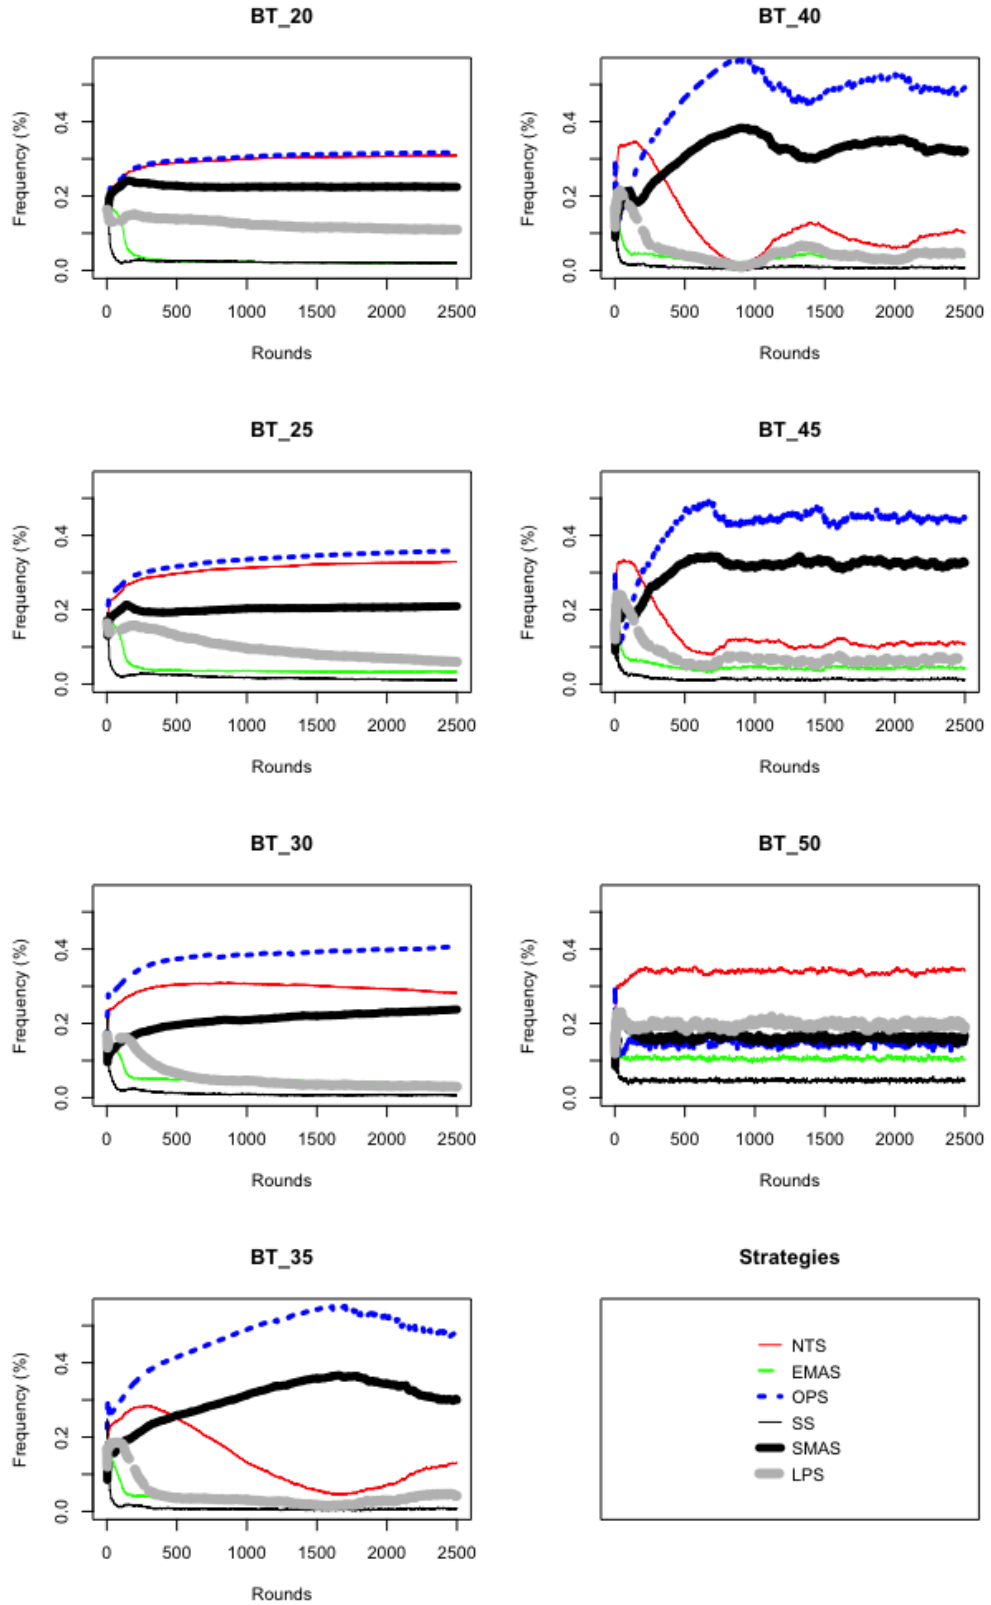


FIGURE 5.33: Frequency of strategies for each treatment, from Treatment 1 (BT_20) to Treatment 7 (BT_50), over time.




FIGURE 5.34: Number of changes of strategies per round over time.




FIGURE 5.35: Number of changes of strategies per round, box plot.

Discussion

In analyzing the results we are especially interested in observing how our concept of rationality, derived from the BDTE and carried out by means of the belief threshold (BT), as explained in Section 3.2.1, impacts market efficiency (CuE , CuH and HaD). Our final goal is to investigate whether the El Farol problem is more efficient when all cognitive agents act rationally.

As we briefly pointed out earlier in this section, interestingly, some unusual patterns emerge from the use of different belief thresholds, specifically in the attendance in T4, T5, and T6. Like Cross et al. (2005), our model, despite its apparent simplicity, was able to reproduce what seems to be a stylized fact (presented in Section 2.2.1) known as volatility clustering, that can be essentially defined as the observation that high-volatility events tend to cluster in time. Volatility clustering resembles the concept of entropy used in a variety of areas such as information and communication theory. Being able to reproduce one or more stylized facts is one of the criteria used to claim that a model is capable of producing realistic results.

Regarding the market efficiency, when we compare the cumulative efficiency (CuE), cumulative happiness (CuH), and happiness distribution among agents (HaD) of the strictly “pure” rational scenario (T7) with the other biased scenarios, namely T1, T2, T3, T4, T5, and T6, we can see that the results of the T7 are notably lower, both individually (HaD) and globally (CuE and CuH). In other words, the results suggest that the El Farol bar problem is not more efficient when all cognitive agents act rationally.

However, it is important to bear in mind that our results were obtained in a particular given setting, with specific configurations in terms of memory size, belief threshold, process for increasing and decreasing beliefs, set of strategies available. Therefore, we consider that our results should not be applied to all markets without carefully consideration. Additionally, to deal with the performance of agents or cognitive agents in similar scenarios, one needs to address the difficult problem of defining what efficiency is as well as how to measure it, since the results heavily depend on that.

Nevertheless, from a cognitive perspective our results suggest that when all cognitive agents act rationally, it does not provide better results, neither individually nor globally. We consider that our work can be used as a starting point towards truly understanding whether markets are more efficient when agents or cognitive agents act rationally. The

next natural step could be the application of the ideas and concepts discussed in this work to more sophisticated and realistic markets, such as those illustrated by other artificial financial markets platforms and models (e.g. some of them presented in Section 2.3.3) beyond the El Farol bar problem.

Despite all the efforts that could be spent, we are aware that it would be difficult (perhaps impossible) to anyone to state that markets are more efficient when all cognitive agents act rationally. However, we could be able to provide more evidence so that we could conclude that by considering that it is well known that we are not rational beings, at least not as defined by the classical economic theories, perhaps markets are as efficient as our bounded and limited rationality allow us to be and, contrary to the general idea, this is not bad.

5.4 Experiment III: The impact of memory size on market efficiency

As presented throughout Section 2.1.1, different memory systems and processes play an important role both in the human life in general and in the decision-making process in particular. The possibilities for investigations in this context are indeed huge. In this experiment we explore one of these possibilities that is the importance of the memory size.

We conducted an experiment with thirty four treatments to explore how the cognitive agents we modelled behave in the context of the El Farol problem, with special focus on **how different memory sizes (MS) impact market efficiency (CuE , CuH and HaD)**. Treatments are defined in terms of the MS , which we gradually increase from 10 to 2500. This is the unique variable we changed between the thirty four treatments. We summarize in Table 5.4 the features of the experiment we conducted.

As we explained earlier, artificial agents randomly select a strategy at the beginning of the simulation. Therefore, as a typical scenario, insofar we increase the number of simulations, we may increase the diversity of the starting conditions. However, taking into account the nature of the algorithm, specifically that all strategies are randomly selected of the set of available strategies by using a uniform distribution together with the changing dynamics provided by the belief threshold concept, we consider that the

TABLE 5.4: Experiment with different memory sizes (MS). For all treatments, belief threshold (BT) = 0.5.

Treatment	Memory Size (MS)	Color in figures
1	10	green
2	20	gray scale
3	30	gray scale
4	40	gray scale
5	50	gray scale
6	60	gray scale
7	70	gray scale
8	80	gray scale
9	90	gray scale
10	100	yellow
11	200	gray scale
12	300	gray scale
13	400	gray scale
14	500	gray scale
15	600	gray scale
16	700	gray scale
17	800	gray scale
18	900	gray scale
19	1000	black
20	1100	gray scale
21	1200	gray scale
22	1300	gray scale
23	1400	gray scale
24	1500	gray scale
25	1600	gray scale
26	1700	gray scale
27	1800	gray scale
28	1900	gray scale
29	2000	blue
30	2100	gray scale
31	2200	gray scale
32	2300	gray scale
33	2400	gray scale
34	2500	red

number of simulations may not affect the results of the experiment. As a result, like for the experiments presented earlier, we have run 100 independent simulations for this experiment.

Results

We first show the results in terms of market efficiency. Figures 5.36 and 5.37 show the CuE (Cumulative Efficiency) results over time and in general (box plot), respectively. Similarly, Figures 5.38 and 5.39 show the CuH (Cumulative Happiness) results over time and in general (box plot), respectively. Last but not least, Figures 5.40 and 5.41 show the HaD (distribution of happiness among agents) results in terms of density and in general (box plot), respectively.

Regarding the dynamics of the strategies, we show in Figure 5.42 the frequency of strategies of some key treatments, namely Treatment 1 ($MS=10$), Treatment 10 ($MS=100$), Treatment 19 ($MS=1000$), Treatment 29 ($MS=2000$), and Treatment 34 ($MS=2500$). Additionally, we stress the dynamics over time of the Noise Trader Strategy (NTS), by compiling its frequency in Figure 5.43. We conclude the presentation of the results by showing in Figures 5.44 and 5.45 the number of changes of strategies per round over time and in general (box plot), respectively.




FIGURE 5.36: CuE (Cumulative Efficiency) over time, zoomed in.

We did not show the results concerning the attendances because we observed that attendance gravitates towards the bar capacity of 60, an expected result that is in line with what was found in previous works (e.g. (Arthur, 1994)). As we increase




FIGURE 5.37: CuE (Cumulative Efficiency), box plot.




FIGURE 5.38: CuH (Cumulative Happiness) over time, zoomed in.




FIGURE 5.39: *CuH* (Cumulative Happiness), box plot.




FIGURE 5.40: *HaD* (distribution of happiness among agents), density.




FIGURE 5.41: *HaD* (distribution of happiness among agents), box plot.

the memory size, *CuH* and *CuE* seem to minimally increase from T1 to T15, with differences in the last round no greater than roughly 0.06, as we can see in Figures 5.36 and 5.38, respectively. Similarly, memory size seem to has a positive impact on *HaD*, increasing it from approximately 980 in T1 to 1080 in T34, as we can see in Figure 5.41. Similarly, regarding the dynamics of the strategies, we can observe that as we increase the memory size, on the one hand, the frequency of the *NTS* increases, as we can see in Figure 5.43, while, on the other hand, the number of changes of strategy per round decreases, as we can see in Figures 5.44 and 5.45.

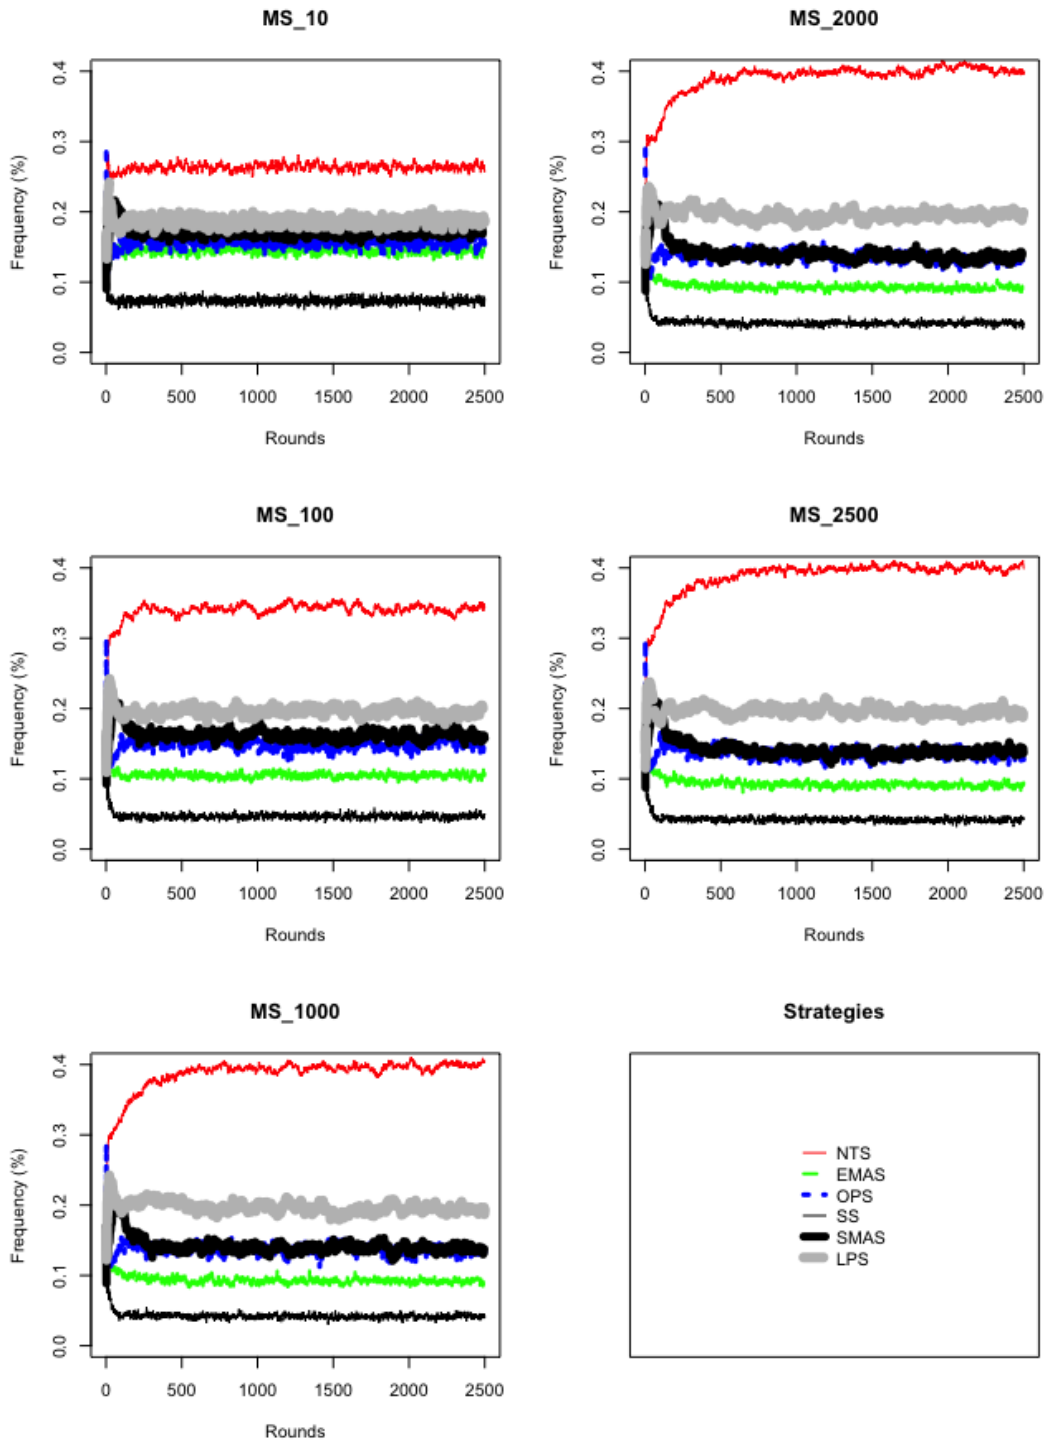


FIGURE 5.42: Frequency of strategies of some treatments, namely Treatment 1 (MS_10), Treatment 10 (MS_100), Treatment 19 (MS_1000), Treatment 29 (MS_2000), and Treatment 34 (MS_2500).




FIGURE 5.43: Frequency of the *NTS* (Noise Trader Strategy) compiled from all treatments.




FIGURE 5.44: Number of changes of strategy per round over time.




FIGURE 5.45: Number of changes of strategy per round, box plot.

Discussion

First of all, it is important to take into account that there is a debate on whether the memory size is relevant or not. For example, in the context of Minority Games (MGs) (presented in Section 2.3.3), Cavagna (1999) applied extensive numerical simulations to show that all features of the MG are completely independent of the memory of the agents. Cavagna states that what really matters is not the agent memory size, but rather the fact that all agents share the same information, regardless of this information being true or false. Challet and Marsili (2000) contradicts the work of Cavagna by quantitatively studying the dynamics of the histories. They state that such dynamics cannot be considered irrelevant. Also, according to Challet and Marsili, memory is not only relevant but one would also be able to quantify the difference between using random and original history.

We would like to stress that in this work the memory size (MS) of the agent consists of both positive and negative instances, resulting from the use of the current strategy in a given number of last rounds. The experience stored in the MS is later used to calculate, via a Bayesian process (i.e., BayesNet algorithm implementation available in the data mining software Weka), the degree of belief in the correctness of the strategy

that is the $b(p,t)$. Similarly, $b(p,t)$ relates to the concept of belief threshold (BT) in the sense that when a $b(p,t) < BT$, in this work $BT=0.5$, the agent must change its current strategy. In this context, it is important not to confuse the concept of memory size (MS) with the concept of window size used by some strategies to generate a forecast for the next attendance. The fact that some strategies may have a window size higher than the MS does not have any effect on the results at all.

From treatment T1 to T34, presented in Table 5.4, we gradually increased the memory size of all cognitive agents. We noticed two different dynamics when analyzing the results, the first concerning treatments from T1 ($MS=10$) to T15 ($MS=600$), and the second concerning treatments from T16 ($MS=700$) to T34 ($MS=2500$). The effects of increasing the memory size on market efficiency, especially in terms of CuH (Cumulative Happiness) and HaD (distribution of happiness among agents), are positive. However, while the benefits are relatively significant from T1 to T15, the results from T16 do not significantly differ from the results of T17. Additionally, from T16 to T34, the benefits of having a larger memory size vanish and somehow stabilize in the sense that as we increase the memory size, market efficiency does not increase.

Surprisingly, when we turned our attention to the aspects concerning the strategies, we found out a quite similar pattern. The frequency of the NTS increases as we increase the memory size, whereas the number of changes of strategy per round decreases. We are not able to claim that the results we found are due to the larger memory size and/or if they are due to one or more secondary effects of having a larger memory size such as the higher frequency of the NTS and/or the lower number of changes of strategy per round. Further experiments with an additional set of treatments (e.g., a treatment comparing agents with and without the NTS) would be necessary to provide more evidence so that one may be able to answer that question properly.

Nevertheless, our results suggest that a larger memory size have a somehow limited effect, increasing market efficiency but only to some extent. The effects of increasing the memory size seem to vanish and stabilize from T15 ($MS=600$) to T16 ($MS=700$). In other words, it means that from T15 to T16, and from T16 to T34, artificial agents would “feel” equally happier both globally (CuH) and individually (HaD) if all cognitive rational agents have a larger memory size. Similarly, the use of the resource, represented in this experiment by the CuE , would be more efficient. Practically speaking, in the context of the El Farol bar, the scenario that seems to yield the better market efficiency

(*CuH*, *HaD*, *CuE*) for the cognitive rational agents ($BT=0.5$) is the one in which all agents have the same memory size, which size is roughly 600.

One of the possible reasons is that our results were obtained with a particular configuration. For example, we assume that all agents use the same Bayesian process regarding the $b(p,t)$. This means that, although we do not know which are the preferences of the agents and, as expected in this sort of experiment, we need to make assumptions, relying on the same process for all artificial agents might be a drawback of our approach, especially with respect to yielding results as realistic as possible (not necessarily our goal in this work). Additionally, in similar scenarios, one needs to address the problem of defining and measuring efficiency, since the results, as in our work, heavily depend on that. As a result of this, it would be difficult, and probably in vain, to compare, either qualitatively or quantitatively, our results with the results of others.

Despite the fact that although our approaches are different and were applied to distinct contexts with different purposes, our results seem to support some previous works. The fact that “remembering” all the experience in using a strategy or, in our case, having a larger memory size, is not necessarily better to make the best possible prediction, is in accordance with the work of Mitra (2005). Similarly, in our work the use of different memory sizes generates different results in terms of efficiency, providing evidence to support the argument of Challet and Marsili (2000) on the relevance of memory size, as opposed to the work of Cavagna (1999). In addition, the fact that the *NTS* is the strategy that yields the best performance is consistent with previous works that show that strategies based on the past (e.g., *SMAS*, *EMAS*) do not perform better than the purely random strategy (e.g., *NTS*) (Biondo et al., 2013).

Even though our results cannot be generalized to other domains, we consider that our work may be used as a starting point towards the understanding of what would be an “optimal” memory size for human agents regarding the performance of their strategy in other economic and financial domains (e.g., stock markets). The next step may be the application of the ideas and concepts discussed in this work to more sophisticated and realistic markets or the incorporation of some findings about how humans use past experience in decision-making (e.g., (Griffin and Tversky, 1992)).

Chapter 6

Conclusion

This thesis presented a novel and interdisciplinary approach to advance the understanding of the financial markets. The main goal of this research work was to explore which contributions the application of a cognitive modeling approach could bring to the understanding of the financial markets, specifically to better explain and understand the behavior of market participants (human agents) in a micro (individual) perspective as well as the behavior of the financial markets in a macro (global) perspective. The contributions of our work are organized in this thesis as follows.

6.1 Contributions

Contributions to the advance of interdisciplinary approaches for the study of economic and financial systems and problems

In Chapter 2 we described the ideas, concepts, definitions, and assumptions of the following three major areas in which this work relies upon: Computational Cognitive Modeling, Financial Markets, and Agent-based Financial Markets.

In the Computational Cognitive Modeling section, we described how some emotions (like happiness, unhappiness, and fear) can be defined and represented as well as described some functions associated with emotions, by giving special attention to the importance of emotions for both the decision-making process and how emotions influence the memory processes of encoding, storing, and retrieving. Then, we presented

the Belief-Desire Theory of Emotions (BDTE), a cognitive emotion theory, and detailed how emotions can be quantitatively defined and qualitatively computed in this context. Subsequently, we focused on the surprise emotion by describing how surprise can be defined, presented some of its functions, and characterized the surprise process proposed by Meyer and colleagues. Furthermore, we described the research on artificial surprise, specifically how surprise can be computed from the cognitive science perspective by the use of the artificial surprise model proposed by Macedo and Cardoso.

The Computational Cognitive Modeling section is a meaningful contribution to the advance in the design and realization of interdisciplinary approaches, particularly to the study of any economic and financial system or problem, by clearly presenting to researchers from other areas such as Economics some important concepts of the cognitive science perspective, informing them about the benefits, challenges and opportunities. For instance, by detailing the beneficial and essential role of emotions for the decision-making process and judgment, this section contributes to the deconstruction of the myth that emotions are in conflict with the idea of rationality.

In the Financial Markets section, we compared the assumptions of the traditional economics theories about market participants and the financial markets, exemplified by the Efficient Market Hypothesis (EMH), such as the claim of rationality and market efficiency, with relevant findings from other research areas such as Behavioral Economics, notably the existence of behavioral biases, i.e., deviations from the so-called rational behavior. Then, we presented the Adaptive Market Hypothesis (AMH), a recent and still under development novel hypothesis that brings together the EMH with behavioral economics by applying the principles of evolution, i.e., innovation, adaptation, competition, and natural selection, to economic and financial interactions. We described the AMH assumptions as well as pointed out some concrete implications of the AMH.

To the best of our knowledge, the AMH is currently the most appropriate hypothesis to understand the behavior of market participants in the financial markets. Like the Computational Cognitive Modeling section, the clear presentation and comparison of different hypothesis that try to explain the behavior of market participants in the financial markets together with a summarization of the recent findings of the AMH constitute meaningful contributions to the advance of interdisciplinary approaches.

We concluded the Financial Markets section by presenting how surprise and related terms such as “less (more) than expected” have been frequently used in the financial

markets. We presented the market consensus idea as well as the importance of the earnings surprise for market prices and behavior. We finish this section by detailing several equations to mathematically compute earnings surprise.

TribeCA: a Generic Conceptual Cognitive Model and implementation

In Chapter 3 we described the proposed conceptual cognitive model named TribeCA (Trading and investing with behavioral-economical Cognitive Agents) and its implementation. The TribeCA has roots in the Belief-Desire Theory of Emotions, the surprise process proposed by Meyer and colleagues, and the artificial surprise model proposed by Macedo and Cardoso.

We then describe four concepts we devise, one to implement rationality in agents and three metrics to measure efficiency. Rationality is implemented through the belief threshold (BT) concept. In this work, artificial cognitive agents have a belief in the correctness of the current strategy, i.e., $b(p,t)$, and when the $b(p,t)$ is lower than BT the agent needs to select another strategy since he/she does not believe anymore in the correctness of the current strategy. Efficiency is measured by the following three metrics: Cumulative Efficiency (CuE), Cumulative Happiness (CuH), and how happiness is distributed among agents (HaD). CuE measures how efficient the resource has been used. CuH measures happiness from the global perspective (like the global wealth concept), whereas HaD measures happiness from the individual perspective (the fairness of the global wealth distribution).

The underlying theory of these four concepts are generic in the sense that they can be applied to other problems. However, to this end, it is necessary to conceive not only a clear definition of what efficiency means but also a way of computing it by taking into account the particularities of the problem or system to be addressed. Therefore, it is important to note that the mathematical equations to measure CuE , CuH , and HaD might be used by other researchers to measure efficiency only in the context of the El Farol bar problem. Even in this situation, some minor adjustments might be needed.

The implementation of the proposed model was integrated into the JABM (Java Agent-Based Modelling toolkit) and JASA (Java Auction Simulator API), two relatively popular tools used in the context of Agent-based Financial Markets, a subfield of Agent-based Computational Economics (ACE). The resulting platform allows the design and realization of a variety of economic and financial experiments with artificial cognitive agents with different agent-based models.

Two case studies presenting the cognitive science perspective of surprise

In Chapter 4, we presented two case studies we carried out in order to explore the application of the cognitive science perspective on surprise, specifically the artificial surprise model proposed by Macedo and Cardoso, to economics and finance. In the first case study, we compared the computation of surprise from the economic and financial perspective to the cognitive science perspective with respect to the idea of market consensus regarding forecasts. In the second case study, we computed the cognitive surprise “felt” by an artificial agent relying on a popular risk management tool known as Value-at-Risk (VaR) historical under two different scenarios, namely a Calm period and a Crash period.

The main contribution of the realization of these two case studies resides in the fact that we have applied in a systematic, clear and easy to reproduce way the ideas, concepts and method of computing artificial surprise in two different economic and financial scenarios. To the best of our knowledge, it is one of the first attempts (possibly the first) to apply a cognitive science perspective of surprise to market consensus, risk, uncertainty, and risk management scenarios. We consider these two case studies as a modest contribution, though a good starting point, towards truly understanding the surprise “felt” by market participants in economics and finance.

Three agent-based experiments in the context of the El Farol bar problem

In Chapter 5, we presented the three agent-based experiments we carried out with the TribeCA in the El Farol bar problem. In the first experiment, we explored the behavior of our artificial cognitive agents in twelve different treatments characterized in terms of set of strategies available, the ability of changing the current strategy, and different memory sizes. In the second experiment, we focused on observing the impact of rationality (BT) on market efficiency during seven treatments in which we gradually increased the BT by a factor of 0.05 from 0.20 (totally biased scenario) to 0.50 (rational scenario). In the third experiment, we investigated the impact of memory size (MS) on market efficiency in thirty-four treatments in which we gradually increased the MS from 10 to 2500. Market efficiency is measured by CuE , CuH , and HaD .

In these three experiments we explore some of the fundamental questions regarding the understanding of the financial markets such as rationality and efficiency. Nevertheless, the use of cognitive modeling approaches in the context of complex systems, especially in

economics and finance, is in its early stages. We consider that the use of relatively simple but powerful tools (e.g., BDTE) together with agent-based modeling and simulation, offers a rich and novel set of possibilities for investigation and to shed light on the behavior of cognitive agents in those contexts.

6.2 Future work

The interdisciplinary exploratory work presented in this thesis has contributed to gain a broad experience on the benefits and challenges of addressing the problem of understanding the behavior of both market participants individually and the financial markets globally from the perspective of cognitive science. In addition to the contributions presented earlier on this chapter, we can foresee several research directions and topics as a continuation of the present work.

Experiments with other agent-based models for studying the financial markets

In this thesis we carried out experiments with the El Farol bar problem. However, the resulting platform we have built can be also used to realize experiments with artificial cognitive agents using other agent-based models. We consider that the order-driven market agent-based model with heterogeneous agents proposed by Chiarella and Iori (presented in Section 2.3.3) is a good starting point. The model of Chiarella and Iori has more resemblances with the real financial markets than the El Farol bar problem. Besides, their model is able to reproduce many of the complex phenomena observed in real stock markets (presented in Section 2.2.1) as well as to generate realistic price dynamics.

Cognitive surprise versus economic and financial indicators: investigating how deviations from consensus generate cognitive surprise

We consider that there are several research questions that may be addressed by the application of the artificial surprise model proposed by Macedo and Cardoso. For example, suppose that the market consensus for the fictional economic indicator α is 100.00. Then, assume the following two scenarios with respect to the release of α : 1) $\alpha = 110$; 2) $\alpha = 90$. In both scenarios, the deviation from the market consensus is

10%. But, does scenario 1 elicit higher surprise than scenario 2 and vice versa? This observation leads to the following research question:

- Do symmetrical or asymmetrical deviations from market consensus generate symmetrical or asymmetrical cognitive surprise?

The application of the previous question to particular scenarios and financial products may also lead to other questions which have the potential to generate novel insights in different areas. Here we briefly present a non exhaustive list of these research questions (presented in a random order), as follows:

- Do symmetrical or asymmetrical deviations from market consensus regarding earnings surprise generate symmetrical or asymmetrical cognitive surprise?
- Do symmetrical or asymmetrical deviations from market consensus regarding different kinds of risk (e.g., credit risk) generate symmetrical or asymmetrical cognitive surprise?

Investigating the relation between leverage, cognitive surprise and fear

Leverage is a tool widely used by market participants in the financial markets that can be created through different financial instruments such as options and futures. The concept of leverage can be thought of as elasticity, the higher the leverage the higher is the potential return (negative or positive) of an investment. For example, considering that a market participant with an initial equity of \$1000 buys 100 shares of company A for a price of \$100 each, the level of leverage is 10, i.e., $(100 \times \$100)/\1000 . Then, assume that in the next trading day shares of company A depreciate (appreciate) 3% in its underlying market. Whilst a market participant with no leverage (level of leverage equals 1) would see its portfolio depreciate (appreciate) 3%, the market participant of our example would see its portfolio depreciate (appreciate) 30% due to its level of leverage ($3\% \times 10$). We can foresee some interesting research questions such as:

- What is the relation between level of leverage and cognitive surprise in the presence of appreciations and depreciations of the underlying asset?
- What is the relation between level of leverage and fear in the presence of appreciations and depreciations of the underlying asset?

For those interested in investigating one or more of these questions, we consider that a good starting point would be to carefully design and apply simple experiments with market participants (human agents), both professionals in the financial markets (e.g., stockbroker) and amateurs (e.g., ordinary market participants, small investors).

Implementing a Cognitive Surprise Indicator

We consider it would be interesting to compare the dynamics of uncertainty and fear indicators with the cognitive surprise. As presented in this work, the computation of cognitive surprise requires the creation of the subjective quantitative beliefs regarding the possible outcomes of a given event. To this end, the subjective quantitative beliefs of market participants about relevant economic and financial events can be explicitly (e.g., self-report) or implicitly (e.g., text-mining, sentiment analysis) extracted. Different Big data tools and technologies together with machine learning techniques offer a very interesting and promising opportunities for investigation (e.g., (O’Leary, 2013; Varian, 2014)). For example, a sentiment analysis technique can be used to extract the forecasts, opinions, mood, etc of market participants with respect to a given indicator or event (e.g., earnings announcement) from different social networks (e.g., (Azar and Lo, 2016; Bollen et al., 2010; Ranco et al., 2015)). As a result, we believe that it is possible to successfully implement a Cognitive Surprise Indicator (CSI). Additionally, we have an intuition that the comparison of the results of different uncertainty (e.g., economic policy uncertainty (Baker et al., 2015)) and fear indicators (e.g., VIX (also known as “the fear index”) (Fleming et al., 1995; Whaley, 1993, 2000, 2009)) with the results of the CSI may provide evidence in favor of the CSI. Perhaps a cognitive surprise indicator might better explain what happens in the financial markets when something truly surprises market participants.

From robo-advisors to cognitive robo-advisors: incorporation of market participants’ behavior and emotions

In the last years there has been a significant increase in the investments in fintech (financial technology) startups and companies. A new class of financial service known as robo-advising is one of the areas with rapid growth. The word robo is used in this context as an umbrella term to refer to any computational system, not necessarily an intelligent system in the sense that it incorporates an artificial intelligence technique or tool, for instance. Robo-advisors (or robo-advisers) are online services that provide automated, algorithmic and computational-based portfolio building and management

with minimal (or zero) human intervention. The idea is to build and manage the investor's portfolio based on his/her goals, time-horizon, budget, preferences, risk profile, etc.

The use of robo-advisors is in its early stages. Some of the pioneer companies like Betterment and Wealthfront currently manage roughly \$3.5 and \$5 billion of assets, respectively, just a tiny fraction of the money being managed by the gigantic financial services industry. The potentialities and benefits of robo-advisors are huge. When compared to traditional brokerages and services, a robo-advisor is a low-cost and transparent (with no hidden fees) alternative solution. We believe that this new class of service will be increasingly used by market participants, especially small investors who would probably not have access to similar services in the traditional financial industry.

As we have stressed in this work, the human decision-making and, therefore, investing is an emotional and extremely complex process. Like others (e.g., (Cabrera-Paniagua et al., 2015; Lo, 2016a)), we consider that one of the problems to be addressed by the robo-advisors is how to incorporate behavioral and emotional aspects into robos. If successful, it would result in, let us say, cognitive robo-advisors. We believe a cognitive robo-advisor could act essentially in two scenarios. First, by being the artificial counterpart of its human agent, acting in the financial markets on behalf of its human agent with similar (ideally improved) behavior and emotions. Second, by acting as an advisor whenever requested or offering advices in moments of turbulences (e.g., high volatility) in the financial markets.

For example, imagine that a human agent is investing in the S&P500 index through its cognitive robo-advisor. Then, in a given day, the S&P500 index falls 3% and the human agent fears losing money and freaks out. In the first scenario, the cognitive robo-advisor could automatically sell the position given a certain level of fear, surprise, etc, calibrated with the fear, surprise, etc of its human agent. In the second scenario, the cognitive robo-advisor could provide advices in order to try to calm the human agent maybe remembering his/her past emotions and the respective results in similar situations. Helping individual market participants to make better economic and financial decisions, particularly in the financial markets, improves not only how human agents allocate their savings in order to achieve their goals but also contributes to have more stable and health financial markets in the end.

During this thesis the author had developed the TribeCA (Trading and investing with behavioral-emotional Cognitive Agents) entrepreneurship project. In addition to this project, this work contributes to the advance of the interdisciplinary knowledge that would be certainly needed to those interested in solving the difficult but fascinating (and probably very lucrative) problem of constructing the so-called cognitive robo-advisors.

Appendix A

TribeCA: design and implementation

In this appendix we briefly introduce the main classes of the JABM as well as of our implementation of the proposed cognitive model (TribeCA) presented in Chapter 3.

A.1 JABM classes

The main JABM classes relevant to our work are shown below. Figure A.1 shows some of the main classes. The Simulation interface (Figure A.1) is used to define the agent-based simulation models. A simulation contains two other important classes, namely Population and SimulationController. The Population class (Figure A.2) represents the population of agents in a simulation, which can be resized dynamically, while class SimulationController is responsible for running a batch of one or more independent runs. Agents are represented by the Agent Interface (Figure A.5). The behavior of each agent (including its possible actions) is encapsulated by its strategy, represented by the Strategy interface. The Agent interface is implemented by the AbstractAgent class (Figure A.6). The list of agents in a simulation represented by the AgentList class (Figure A.7).

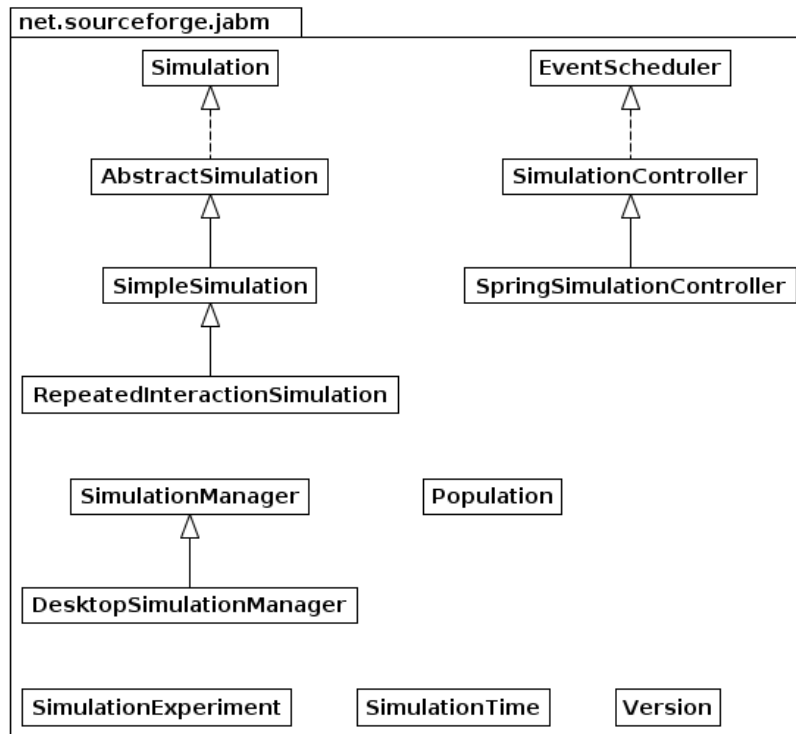


FIGURE A.1: JABM main classes.

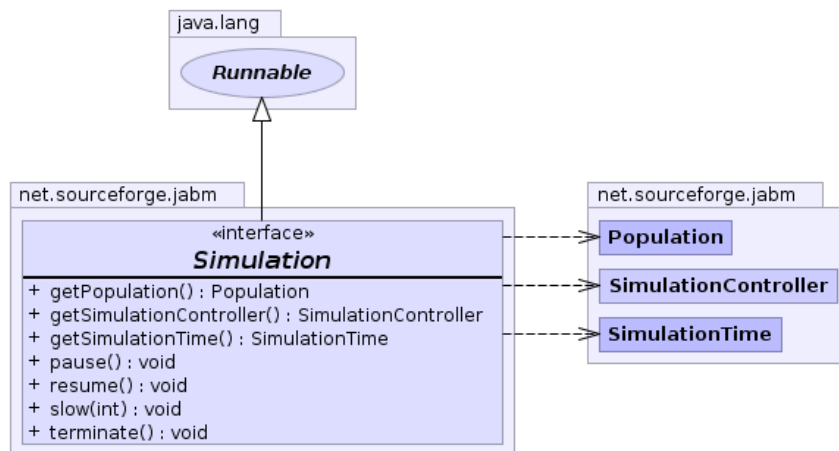


FIGURE A.2: JABM: Simulation interface.

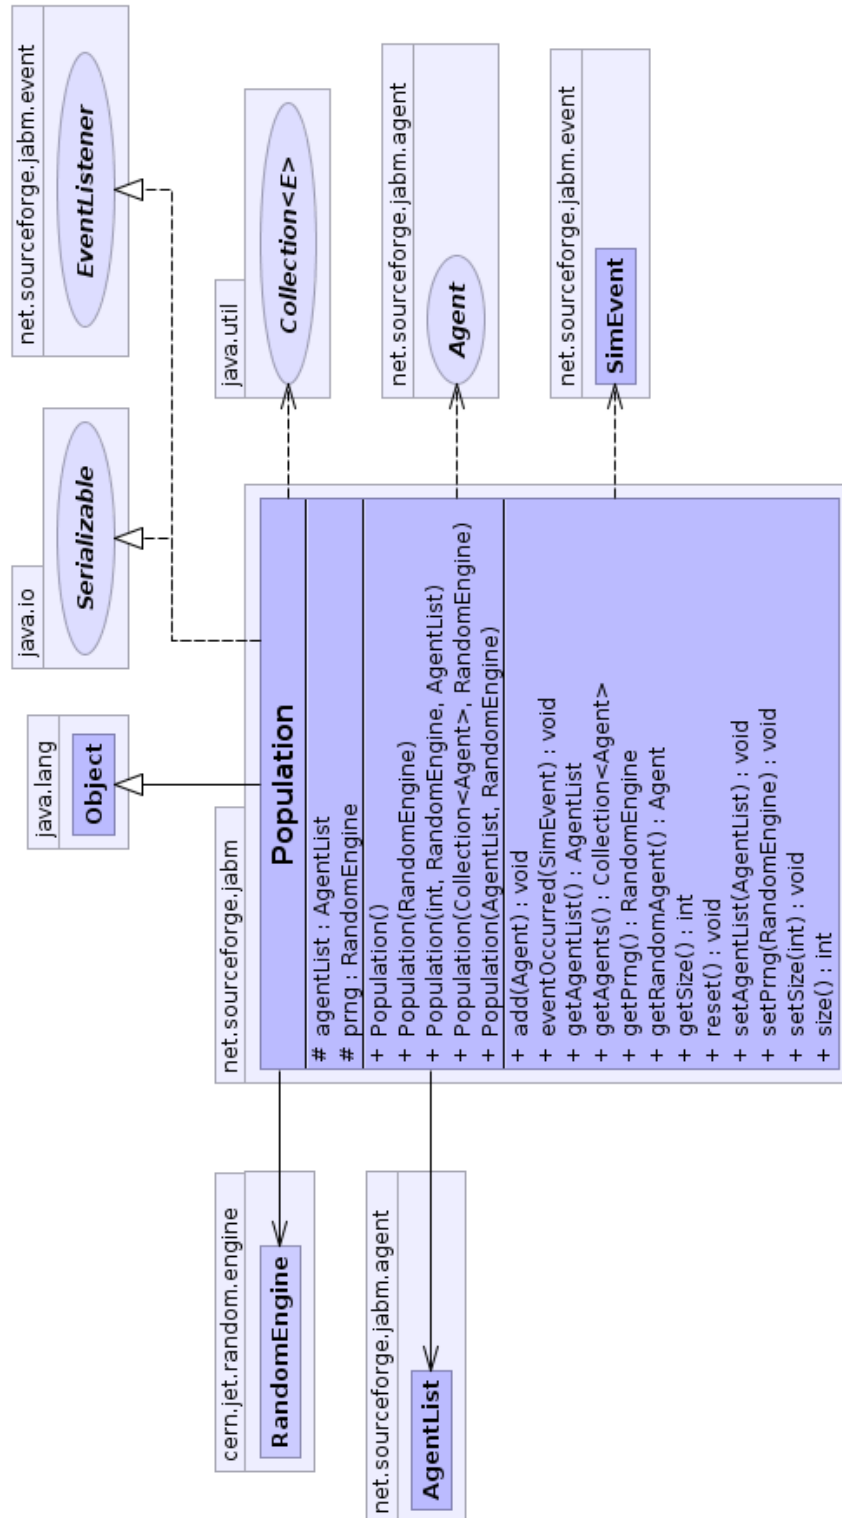


FIGURE A.3: JABM: Population class.

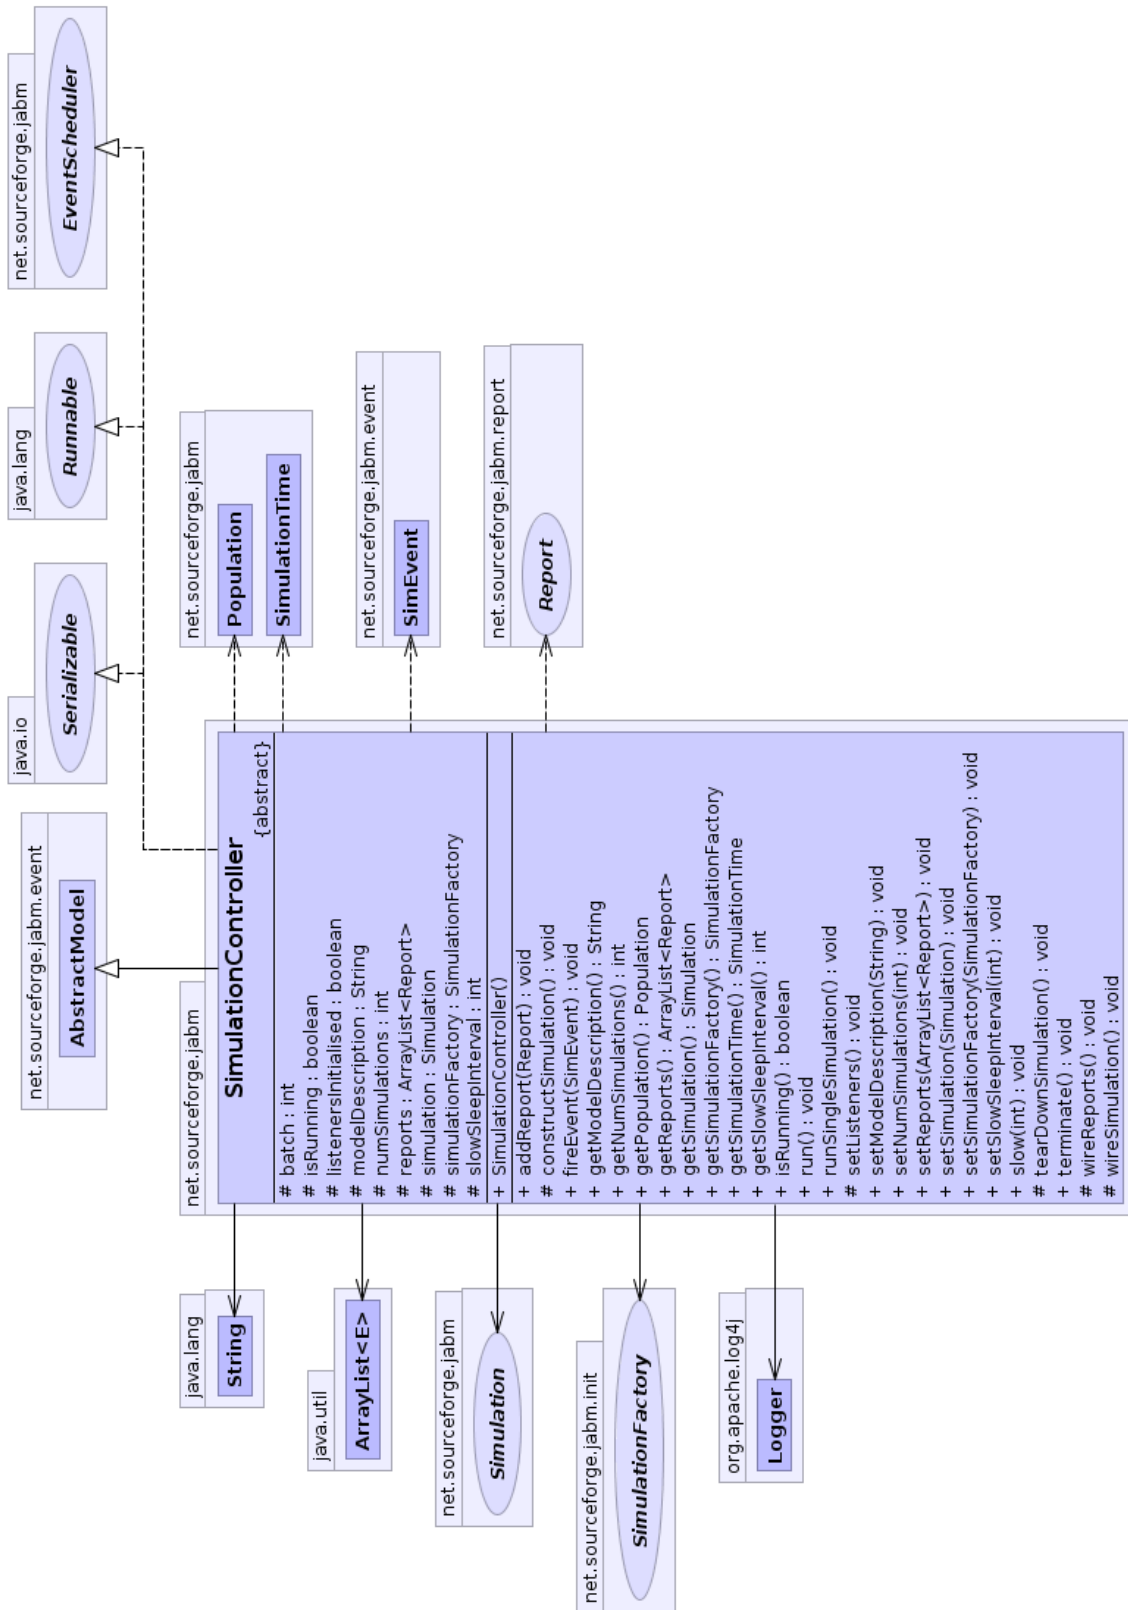


FIGURE A.4: JABM: SimulationController class.

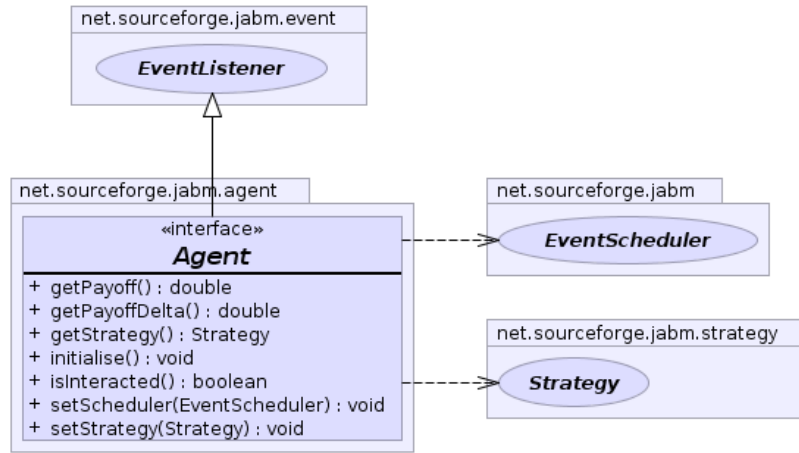


FIGURE A.5: JABM: Agent interface.

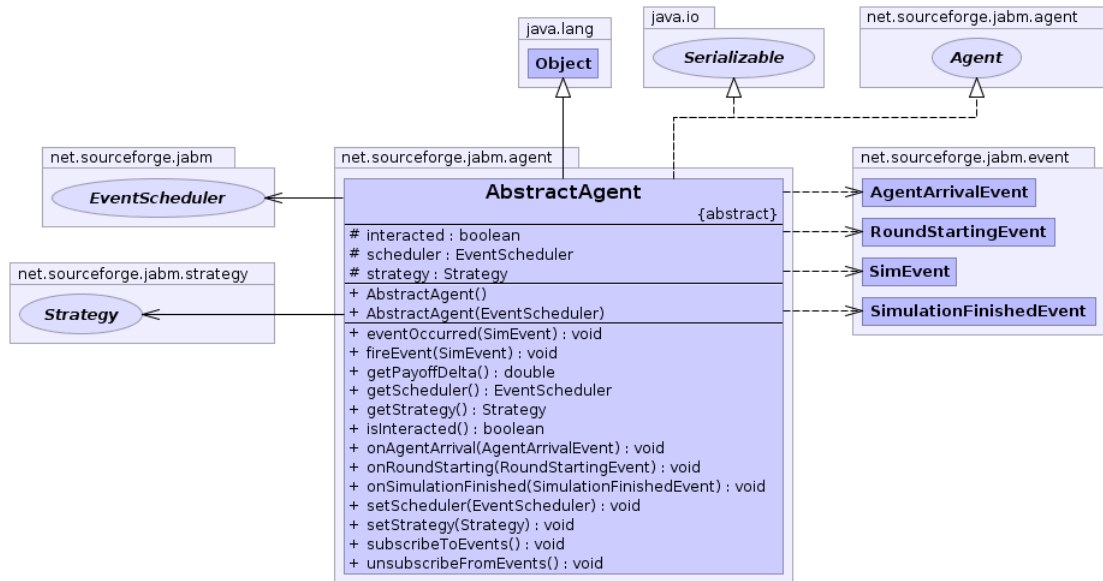


FIGURE A.6: JABM: AbstractAgent class.

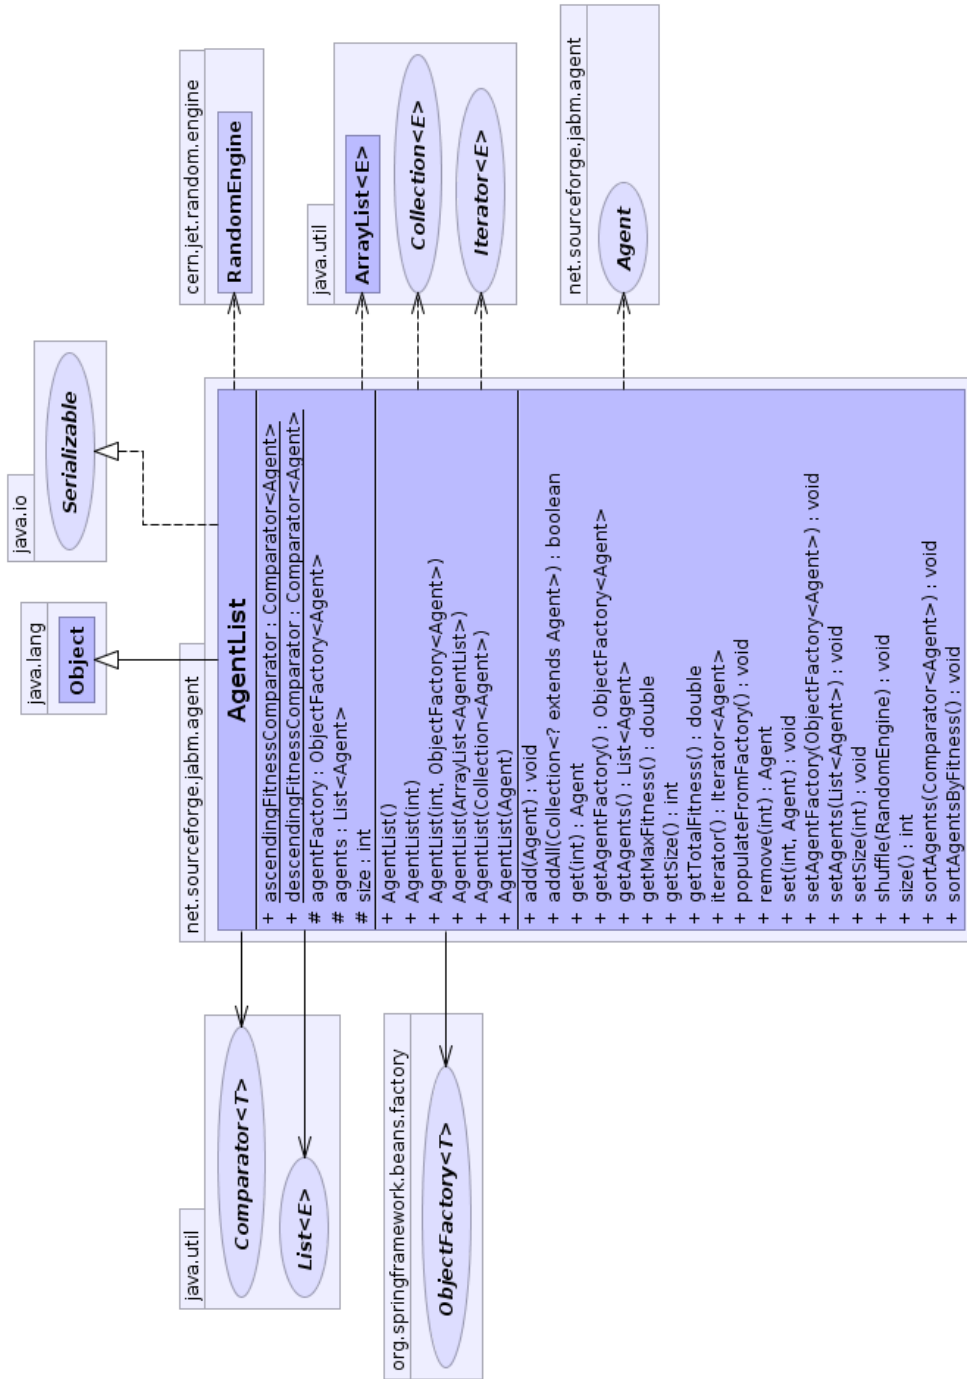


FIGURE A.7: JABM: AgentList class.

A.2 Cognitive model classes

The main classes of the implemented conceptual cognitive model are shown . Figure A.8 shows the classes that are used to represent the artificial cognitive agent. Each agent is represented by the Agent class. Each Agent class has a BBC and BDC comparators, a working memory and a long-term memory. The long-term memory stores a set of emotional memories or episodics. Figure A.9 shows the BBC and BDC classes. For each new belief, while the BDC yields either a DesireFullfilmentSignal or a DesireFrustrationSignal, the BBC yields either a BeliefConfirmationSignal or a BeliefDisconfirmationSignal.

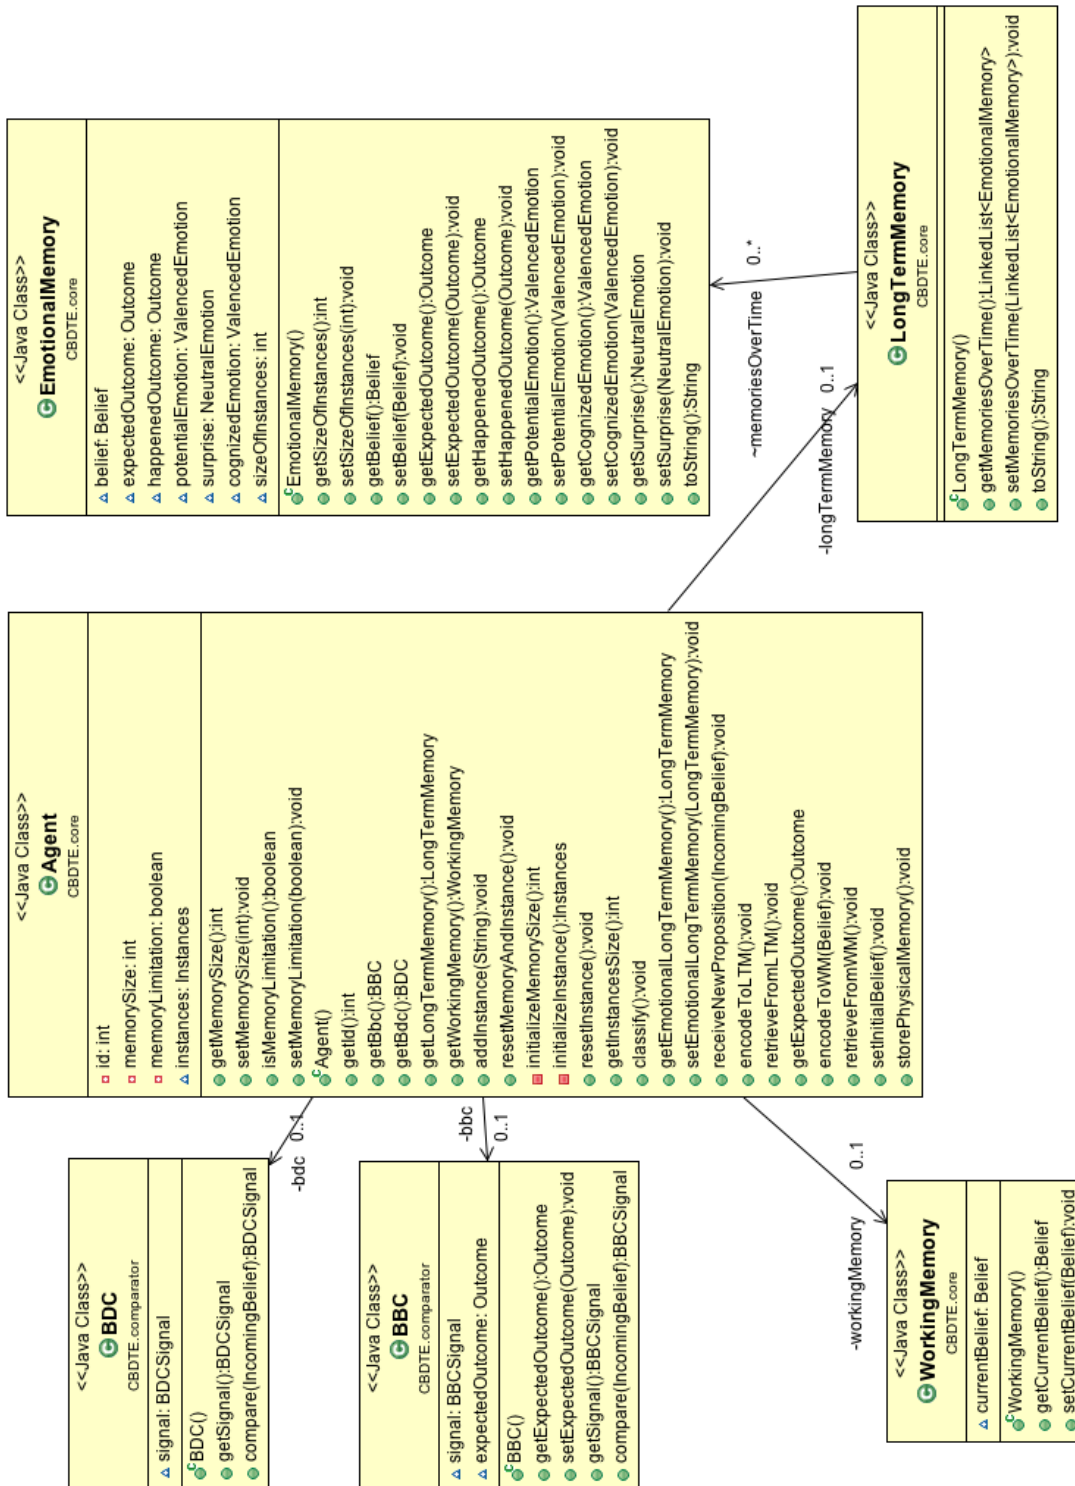


FIGURE A.8: Conceptual model: main classes of the artificial cognitive agent.

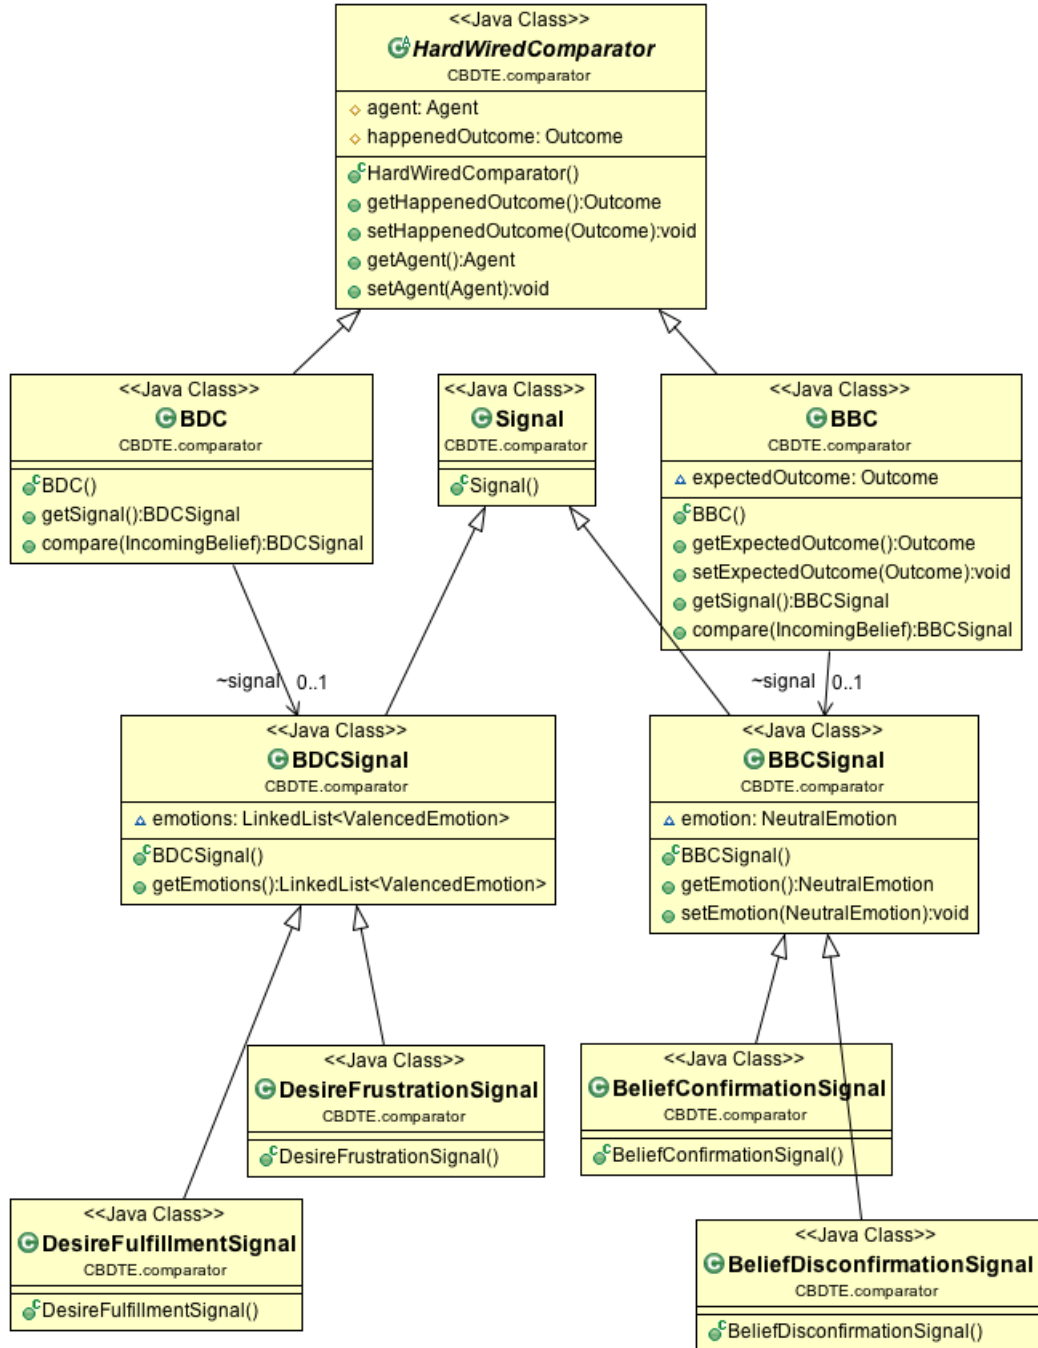


FIGURE A.9: Conceptual model: BBC and BDC classes.

Bibliography

- Ackert, L. F., Church, B. K., and Ely, K. (2008). Biases in Individual Forecasts: Experimental Evidence. *Journal of Behavioral Finance*, 9(2):53–61.
- Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.
- Akerlof, G. A. and Shiller, R. J. (2009). *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton University Press.
- Alexander, C. and Sarabia, J. M. (2012). Quantile Uncertainty and Value-at-Risk Model Risk. *Risk Analysis*, 32(8):1293–1308.
- Alfi, V., Cristelli, M., Pietronero, L., and Zaccaria, A. (2009a). Minimal agent based model for financial markets i. *The European Physical Journal B - Condensed Matter and Complex Systems*, 67(3):385–397.
- Alfi, V., Cristelli, M., Pietronero, L., and Zaccaria, A. (2009b). Minimal agent based model for financial markets II. *The European Physical Journal B - Condensed Matter and Complex Systems*, 67(3):399–417.
- Andersson, M., Hedesstrom, M., and Garling, T. (2014). A Social-Psychological Perspective on Herding in Stock Markets. *Journal of Behavioral Finance*, 15(3):226–234.
- Antoniou, A., Ergul, N., Holmes, P., and Priestley, R. (1997). Technical analysis, trading volume and market efficiency: evidence from an emerging market. *Applied Financial Economics*, 7(4):361.
- Apergis, N., Artikis, P. G., and Kyriazis, D. (2015). Does stock market liquidity explain real economic activity? New evidence from two large European stock markets. *Journal of International Financial Markets, Institutions and Money*, 38:42–64.

- Ariely, D. (2010). *Predictably Irrational, Revised and Expanded Edition: The Hidden Forces That Shape Our Decisions*. Harper Perennial, New York, NY, 1 exp rev edition edition.
- Arthur, W. (1993). On designing economic agents that behave like human agents. *Journal of Evolutionary Economics*, 3(1):1–22.
- Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American Economic Review*, 84(2):406–411.
- Arthur, W. B., Holl, J. H., Lebaron, B., Palmer, R., and Tayler, P. (1997). Asset pricing under endogenous expectations in an artificial stock market. pages 15–44.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228.
- Aruoba, S. B., Diebold, F. X., and Scotti, C. (2009). Real-Time Measurement of Business Conditions. *Journal of Business & Economic Statistics*, 27(4):417–427.
- Assenza, T., Delli Gatti, D., and Grazzini, J. (2015). Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control*, 50:5–28.
- Axelrod, R. (1997). Advancing the art of simulation in the social sciences. *Complex*, 3(2):16–22.
- Axmacher, N., Cohen, M. X., Fell, J., Haupt, S., Dümpelmann, M., Elger, C. E., Schlaepfer, T. E., Lenartz, D., Sturm, V., and Ranganath, C. (2010). Intracranial EEG Correlates of Expectancy and Memory Formation in the Human Hippocampus and Nucleus Accumbens. *Neuron*, 65(4):541–549.
- Azar, P. D. and Lo, A. W. (2016). The Wisdom of Twitter Crowds: Predicting Stock Market Reactions to FOMC Meetings via Twitter Feeds. *The Journal of Portfolio Management*, 42(5):123–134.
- Azariadis, C. (1981). Self-fulfilling prophecies. *Journal of Economic Theory*, 25(3):380–396.
- Baccan, D., Macedo, L., and Sbruzzi, E. (2014a). Is the El Farol more efficient when cognitive rational agents have a larger memory size? In *IEEE International Conference*

-
- on Systems, Man, and Cybernetics (SMC) 2014*, pages 39–44, San Diego, California, USA. IEEE.
- Baccan, D., Macedo, L., and Sbruzzi, E. (2014b). Towards modeling surprise in economics and finance: a cognitive science perspective. In *STAIRS 2014*, volume 264 of *Frontiers in Artificial Intelligence and Applications*, pages 31–40, Prague, Czech Republic.
- Baccan, D. D. and Macedo, L. (2013a). Are markets more efficient when cognitive agents act rationally? evidence from the el farol. In *16th Portuguese Conference on Artificial Intelligence*, pages 438–449, Azores, Portugal.
- Baccan, D. D. and Macedo, L. (2013b). Revisiting the el farol problem: A cognitive modeling approach. In Giardini, F. and Amblard, F., editors, *Multi-Agent-Based Simulation XIII*, number 7838 in *Lecture Notes in Computer Science*, pages 56–68. Springer Berlin Heidelberg.
- Baccan, D. D., Sbruzzi, E., and Macedo, L. M. M. L. (2015). On the Cognitive Surprise in Risk Management: An Analysis of the Value-at-Risk (VaR) Historical. In *17th Portuguese Conference on Artificial Intelligence*, pages 402–413, Coimbra, Portugal.
- Baddeley, A., Eysenck, M., and Anderson, M. C. (2009). *Memory*. Psychology Press, 1 edition.
- Bak, P., Paczuski, M., and Shubik, M. (1997). Price variations in a stock market with many agents. *Physica A: Statistical Mechanics and its Applications*, 246(3-4):430–453.
- Baker, S. R., Bloom, N., and Davis, S. J. (2015). Measuring Economic Policy Uncertainty. Working Paper 21633, National Bureau of Economic Research.
- Baldi, P. and Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*, 23(5):649–666.
- Ball, R. and Brown, P. (1968). An Empirical Evaluation of Accounting Income Numbers. *Journal of Accounting Research*, 6(2):159–178.
- Banerjee, A. and Marcellino, M. (2006). Are there any reliable leading indicators for US inflation and GDP growth? *International Journal of Forecasting*, 22(1):137–151.

- Bar-Eli, M., Avugos, S., and Raab, M. (2006). Twenty years of “hot hand” research: Review and critique. *Psychology of Sport and Exercise*, 7(6):525–553.
- Bar-Hillel, M. and Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics*, 12(4):428–454.
- Barber, B. M., Lee, Y.-T., Liu, Y.-J., and Odean, T. (2009). Just How Much Do Individual Investors Lose by Trading? *Review of Financial Studies*, 22(2):609–632.
- Barber, B. M. and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics*, 116(1):261–292.
- Barberis, N. (2013). The Psychology of Tail Events: Progress and Challenges. *American Economic Review*, 103(3):611–16.
- Barberis, N., Huang, M., and Santos, T. (2001). Prospect Theory and Asset Prices. *The Quarterly Journal of Economics*, 116(1):1–53.
- Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343.
- Barberis, N. and Thaler, R. (2002). A Survey of Behavioral Finance. *National Bureau of Economic Research Working Paper Series*, No. 9222:71.
- Baron, J. M. and Bluck, S. (2009). Autobiographical memory sharing in everyday life: Characteristics of a good story. *International Journal of Behavioral Development*, 33(2):105–117.
- Barro, R. J. (1990). The Stock Market and Investment. *Review of Financial Studies*, 3(1):115–131.
- Bartlett, F. C. (1932). *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, New York, 1 edition.
- Bartov, E., Givoly, D., and Hayn, C. (2002). The rewards to meeting or beating earnings expectations. *Journal of Accounting and Economics*, 33(2):173–204.
- Batten, D. (2000). *Discovering Artificial Economics: How Agents Learn and Economies Evolve*. Westview Press, 0 edition.

- Beaman, A., Pushkar, D., Etezadi, S., Bye, D., and Conway, M. (2007). Autobiographical memory specificity predicts social problem-solving ability in old and young adults. *The Quarterly Journal of Experimental Psychology*, 60(9):1275.
- Beber, A., Brandt, M. W., and Kavajecz, K. A. (2011). What Does Equity Sector Orderflow Tell Us About the Economy? *Review of Financial Studies*, 24(11):3688–3730.
- Bechara, A. (2004). The role of emotion in decision-making: Evidence from neurological patients with orbitofrontal damage. *Brain and Cognition*, 55(1):30–40.
- Bechara, A. and Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior*, 52(2):336–372.
- Bechara, A. and Reimann, M. (2010). The somatic marker framework as a neural theory of decision-making: Review, conceptual comparisons, and future neuroeconomic research. *Journal of Economic Psychology*, In Press, Accepted Manuscript.
- Beckmann, D., Menkhoff, L., and Suto, M. (2008). Does culture influence asset managers' views and behavior? *Journal of Economic Behavior & Organization*, 67(3-4):624–643.
- Bell, D. E. (1983). Risk premiums for decision regret. *Management Science*, 29(10):1156–1166.
- Benartzi, S. and Thaler, R. H. (1995). Myopic Loss Aversion and the Equity Premium Puzzle. *The Quarterly Journal of Economics*, 110(1):73–92.
- Benos, A. V. (1998). Aggressiveness and survival of overconfident traders. *Journal of Financial Markets*, 1(3-4):353–383.
- Berg, J. E., Nelson, F. D., and Rietz, T. A. (2008). Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2):285–300.
- Berkowitz, J. and O'Brien, J. (2002). How Accurate Are Value-at-Risk Models at Commercial Banks? *Journal of Finance*, 57(3):1093–1111.
- Bernanke, B. S. (1983). Non-Monetary Effects of the Financial Crisis in the Propagation of the Great Depression. Working Paper 1054, National Bureau of Economic Research.

- Bernard, V. L. and Thomas, J. K. (1989). Post-earnings-announcement drift: Delayed price response or risk premium? *Journal of Accounting Research*, 27:1–36.
- Bessembinder, H. and Chan, K. (1995). The profitability of technical trading rules in the asian stock markets. *Pacific-Basin Finance Journal*, 3(2-3):257–284.
- Biais, B., Glosten, L., and Spatt, C. (2005). Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2):217–264.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5):992–1026.
- Binder, J. R. and Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, 15(11):527–536.
- Biondo, A. E., Pluchino, A., Rapisarda, A., and Helbing, D. (2013). Are random trading strategies more successful than technical ones? *PLoS ONE*, 8(7):e68344.
- Black, F. (1986). Noise. *Journal of Finance*, 41(3):529–543.
- Blanchard, O. and Johnson, D. R. (2012). *Macroeconomics*. Pearson, 6th edition edition.
- Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.
- Bluck, S. and Alea, N. (2009). Thinking and talking about the past: Why remember? *Applied Cognitive Psychology*, 23(8):1089–1104.
- Blume, L. and Easley, D. (1992). Evolution and market behavior. *Journal of Economic Theory*, 58(1):9–40.
- Blume, L., Easley, D., Kleinberg, J., Kleinberg, R., and Tardos, (2015). Introduction to computer science and economic theory. *Journal of Economic Theory*, 156:1–13.
- Blume, L., Easley, D., and O’Hara, M. (1994). Market statistics and technical analysis: The role of volume. *Journal of Finance*, 49(1):153–81.
- Boden, M. (1995). Creativity and unpredictability. *Stanford Humanities Review*, 4(2):123–139.

- Bodie, Z., Kane, A., and Marcus, A. (2004). *Investments*. McGraw-Hill/Irwin, 6 edition.
- Boehner, K., DePaula, R., Dourish, P., and Sengers, P. (2007). How emotion is made and measured. *International Journal of Human-Computer Studies*, 65(4):275–291.
- Boer-Sorban, K. (2008). *Agent-Based Simulation of Financial Markets*. PhD thesis, Erasmus University Rotterdam.
- Bollen, J., Mao, H., and Zeng, X.-J. (2010). Twitter mood predicts the stock market. *1010.3003*.
- Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 99 Suppl 3:7280–7287.
- Bond, P., Edmans, A., and Goldstein, I. (2011). The Real Effects of Financial Markets. Working Paper 17719, National Bureau of Economic Research.
- Boswijk, H. P., Hommes, C. H., and Manzan, S. (2007). Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control*, 31(6):1938–1970.
- Bouchaud, J. (2008). Economics needs a scientific revolution. *Nature*, 455(7217):1181.
- Bouchaud, J., Mezard, M., and Potters, M. (2002). Statistical properties of stock order books: empirical results and models. *Quantitative Finance*, 2(4):251–256.
- Bower, G. H. (1981). Mood and memory. *American Psychologist*, 36(2):129–148.
- Brand, M., Labudda, K., and Markowitsch, H. J. (2006). Neuropsychological correlates of decision-making in ambiguous and risky situations. *Neural Networks*, 19(8):1266–1276.
- Brandouy, O. (2005). Stock markets as Minority Games: cognitive heterogeneity and equilibrium emergence. *Physica A: Statistical Mechanics and its Applications*, 349(1–2):302–328.
- Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information.
- Bratman, M., Israel, D., and Pollack, M. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(3):349–355.

- Brennan, T. J. and Lo, A. W. (2011). The Origin of Behavior. *The Quarterly Journal of Finance*, 01:55.
- Brennan, T. J. and Lo, A. W. (2012). An evolutionary model of bounded rationality and intelligence. *PLoS ONE*, 7(11):e50310.
- Brock, W., Lakonishok, J., and LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, 47(5):1731–64.
- Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8-9):1235–1274.
- Brown, L. D. (2001). A temporal analysis of earnings surprises: Profits versus losses. *Journal of Accounting Research*, 39(2):221–241.
- Brown, L. D. (2003). Small negative surprises: frequency and consequence. *International Journal of Forecasting*, 19(1):149–159.
- Brown, R. and Kulik, J. (1977). Flashbulb memories. *Cognition*, 5(1):73–99.
- Brunnermeier, M. K. (2009). Deciphering the Liquidity and Credit Crunch 2007-2008. *Journal of Economic Perspectives*, 23(1):77–100.
- Buchanan, M. (2009). Meltdown modelling: Could agent-based computer models prevent another financial crisis? *Nature*, 460(6):680–682.
- Buchanan, T. W. (2007). Retrieval of emotional memories. *Psychological bulletin*, 133(5):761–779. PMID: 17723029 PMCID: 2265099.
- Buffet, W. (1996). Annual report berkshire hathaway.
- Caballero, R. J. (2010). Macroeconomics after the Crisis: Time to Deal with the Pretense-of-Knowledge Syndrome. *Journal of Economic Perspectives*, 24(4):85–102.
- Cabrera-Paniagua, D., Cubillos, C., Vicari, R., and Urra, E. (2015). Decision-making system for stock exchange market using artificial emotions. *Expert Systems with Applications*, 42(20):7070–7083.
- Caginalp, G., McCabe, K., and Porter, D. (2003). The Foundations of Experimental Economics and Applications to Behavioral Finance: The Contributions of Nobel Laureate Vernon Smith. *Journal of Behavioral Finance*, 4(1):3–6.

- Caldarelli, G., Marsili, M., and Zhang, Y. (1997). A prototype model of stock exchange. *Europhysics Letters (EPL)*, 40(5):479–484.
- Camerer, C. (1995). Individual decision making. In *Handbook of Experimental Economics*, pages 587–703.
- Camerer, C., Loewenstein, G., and Prelec, D. (2005). Neuroeconomics: How Neuroscience Can Inform Economics. *Journal of Economic Literature*, 43(1):9–64.
- Camerer, C. F. (2014). Behavioral economics. *Current Biology*, 24(18):R867–R871.
- Camerer, C. F., Loewenstein, G., and Rabin, M., editors (2003). *Advances in Behavioral Economics*. Princeton University Press, Princeton, NJ, first edition, sixth printing edition edition.
- Campbell, J. Y., Lo, A. W., MacKinlay, A. C., and Lo, A. Y. (1996). *The Econometrics of Financial Markets*. Princeton University Press.
- Campbell, J. Y. and Shiller, R. J. (1988a). The dividend-price ratio and expectations of future dividends and discount factors. *Rev. Financ. Stud.*, 1(3):195–228.
- Campbell, J. Y. and Shiller, R. J. (1988b). Stock Prices, Earnings, and Expected Dividends. *The Journal of Finance*, 43(3):661–676.
- Castro, P. A. L. (2009). *Uma Arquitetura para Administração Automatizada de Ativos Baseada em Agentes Competitivos*. PhD thesis, University of São Paulo, São Paulo.
- Castro, P. A. L. and Sichman, J. S. (2009). AgEx: a financial market simulation tool for software agents. In *Enterprise Information Systems*, pages 704–715.
- Cavagna, A. (1999). Irrelevance of memory in the minority game. *Physical Review E*, 59(4):R3783–R3786.
- Cavedini, P., Riboldi, G., Keller, R., D’Annunzi, A., and Bellodi, L. (2002). Frontal lobe dysfunction in pathological gambling patients. *Biological Psychiatry*, 51(4):334–341.
- Cavenett, T. and Nixon, R. D. (2006). The effect of arousal on memory for emotionally-relevant information: A study of skydivers. *Behaviour Research and Therapy*, 44(10):1461–1469.
- Chakraborti, A., Toke, I. M., Patriarca, M., and Abergel, F. (2011a). Econophysics review: I. empirical facts. *Quantitative Finance*, 11(7):991–1012.

- Chakraborti, A., Toke, I. M., Patriarca, M., and Abergel, F. (2011b). Econophysics review: II. Agent-based models. *Quantitative Finance*, 11(7):1013–1041.
- Challet, D. (2016). Regrets, learning and wisdom. *arXiv:1605.01052 [physics, q-fin]*. arXiv: 1605.01052.
- Challet, D., Chessa, A., Marsili, M., and Zhang, Y.-C. (2001a). From minority games to real markets. *Quantitative Finance*, 1(1):168–176.
- Challet, D. and Marsili, M. (2000). Relevance of memory in minority games. *Physical Review E*, 62(2):1862–1868.
- Challet, D., Marsili, M., and Zhang, Y. (2000). Modeling market mechanism with minority game. *Physica A: Statistical Mechanics and its Applications*, 276(1-2):284–315.
- Challet, D., Marsili, M., and Zhang, Y.-C. (2001b). Stylized facts of financial markets and market crashes in Minority Games. *Physica A: Statistical Mechanics and its Applications*, 294(3–4):514–524.
- Challet, D. and Zhang, Y. C. (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical and Theoretical Physics*, 246(3-4):407–418.
- Chambers, D. and Dimson, E. (2013). Retrospectives: John Maynard Keynes, Investment Innovator. *Journal of Economic Perspectives*, 27(3):213–28.
- Chan, N. T., LeBaron, B., Lo, A. W., and Poggio, T. (1999). Agent-Based Models of Financial Markets: A Comparison with Experimental Markets. *Working Paper*.
- Chang, H.-J. (2012). *23 Things They Don't Tell You About Capitalism*. Bloomsbury Press, 1 edition.
- Chapman, T. (2003). Bringing economics into the laboratory. *Quantitative Finance*, 3(4):C58–C60.
- Chen, S.-H. (2008). Software-Agent Designs in Economics: An Interdisciplinary. *IEEE Computational Intelligence Magazine*, 3(4):18–22.
- Chen, S.-H. (2012). Varieties of agents in agent-based computational economics: A historical and an interdisciplinary perspective. *Journal of Economic Dynamics and Control*, 36(1):1–25.

- Chen, S.-S. (2011). Lack of consumer confidence and stock returns. *Journal of Empirical Finance*, 18(2):225–236.
- Chen, T. (2013). Do Investors Herd in Global Stock Markets? *Journal of Behavioral Finance*, 14(3):230–239.
- Chiarella, C. and Iori, G. (2002). A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2(5):346.
- Chicaroli, R. and Valls Pereira, P. L. (2015). Predictability of Equity Models. *Journal of Forecasting*, 34(6):427–440.
- Chiodo, A., Guidolin, M., Owyang, M. T., and Shimoji, M. (2003). Subjective probabilities: psychological evidence and economic applications. Technical report, Federal Reserve Bank of St. Louis.
- Chordia, T., Goyal, A., Sadka, G., Sadka, R., and Shivakumar, L. (2009). Liquidity and the Post-Earnings-Announcement Drift. *Financial Analysts Journal*, 65(4):18–32.
- Chordia, T. and Shivakumar, L. (2006). Earnings and price momentum. *Journal of Financial Economics*, 80(3):627–656.
- Christoffersen, P. and Pelletier, D. (2004). Backtesting Value-at-Risk: A Duration-Based Approach. *Journal of Financial Econometrics*, 2(1):84–108.
- Chui, A. C., Titman, S., and Wei, K. J. (2010). Individualism and Momentum around the World. *The Journal of Finance*, 65(1):361–392.
- Clark, L., Studer, B., Bruss, J., Tranel, D., and Bechara, A. (2014). Damage to insula abolishes cognitive distortions during simulated gambling. *Proceedings of the National Academy of Sciences*, page 201322295.
- Clarke, R. G., Krase, S., and Statman, M. (1994). Tracking errors, regret, and tactical asset allocation. *The Journal of Portfolio Management*, 20(3):16–24.
- Cochrane, J. (2005). Financial Markets and the Real Economy. NBER Working Paper 11193, National Bureau of Economic Research, Inc.
- Cochrane, J. H. (2008). The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies*, 21(4):1533–1575.

- Cochrane, J. H. (2011). Presidential Address: Discount Rates. *The Journal of Finance*, 66(4):1047–1108.
- Cohen, G. and Conway, M. A. (2008). *Memory in the Real World*. Psychology Press, 3 edition.
- Collins, A. and Loftus, E. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6):428, 407.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223.
- Cont, R. (2007). Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models. In Teysriere, D. G. and Kirman, P. A. P., editors, *Long Memory in Economics*, pages 289–309. Springer Berlin Heidelberg.
- Conway, A. R. A., Skitka, L. J., Hemmerich, J. A., and Kershaw, T. C. (2009). Flashbulb memory for 11 september 2001. *Applied Cognitive Psychology*, 23(5):605–623.
- Conway, M. A. (2005). Memory and the self. *Journal of Memory and Language*, 53(4):594–628.
- Conway, M. A. (2009). Episodic memories. *Neuropsychologia*, 47(11):2305–2313.
- Conway, M. A. and Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological Review*, 107(2):261–288. PMID: 10789197.
- Copeland, T., Dolgoff, A., and Moel, A. (2004). The role of expectations in explaining the Cross-Section of stock returns. *Review of Accounting Studies*, 9(2):149–188.
- Cosmides, L. and Tooby, J. (2013). Evolutionary Psychology: New Perspectives on Cognition and Motivation. *Annual Review of Psychology*, 64(1):201–229.
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1):185, 87.
- Cowan, N. (2008a). Sensory Memory. In *Learning and Memory: A Comprehensive Reference*, pages 23–32. Academic Press, Oxford.

- Cowan, N. (2008b). What are the differences between long-term, short-term, and working memory? In *Essence of Memory*, volume Volume 169, pages 323–338. Elsevier.
- Cristelli, M. (2011). *Complexity in Financial Markets: Modeling Psychological Behavior in Agent-Based Models and Order Book Models*. PhD thesis, Sapienza - University of Rome.
- Cross, R., Grinfeld, M., Lamba, H., and Seaman, T. (2005). A threshold model of investor psychology. *Physica A: Statistical Mechanics and its Applications*, 354:463–478.
- Cutler, D. M., Poterba, J. M., and Summers, L. H. (1989). What moves stock prices? *Journal of Portfolio Management*.
- Cutler, D. M., Poterba, J. M., and Summers, L. H. (1991). Speculative dynamics. *Review of Economic Studies*, 58(3):529–46.
- Dalgleish, T. (2004). The emotional brain. *Nature Reviews Neuroscience*, 5(7):583–589.
- Dalgleish, T. and Power, M. (1999). *Handbook of Cognition and Emotion*. Wiley, 1 edition.
- Damasio, A. R. (1994). *Descartes Error Emotion, Reason, and the Human Brain*. Grosset/Putnam, New York, 1 edition.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 351(1346):1413–1420.
- Damasio, A. R. (1999). *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Harcourt, New York, 1 edition.
- Damasio, A. R., Tranel, D., and Damasio, H. C. (1991). Somatic markers and the guidance of behavior: Theory and preliminary testing. In *Frontal lobe function and dysfunction*, pages 217–229. Oxford University Press US, New York.
- Daniel, K. and Titman, S. (1999). Market Efficiency in an Irrational World. *Financial Analysts Journal*, 55(6):28–40.
- Darwin, C. (1859). *The origin of species by means of natural selection: or, The preservation of favored races in the struggle for life*. John Murray, London.

- Darwin, C. (1872). *The Expression of the Emotions in Man and Animals*. John Murray, London.
- David Cabedo, J. and Moya, I. (2003). Estimating oil price Value at Risk using the historical simulation approach. *Energy Economics*, 25(3):239–253.
- Davidsson, P. (2000). Multi Agent Based Simulation: Beyond Social Simulation. In Moss, S. and Davidsson, P., editors, *Multi-Agent-Based Simulation*, number 1979 in Lecture Notes in Computer Science, pages 97–107. Springer Berlin Heidelberg.
- Davis, M. and Whalen, P. J. (2001). The amygdala: vigilance and emotion. *Molecular Psychiatry*, 6(1):13–34.
- De Bondt, W. F. and Thaler, R. (1985). Does the stock market overreact? *Journal of Finance*, 40(3):793–805.
- De Bondt, W. F. M. and Thaler, R. H. (1987). Further Evidence On Investor Overreaction and Stock Market Seasonality. *The Journal of Finance*, 42(3):557–581.
- De Martino, B., Camerer, C. F., and Adolphs, R. (2010). Amygdala damage eliminates monetary loss aversion. *Proceedings of the National Academy of Sciences of the United States of America*.
- DeLong, J. B. (1991). The survival of noise traders in financial markets. *Journal of Business*, 64(1):1–19.
- DeLong, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–38.
- Dopfer, K., Foster, J., and Potts, J. (2004). Micro-meso-macro. *Journal of Evolutionary Economics*, 14(3):263–279.
- Drogoul, A., Vanbergue, D., and Meurisse, T. (2003). Multi-agent Based Simulation: Where Are the Agents? In Sichman, J. S., Bousquet, F., and Davidsson, P., editors, *Multi-Agent-Based Simulation II*, number 2581 in Lecture Notes in Computer Science, pages 1–15. Springer Berlin Heidelberg.
- Duffie, D. and Pan, J. (1997). An Overview of Value at Risk. *The Journal of Derivatives*, 4(3):7–49.
- Dunn, B. D., Dalgleish, T., and Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience & Biobehavioral Reviews*, 30(2):239–271.

-
- Easley, D. and O'Hara, M. (1995). Chapter 12 market microstructure. In R.A. Jarrow, V. M. a. W. Z., editor, *Handbooks in Operations Research and Management Science*, volume Volume 9, pages 357–383. Elsevier.
- Edwards, R. and Magee, J. (2010). *Technical Analysis of Stock Trends*. Snowball Publishing.
- Eguiluz, V. M. and Zimmermann, M. G. (2000). Transmission of Information and Herd Behavior: An Application to Financial Markets. *Physical Review Letters*, 85(26):5659–5662.
- Eich, E. (1980). The cue-dependent nature of state-dependent retrieval. *Memory & Cognition*, 8(2):157–173. PMID: 7382817.
- Eich, E. (1995). Searching for mood dependent memory. *Psychological Science*, 6(2):67–75.
- Eich, E. and Macaulay, D. (2000). Are real moods required to reveal Mood-Congruent and Mood-Dependent memory? *Psychological Science*, 11(3):244–248.
- Ekman, P. (1992). An argument for basic emotions. *Cognition & Emotion*, 6(3):169.
- Ekman, P. (1999). Basic emotions. In *The Handbook of Cognition and Emotion*, pages 45–60. John Wiley & Sons, Ltd, Sussex, UK.
- Ekman, P. and Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology*, 17(2):124–129. PMID: 5542557.
- Ekman, P. and Rosenberg, E. L. (1998). *What the Face Reveals: Basic and Applied Studies of Spontaneous Expression Using the Facial Action Coding System (FACS)*. Oxford University Press, USA, 1 edition.
- Eysenck, M. W. (2001). *Principles of Cognitive Psychology*. Psychology Press Ltd, 2nd revised edition edition.
- Fagiolo, G., Moneta, A., and Windrum, P. (2006). Confronting Agent-Based Models with Data: Methodological Issues and Open Problems. In Bruun, C., editor, *Advances in Artificial Economics*, volume 584 of *Lecture Notes in Economics and Mathematical Systems*, pages 255–267. Springer Berlin Heidelberg.

- Fagiolo, G., Moneta, A., and Windrum, P. (2007). A Critical Guide to Empirical Validation of Agent-Based Models in Economics: Methodologies, Procedures, and Open Problems. *Computational Economics*, 30(3):195–226.
- Fama, E. F. (1965). Random Walks in Stock-Market Prices. *FINANCIAL ANALYSTS JOURNAL*, 21:55–59.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417.
- Fama, E. F. (1991). Efficient capital markets: II. *Journal of Finance*, 46(5):1575–617.
- Fama, E. F. (1998). Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics*, 49(3):283–306.
- Fama, E. F. (2011). My life in finance. *Annual Review of Financial Economics*, 3(1):1–15.
- Fama, E. F. (2014). Two pillars of asset pricing. *American Economic Review*, 104(6):1467–1485.
- Fama, E. F. and French, K. R. (1988). Dividend yields and expected stock returns. *Journal of Financial Economics*, 22(1):3–25.
- Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.
- Farmer, J. D. (2002). Market force, ecology and evolution. *Industrial and Corporate Change*, 11:895–953.
- Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256):685–686.
- Farmer, J. D., Gallegati, M., Hommes, C., Kirman, A., Ormerod, P., Cincotti, S., Sanchez, A., and Helbing, D. (2012). A complex systems approach to constructing better models for managing financial markets and the economy. *The European Physical Journal Special Topics*, 214(1):295–324.
- Farmer, J. D. and Lo, A. W. (1999). Frontiers of finance: Evolution and efficient markets. *Proceedings of the National Academy of Sciences of the United States of America*, 96(18):9991–9992. PMID: 10468547 PMCID: 33727.

- Farmer, J. D., Patelli, P., and Zovko, I. I. (2005). The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America*, 102(6):2254–2259.
- Farmer, R. E. (2012). The stock market crash of 2008 caused the Great Recession: Theory and evidence. *Journal of Economic Dynamics and Control*, 36(5):693–707.
- Fehr, E. and Rangel, A. (2011). Neuroeconomic Foundations of Economic Choice - Recent Advances. *Journal of Economic Perspectives*, 25(4):3–30.
- Fischer, S. and Merton, R. C. (1984). Macroeconomics and Finance: The Role of the Stock Market. Working Paper 1291, National Bureau of Economic Research.
- Fleming, J., Ostdiek, B., and Whaley, R. E. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15(3):265–302.
- Foroni, I. and Agliari, A. (2008). Complex Price Dynamics in a Financial Market with Imitation. *Computational Economics*, 32(1-2):21–36.
- Foster, F. D. and Warren, G. J. (2015). Why Might Investors Choose Active Management? *Journal of Behavioral Finance*, 16(1):20–39.
- Foster, F. D. and Warren, G. J. (2016). Interviews with Institutional Investors: The How and Why of Active Investing. *Journal of Behavioral Finance*, 17(1):60–84.
- Foster, G., Olsen, C., and Shevlin, T. (1984). Earnings Releases, Anomalies, and the Behavior of Security Returns. *The Accounting Review*, 59(4):574–603. ArticleType: research-article / Full publication date: Oct., 1984 / Copyright © 1984 American Accounting Association.
- Foster, J. (2005). From simplistic to complex systems in economics. *Cambridge Journal of Economics*, 29(6):873–892.
- Frank, R. H. (2005). *Microeconomics and Behavior*. McGraw-Hill/Irwin, 6 edition.
- Frank, R. H. (2008). *The Economic Naturalist: In Search of Explanations for Everyday Enigmas*. Basic Books.
- Frankfurter, G. M. and McGoun, E. G. (2000). Market Efficiency or Behavioral Finance: The Nature of the Debate. *Journal of Psychology and Financial Markets*, 1(3-4):200–210.

- Friedman, M. and Savage, L. (1948). The Utility Analysis of Choices Involving Risk. *The Journal of Political Economy*, 56(4):304, 279.
- Friedman, N. and Halpern, J. Y. (1997). Modeling belief in dynamic systems, part I: Foundations. *Artificial Intelligence*, 95(2):257–316.
- Frijda, N. H. (1986). *The emotions*. Cambridge University Press.
- Frijda, N. H., Manstead, A. S. R., and Bem, S. (2000). *Emotions and Beliefs: How Feelings Influence Thoughts*. Cambridge University Press, 1 edition.
- Gao, L. and Schmidt, U. (2005). Self is never neutral: Why economic agents behave irrationally. *Journal of Behavioral Finance*, 6(1):27–37.
- Gärdenfors, P. (1978). On the logic of relevance. *Synthese*, 37(3):351–367.
- Gärdenfors, P. (1992). Belief revision: An introduction. In *Belief Revision*, Cambridge Tracts in Theoretical Computer Science, pages 1–20. Cambridge University Press.
- Gärdenfors, P. (1994). The Role of Expectations in Reasoning. In *International Conference Logic at Work on Knowledge Representation and Reasoning Under Uncertainty, Logic at Work*, pages 1–16. Springer-Verlag. ACM ID: 745051.
- Gatti, D., Gaffeo, E., and Gallegati, M. (2010). Complex agent-based macroeconomics: a manifesto for a new paradigm. *Journal of Economic Interaction and Coordination*, 5(2):111–135.
- Geithner, T. (2012). Financial Crisis Amnesia. *Wall Street Journal*.
- Gervais, S. and Odean, T. (2001). Learning to be overconfident. *Review of Financial Studies*, 14(1):1–27.
- Getchell, A. (2008). Agent-Based modeling. Working paper, Department of Physics, UC Davis.
- Getmansky, M., Lee, P. A., and Lo, A. W. (2015). Hedge Funds: A Dynamic Industry In Transition. Working Paper 21449, National Bureau of Economic Research.
- Giardina, I. and Bouchaud, J.-P. (2003). Bubbles, crashes and intermittency in agent based market models. *The European Physical Journal B - Condensed Matter and Complex Systems*, 31(3):421–437.

- Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic Decision Making. *Annual Review of Psychology*, 62(1):451–482.
- Gilboa, I., Postlewaite, A. W., and Schmeidler, D. (2008). Probability and uncertainty in economic modeling. *Journal of Economic Perspectives*, 22(3):173–188.
- Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3):295–314.
- Glimcher, P. W. and Fehr, E. (2013). *Neuroeconomics: Decision Making and the Brain*. Academic Press.
- Gode, D. K. and Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1):119–37.
- Goodwin, D. W., Crane, J. B., and Guze, S. B. (1969). Alcoholic "blackouts": a review and clinical study of 100 alcoholics. *The American Journal of Psychiatry*, 126(2):191–198. PMID: 5804804.
- Gorton, G. (2009). Information, Liquidity, and the (Ongoing) Panic of 2007. *American Economic Review*, 99(2):567–572.
- Gorton, G. B. (2008). The Panic of 2007. Working Paper 14358, National Bureau of Economic Research.
- Gourieroux, C., Laurent, J., and Scaillet, O. (2000). Sensitivity analysis of Values at Risk. *Journal of Empirical Finance*, 7(3–4):225–245.
- Graham, B. (1976). A Conversation with Benjamin Graham. *Financial Analysts Journal*, 32(5):20–23.
- Graham, B. (1986). *The Intelligent Investor: A Book of Practical Counsel*. Harper & Row Publishers, 4th edition.
- Greenwood, J. and Smith, B. D. (1997). Financial markets in development, and the development of financial markets. *Journal of Economic Dynamics and Control*, 21(1):145–181.
- Griffin, D. and Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3):411–435.

- Grossman, S. (1976). On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance*, 31(2):573–585.
- Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3):393–408.
- Grune, L. and Semmler, W. (2008). Asset pricing with loss aversion. *Journal of Economic Dynamics and Control*, 32(10):3253–3274.
- Gualdi, S., Tarzia, M., Zamponi, F., and Bouchaud, J.-P. (2015). Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, 50:29–61.
- Gupta, R., Duff, M. C., Denburg, N. L., Cohen, N. J., Bechara, A., and Tranel, D. (2009). Declarative memory is critical for sustained advantageous complex decision-making. *Neuropsychologia*, 47(7):1686–1693.
- Gupta, R., Kosciak, T. R., Bechara, A., and Tranel, D. (2011). The amygdala and decision making. *Neuropsychologia*, 49(4):760–766.
- Gutknecht, O. and Ferber, J. (2001). The MADKIT agent platform architecture. In *Revised Papers from the International Workshop on Infrastructure for Multi-Agent Systems: Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems*, pages 48–55. Springer-Verlag.
- Hagstrom, R. G. (2005). *The Warren Buffett Way*. Wiley, 2 edition.
- Halbleib, R. and Pohlmeier, W. (2012). Improving the value at risk forecasts: Theory and evidence from the financial crisis. *Journal of Economic Dynamics and Control*, 36(8):1212–1228.
- Hall, R. E. (2010). Why Does the Economy Fall to Pieces after a Financial Crisis? *Journal of Economic Perspectives*, 24(4):3–20.
- Hardt, O., Nader, K., and Nadel, L. (2013). Decay happens: the role of active forgetting in memory. *Trends in Cognitive Sciences*, 17(3):111–120.
- Harper, D. A. and Lewis, P. (2012). New perspectives on emergence in economics. *Journal of Economic Behavior & Organization*, 82(2–3):329–337.
- Hayek, F. (1945). The use of knowledge in society. *American Economic Review*, 35(4):519–530.

-
- Hellwig, M. F. (1980). On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3):477–498.
- Hendricks, D. (1996). Evaluation of value-at-risk models using historical data. *Economic Policy Review*, (Apr):39–69.
- Hertwig, R., Barron, G., Weber, E. U., and Erev, I. (2004). Decisions from Experience and the Effect of Rare Events in Risky Choice. *Psychological Science*, 15(8):534–539.
- Hirshleifer, D. (2001). Investor psychology and asset pricing. *The Journal of Finance*, 56(4):1533–1597.
- Hirshleifer, D. (2015). Behavioral Finance. *Annual Review of Financial Economics*, 7(1):133–159.
- Hirshleifer, D. and Hong Teoh, S. (2003). Herd Behaviour and Cascading in Capital Markets: a Review and Synthesis. *European Financial Management*, 9(1):25–66.
- Hirshleifer, D. and Luo, G. Y. (2001). On the survival of overconfident traders in a competitive securities market. *Journal of Financial Markets*, 4(1):73–84.
- Hoffmann, A., Delre, S., Eije, J., and Jager, W. (2006). Stock price dynamics in artificial Multi-Agent stock markets. In *Artificial Economics*, pages 191–201.
- Holland, A. C. and Kensinger, E. A. (2010). Emotion and autobiographical memory. *Physics of Life Reviews*, 7(1):88–131.
- Holland, J. H. and Miller, J. H. (1991). Artificial Adaptive Agents in Economic Theory. *American Economic Review*, 81(2):365–71.
- Hommes, C. (2001). Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance*, 1(1):149–167.
- Hommes, C. and Iori, G. (2015). Introduction special issue crises and complexity. *Journal of Economic Dynamics and Control*, 50:1–4.
- Horn, C. F., Ivens, B. S., Ohneberg, M., and Brem, A. (2014). Prediction Markets – A literature review 2014. *The Journal of Prediction Markets*, 8(2):89–126.
- Hull, J. C. (2005). *Options, Futures and Other Derivatives*. Prentice Hall, Upper Saddle River, NJ, 6th edition edition.

- Itti, L. and Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49(10):1295–1306.
- Ivashina, V. and Scharfstein, D. (2010). Bank lending during the financial crisis of 2008. *Journal of Financial Economics*, 97(3):319–338.
- James, W. (1884). What is an emotion? *Mind*, 9:188–205.
- James, W. (1890). *The principles of psychology*. Holt, New York.
- Janssen, M. A. and Ostrom, E. (2006). Empirically Based, Agent-based models. *Ecology and Society*, 11(2).
- Jen, F. C. (1970). Random Walks and Technical Theories: Some Additional Evidence: Discussion. *The Journal of Finance*, 25(2):495–499. ArticleType: primary_article / Issue Title: Papers and Proceedings of the Twenty-Eighth Annual Meeting of the American Finance Association New York, N.Y. December, 28-30, 1969 / Full publication date: May, 1970 / Copyright © 1970 American Finance Association.
- Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6(2-3):95–101.
- Jordan, S. J. Is momentum a self-fulfilling prophecy? *Quantitative Finance*, 0(0):1–12.
- Jorion, P. (2006). *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill; 3rd edition.
- Kaestner, M. (2006a). Anomalous price behavior following earnings surprises: Does representativeness cause overreaction? *Revue Finance*, 27(2).
- Kaestner, M. (2006b). Investors' misreaction to unexpected earnings: Evidence of simultaneous overreaction and underreaction. *Behavioral Finance*, 3.
- Kahana, M., Howard, M., Polyn, S., and Byrne, J. H. (2008). Associative retrieval processes in episodic memory. In *Learning and Memory: A Comprehensive Reference*, pages 1–24. Academic Press, Oxford.
- Kahneman, D. (2009). The myth of risk attitudes. *The Journal of Portfolio Management*, 36(1):1.
- Kahneman, D. and Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454.

- Kahneman, D. and Tversky, A. (1979a). Intuitive prediction: Biases and corrective procedures. *Management Science*, 12:313–327.
- Kahneman, D. and Tversky, A. (1979b). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.
- Karino, Y. and Kawagoe, T. (2009). Bubble and crash in the artificial financial market. In *Artificial Economics*, pages 159–170.
- Kaszniak, R. and McNichols, M. F. (2002). Does meeting earnings expectations matter? evidence from analyst forecast revisions and share prices. *Journal of Accounting Research*, 40(3):727–759.
- Kawata, R. and Kijima, M. (2007). Value-at-risk in a market subject to regime switching. *Quantitative Finance*, 7(6):609–619.
- Kearns, M. and Ortiz, L. (2003). The Penn-Lehman automated trading project. *IEEE Intelligent Systems*, 18(6):22–31.
- Kelemen, W. L. and Creeley, C. E. (2003). State-dependent memory effects using caffeine and placebo do not extend to metamemory. *The Journal of General Psychology*, 130(1):70–86. PMID: 12635857.
- Kelly, M. (1997). Do noise traders influence stock prices? *Journal of Money, Credit and Banking*, 29(3):351–363.
- Kensinger, E. A. (2009a). How emotion affects older adults' memories for event details. *Memory (Hove, England)*, 17(2):208–219. PMID: 18608972.
- Kensinger, E. A. (2009b). Remembering the details: Effects of emotion. *Emotion Review*, 1(2):99–113. PMID: 19421427.
- Kensinger, E. A., Krendl, A. C., and Corkin, S. (2005). Memories of an emotional and a nonemotional event: Effects of aging and delay interval. *Experimental Aging Research: An International Journal Devoted to the Scientific Study of the Aging Process*, 32(1):23.
- Keynes, J. M. (1936). *The General Theory Of Employment, Interest, And Money*. Macmillan, London.
- Kihlstrom, J. F. (2009). "So that we might have roses in december": The functions of autobiographical memory. *Applied Cognitive Psychology*, 23(8):1179–1192.

- Kindleberger, C. P., Aliber, R., and Solow, R. (2005). *Manias, Panics, and Crashes: A History of Financial Crises*. Wiley, Hoboken, N.J, 5 edition edition.
- Kirman, A. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, 108(1):137–156.
- Kirman, A. (2006). Heterogeneity in economics. *Journal of Economic Interaction and Coordination*, 1(1):89–117.
- Kirman, A. (2010). The Economic Crisis is a Crisis for Economic Theory. *CEISifo Economic Studies*, 56(4):498–535.
- Klimek, P., Poledna, S., Doyne Farmer, J., and Thurner, S. (2015). To bail-out or to bail-in? Answers from an agent-based model. *Journal of Economic Dynamics and Control*, 50:144–154.
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Boston and New York.
- Kodia, Z. and Said, L. (2009). Multi-agent simulation of investor cognitive behavior in stock market. In *7th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2009)*, pages 90–99.
- Kodia, Z., Said, L., and Ghedira, K. (2010a). A multi-agent based pricing: a virtual stock market simulation. In *8th International Conference of Modeling and Simulation*, Tunisia.
- Kodia, Z., Said, L. B., and Ghedira, K. (2009). SiSMar: social multi-agent based simulation of stock market. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, pages 1345–1346, Budapest, Hungary. International Foundation for Autonomous Agents and Multiagent Systems.
- Kodia, Z., Said, L. B., and Ghedira, K. (2010b). Towards a new cognitive modeling approach for multi-agent based simulation of stock market dynamics. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1*, pages 1565–1566, Toronto, Canada. International Foundation for Autonomous Agents and Multiagent Systems.
- Kothari, S., Lewellen, J., and Warner, J. B. (2006). Stock returns, aggregate earnings surprises, and behavioral finance. *Journal of Financial Economics*, 79(3):537–568.

- Krajbich, I., Oud, B., and Fehr, E. (2014). Benefits of Neuroeconomic Modeling: New Policy Interventions and Predictors of Preference. *American Economic Review*, 104(5):501–506.
- Kwon, K.-Y. and Kish, R. J. (2002). A comparative study of technical trading strategies and return predictability: an extension of using NYSE and NASDAQ indices. *The Quarterly Review of Economics and Finance*, 42(3):611–631.
- Ladley, D. (2012). Zero intelligence in economics and finance. *The Knowledge Engineering Review*, 27(Special Issue 02):273–286.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–77.
- Landefeld, J. S., Seskin, E. P., and Fraumeni, B. M. (2008). Taking the pulse of the economy: Measuring GDP. *Journal of Economic Perspectives*, 22(2):193–216.
- Lazarus, R. S. (1991). *Emotion and Adaptation*. Oxford University Press, New York.
- LeBaron, B. (1999). Technical trading rule profitability and foreign exchange intervention. *Journal of International Economics*, 49(1):125–143.
- LeBaron, B. (2000). Agent-based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control*, 24(5-7):679–702.
- LeBaron, B. (2001). A builder’s guide to agent-based financial markets. *Quantitative Finance*, 1(2):254.
- LeBaron, B. (2002). Building the santa fe artificial stock market. *School of International Economics and Finance*, 21:1117–1147.
- LeBaron, B. (2006). Agent-based computational finance. volume 2 of *Handbook of Computational Economics*, pages 1187–1233. Elsevier.
- LeDoux, J. E. (2000). Emotion Circuits in the Brain. *Annual Review of Neuroscience*, 23(1):155–184.
- Lee, C.-F., Tsai, C.-M., and Lee, A. C. (2013). Asset pricing with disequilibrium price adjustment: theory and empirical evidence. *Quantitative Finance*, 13(2):227–239.

- Lejarraga, T., Pachur, T., Frey, R., and Hertwig, R. (2016). Decisions from Experience: From Monetary to Medical Gambles. *Journal of Behavioral Decision Making*, 29(1):67–77.
- Lempérière, Y., Deremble, C., Seager, P., Potters, M., and Bouchaud, J. P. (2014). Two centuries of trend following. *arXiv:1404.3274 [q-fin]*. arXiv: 1404.3274.
- Lengnick, M. (2013). Agent-based macroeconomics: A baseline model. *Journal of Economic Behavior & Organization*, 86:102–120.
- Lettau, M. (1997). Explaining the facts with adaptive agents: The case of mutual fund flows. *Journal of Economic Dynamics and Control*, 21(7):1117–1147.
- Levine, L. J., Lench, H. C., and Safer, M. A. (2009). Functions of remembering and misremembering emotion. *Applied Cognitive Psychology*, 23(8):1059–1075.
- Lewis, M. (2011). *The Big Short: Inside the Doomsday Machine*. W. W. Norton & Company, reprint edition.
- Lewis, P., Critchley, H., Smith, A., and Dolan, R. (2005). Brain mechanisms for mood congruent memory facilitation. *NeuroImage*, 25(4):1214–1223.
- Lewis, P. A. and Critchley, H. D. (2003). Mood-dependent memory. *Trends in Cognitive Sciences*, 7(10):431–433.
- Lichtenstein, S., Fischhoff, B., and Phillips, L. D. (1982). Calibration of probabilities: the state of the art to 1980. In *Judgment under uncertainty: heuristics and biases*, pages 306–334. Cambridge University Press, Cambridge.
- Lo, A. (2016a). Imagine if Robo Advisers Could Do Emotions. *Blog Wall Street Journal*.
- Lo, A. and Mueller, M. (2010). WARNING: physics envy may be hazardous to your wealth! *Journal of Investment Management*, 8:13–63.
- Lo, A. W. (1999). The three p’s of total risk management. *Financial Analysts Journal*, 55:13–26.
- Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, pages 15–29.
- Lo, A. W. (2005). Reconciling efficient markets with behavioral finance: The adaptive markets hypothesis. *Journal of Investment Consulting*, 7:21–44.

- Lo, A. W. (2010). Survival of the richest. *The Journal of Portfolio Management*, 36(3):4.
- Lo, A. W. (2012). Adaptive Markets and the New World Order. *Financial Analysts Journal*, 68(2):18–29.
- Lo, A. W. (2016b). What Is an Index? *The Journal of Portfolio Management*, 42(2):21–36.
- Lo, A. W., Mamaysky, H., and Wang, J. (2000). Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation. *Journal of Finance*, 55:1705–1765.
- Loewenstein, G., Rick, S., and Cohen, J. D. (2008). Neuroeconomics. *Annual Review of Psychology*, 59(1):647–672.
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., and Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, 127(2):267–286.
- Lorini, E. and Castelfranchi, C. (2006). The unexpected aspects of surprise. In *International Journal of Pattern Recognition and Artificial Intelligence*, pages 817–835.
- Lorini, E. and Castelfranchi, C. (2007). The cognitive structure of surprise: looking for basic principles. *Topoi: An International Review of Philosophy*, 26:133–149.
- Lovric, M. (2011a). *Behavioral Finance and Agent-Based Artificial Markets*. PhD thesis, Erasmus University Rotterdam.
- Lovric, M. (2011b). *Behavioral Finance and Agent-Based Artificial Markets*. PhD thesis, Erasmus University Rotterdam.
- Lu, Y., Wang, J., and Niu, H. (2015). Agent-based financial dynamics model from stochastic interacting epidemic system and complexity analysis. *Physics Letters A*, 379(14–15):1023–1031.
- Lux, T. (1998). The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. *Journal of Economic Behavior & Organization*, 33(2):143–165.
- Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719):498–500.

- Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: a microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance*, 03(04):675–702.
- Lynch, P. (1994). *Beating the Street*. Simon & Schuster, revised edition.
- Macal, C. M. and North, M. J. (2005). Tutorial on agent-based modeling and simulation. In *Proceedings of the 37th conference on Winter simulation*, pages 2–15, Orlando, Florida. Winter Simulation Conference.
- Macal, C. M. and North, M. J. (2006). Tutorial on agent-based modeling and simulation part 2: how to model with agents. In *Proceedings of the 38th conference on Winter simulation*, pages 73–83, Monterey, California. Winter Simulation Conference.
- Macedo, L. (2006). *The Exploration of Unknown Environments by Affective Agents*. PhD thesis, University of Coimbra.
- Macedo, L. (2010). A surprise-based selective attention agent for travel information. In *6th Workshop on Agents in Traffic and Transportation, Autonomous Agents and Multiagent Systems*, Toronto, Canada.
- Macedo, L. (2013). Arguments for a Computational Model for Forms of Selective Attention based on Cognitive and Affective Feelings. In *5th International Conference on Affective Computing and Intelligent Interaction (ACII 2013)*.
- Macedo, L. and Cardoso, A. (2001). Creativity and Surprise. *AISB 01 Symposium on Creativity in Arts and Science*, pages 84–92.
- Macedo, L. and Cardoso, A. (2004). Exploration of unknown environments with motivational agents. In *Third International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 328–335, Los Alamitos, CA, USA. IEEE Computer Society.
- Macedo, L., Cardoso, A., and Reisenzein, R. (2006). A surprise-based agent architecture. In *18th European Meeting on Cybernetics and Systems Research*, Vienna.
- Macedo, L., Cardoso, A., Reisenzein, R., Lorini, E., and Castelfranchi, C. (2009). Artificial surprise. In *Handbook of Research on Synthetic Emotions and Sociable Robotics: New Applications in Affective Computing and Artificial Intelligence*, pages 267–291.

-
- Macedo, L., Reizenzein, R., and Cardoso, A. (2004). Modeling forms of surprise in artificial agents: empirical and theoretical study of surprise functions. In *26th Annual Conference of the Cognitive Science Society*, pages 588–593.
- Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3(3):205–258.
- Maia, T. V. and McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: what participants really know in the Iowa gambling task. *Proceedings of the National Academy of Sciences of the United States of America*, 101(45):16075–16080.
- Maia, T. V. and McClelland, J. L. (2005). The somatic marker hypothesis: still many questions but no answers: Response to bechara et al. *Trends in Cognitive Sciences*, 9(4):162–164.
- Makarov, I. and Rytchkov, O. (2012). Forecasting the forecasts of others: Implications for asset pricing. *Journal of Economic Theory*, 147(3):941–966.
- Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1):59–82.
- Malkiel, B. G. (2005). Reflections on the efficient market hypothesis: 30 years later. *The Financial Review*, 40(1).
- Malkiel, B. G. (2016). *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing*. W. W. Norton & Company, 11 edition edition.
- Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, 36(4):394–419.
- Mandelbrot, B. B. (1997). The variation of certain speculative prices. In *Fractals and Scaling in Finance*, pages 371–418. Springer New York.
- Mankiw, N. (2014). *Principles of Macroeconomics*. Cengage Learning.
- Marcellino, M. (2006). Chapter 16 Leading Indicators. In G. Elliott, C. W. J. G. a. A. T., editor, *Handbook of Economic Forecasting*, volume 1, pages 879–960. Elsevier.

- Markose, S., Arifovic, J., and Sunder, S. (2007). Advances in experimental and agent-based modelling: Asset markets, economic networks, computational mechanism design and evolutionary game dynamics. *Journal of Economic Dynamics and Control*, 31(6):1801–1807.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7:77–91.
- Martinho, C. and Paiva, A. (2006). Using Anticipation to Create Believable Behaviour. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1, AAAI'06*, pages 175–180, Boston, Massachusetts. AAAI Press.
- Martino, B. D., Kumaran, D., Seymour, B., and Dolan, R. J. (2006). Frames, Biases, and Rational Decision-Making in the Human Brain. *Science*, 313(5787):684–687.
- Mas-Colell, A., Whinston, M. D., and Green, J. (1995). *Microeconomic Theory*. Oxford University Press, USA, New York, NY.
- Maslov, S. (2000). Simple model of a limit order-driven market. *Physica A: Statistical Mechanics and its Applications*, 278(3–4):571–578.
- Matsumoto, M. and Nishimura, T. (1998). Mersenne Twister: A 623-dimensionally Equidistributed Uniform Pseudo-random Number Generator. *ACM Trans. Model. Comput. Simul.*, 8(1):3–30.
- Mauss, I. . and Robinson, M. . (2009). Measures of emotion: A review. *Cognition and Emotion*, 23:209–237.
- May, R. M., Levin, S. A., and Sugihara, G. (2008). Complex systems: Ecology for bankers. *Nature*, 451(7181):893–895.
- Mayer, J. D., McCormick, L. J., and Strong, S. E. (1995). Mood-Congruent memory and natural mood: New evidence. *Personality and Social Psychology Bulletin*, 21(7):736–746.
- McGaugh, J. L. (2015). Consolidating Memories. *Annual Review of Psychology*, 66(1):1–24.
- Meder, B., Lec, F. L., and Osman, M. (2013). Decision making in uncertain times: what can cognitive and decision sciences say about or learn from economic crises? *Trends in Cognitive Sciences*, 17(6):257–260.

- Mellers, B. A., Schwartz, A., and Cooke, A. D. J. (1998). Judgment and Decision Making. *Annual Review of Psychology*, 49(1):447–477.
- Melvin, M. and Taylor, M. P. (2009). The global financial crisis: Causes, threats and opportunities. Introduction and overview. *Journal of International Money and Finance*, 28(8):1243–1245.
- Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2012). Currency momentum strategies. *Journal of Financial Economics*, 106(3):660–684.
- Merton, R. K. (1948). The Self-Fulfilling Prophecy. *The Antioch Review*, 8(2):193–210.
- Meyer, W. U., Reisenzein, R., and Schutzwohl, A. (1997). Toward a process analysis of emotions: The case of surprise. In *Motivation and Emotion*, volume 21, pages 251–274.
- Miceli, M. and Castelfranchi, C. (2015). *Expectancy and Emotion*. Oxford University Press.
- Miles, C. and Hardman, E. (1998). State-dependent memory produced by aerobic exercise. *Ergonomics*, 41:20–28.
- Miller, G. A. (2003). The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences*, 7(3):141–144.
- Minsky, M. (1974). A framework for representing knowledge. Technical report, Massachusetts Institute of Technology.
- Mitra, K. (2005). Is more data better? *Journal of Economic Behavior & Organization*, 56(2):263–272.
- Mullainathan, S. and Thaler, R. H. (2000). Behavioral economics. In *National Bureau of Economic Research*.
- Murphy, J. J. (1999). *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*. New York Institute of Finance.
- Neely, C. J. and Weller, P. A. (1999). Technical trading rules in the European Monetary System. *Journal of International Money and Finance*, 18(3):429–458.

- Neely, C. J., Weller, P. A., and Ulrich, J. M. (2009). The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market. *Journal of Financial and Quantitative Analysis*, 44(02):467–488.
- Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, 9(1):11–15.
- Nikolai, C. and Madey, G. (2009). Tools of the trade: A survey of various agent based modeling platforms. *Journal of Artificial Societies and Social Simulation*, 12.
- Noordewier, M. K. and Breugelmans, S. M. (2013). On the valence of surprise. *Cognition & Emotion*, 27(7):1326–1334.
- Oaksford, M. and Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences*, 5(8):349–357.
- Oberlechner, T. (2001). Importance of technical and fundamental analysis in the european foreign exchange market. *International Journal of Finance & Economics*, 6(1):81–93.
- Ochsner, K. N. and Gross, J. J. (2005). The cognitive control of emotion. *Trends in Cognitive Sciences*, 9(5):242–249.
- Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*, 53(5):1775–1798.
- Ohanian, L. E. (2010). The Economic Crisis from a Neoclassical Perspective. *Journal of Economic Perspectives*, 24(4):45–66.
- O’Leary, D. (2013). Artificial Intelligence and Big Data. *IEEE Intelligent Systems*, 28(2):96–99.
- Ortony, A., Clore, G. L., and Collins, A. (1998). *The cognitive structure of emotions*. Cambridge University Press.
- Ortony, A. and Partridge, D. (1987). Surprisingness and expectation failure: what’s the difference? In *Proceedings of the 10th international joint conference on Artificial intelligence - Volume 1*, pages 106–108, Milan, Italy. Morgan Kaufmann Publishers Inc.
- Ortony, A. and Turner, T. J. (1990). What’s basic about basic emotions? *Psychological Review*, 97(3):315–331. PMID: 1669960.

- Pagan, A. (1996). The econometrics of financial markets. *Journal of Empirical Finance*, 3(1):15–102.
- Palmer, R. G., Arthur, W. B., Holland, J. H., LeBaron, B., and Tayler, P. (1994). Artificial economic life: a simple model of a stockmarket. *Phys. D*, 75(1-3):264–274.
- Palomino, F. (1996). Noise Trading in Small Markets. *The Journal of Finance*, 51(4):1537–1550.
- Pereira, C. d. C., Mauri, A., and Tettamanzi, A. G. B. (2009). Cognitive-Agent-Based Modeling of a Financial Market. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 02*, pages 20–27. IEEE Computer Society.
- Peters, R. and McGee, R. (1982). Cigarette smoking and state-dependent memory. *Psychopharmacology*, 76(3):232–235.
- Peterson, R. L. (2007). Affect and Financial Decision-Making: How Neuroscience Can Inform Market Participants. *Journal of Behavioral Finance*, 8(2):70–78.
- Phelps, E. A. (2006). Emotion and Cognition: Insights from Studies of the Human Amygdala. *Annual Review of Psychology*, 57(1):27–53.
- Phelps, S. (2012). Applying Dependency Injection to Agent-Based Modelling: the JABM Toolkit. Technical report, Working Papers WP056-12, Centre for Computational Finance and Economic Agents, University of Essex.
- Phelps, S. and Ng, W. L. (2014). A Simulation Analysis of Herding and Unifractal Scaling Behaviour. *Intelligent Systems in Accounting, Finance and Management*, 21(1):39–58.
- Phelps, S. G. (2007). *Evolutionary Mechanism Design*. PhD thesis, University of Liverpool.
- Picard, R. W. (1997). *Affective computing*. MIT Press.
- Picard, R. W. (2003). Affective computing: challenges. *International Journal of Human-Computer Studies*, 59(1-2):55–64.
- Picard, R. W., Wexelblat, A., and Nass, C. I. N. I. C. I. (2002). Future interfaces: social and emotional. In *CHI '02 extended abstracts on Human factors in computing systems*, pages 698–699, Minneapolis, Minnesota, USA. ACM.

- Plutchik, R. (1980). *Emotion: A Psychoevolutionary Synthesis*. Harpercollins College Div.
- Plutchik, R. (1982). A psychoevolutionary theory of emotions. *Social Science Information*, 21(4-5):529–553.
- Prechter, R. R. (2001). Unconscious Herding Behavior as the Psychological Basis of Financial Market Trends and Patterns. *Journal of Psychology and Financial Markets*, 2(3):120–125.
- Quartz, S. R. (2009). Reason, emotion and decision-making: risk and reward computation with feeling. *Trends in Cognitive Sciences*, 13(5):209–215.
- Raafat, R. M., Chater, N., and Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, 13(10):420–428.
- Rabin, M. (1998). Psychology and economics. *Journal of Economic Literature*, 36(1):11–46.
- Rabin, M. (2002). A perspective on psychology and economics. *European Economic Review*, 46:657–685.
- Ranco, G., Aleksovski, D., Caldarelli, G., Grčar, M., and Mozetič, I. (2015). The Effects of Twitter Sentiment on Stock Price Returns. *PLOS ONE*, 10(9):e0138441.
- Ranganath, C. and Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nat Rev Neurosci*, 4(3):193–202.
- Rao, A. and Georgeff, M. (1995). BDI-agents: from theory to practice. In *Proceedings of the First Intl. Conference on Multiagent Systems*.
- Rao, A. S. and Georgeff, M. P. (1998). Modeling rational agents with a BDI-architecture. In *Readings in agents*, pages 317–328. Morgan Kaufmann Publishers Inc.
- Ray, R. (2006). Prediction Markets and the Financial "Wisdom of Crowds". *Journal of Behavioral Finance*, 7(1):2–4.
- Reinhart, C. M. and Rogoff, K. (2009a). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press, 1st edition edition.

-
- Reinhart, C. M. and Rogoff, K. S. (2009b). The Aftermath of Financial Crises. *American Economic Review*, 99(2):466–472.
- Reinhart, C. M. and Rogoff, K. S. (2014). Recovery from Financial Crises: Evidence from 100 Episodes. *American Economic Review*, 104(5):50–55.
- Reisenzein, R. (2007). What is a definition of emotion? and are emotions mental-behavioral processes? *Social Science Information*, 46(3):424–428.
- Reisenzein, R. (2009). Emotions as metarepresentational states of mind: Naturalizing the belief-desire theory of emotion. *Cognitive Systems Research*, 10(1):6–20.
- Reisenzein, R., Hudlicka, E., Dastani, M., Gratch, J., Hindriks, K., Lorini, E., and Meyer, J. (2013). Computational Modeling of Emotion: Towards Improving the Inter- and Intradisciplinary Exchange. *IEEE Transactions on Affective Computing*, Early Access Online.
- Rumsfeld, D. (2002). US Department of Defense News Briefing.
- Russell, J. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39:1178, 1161.
- Russell, J. and Barrett, L. (1999). Core affect, prototypical emotional episodes, and other things called emotion: dissecting the elephant. *Journal of personality and social psychology*, 76(5):819, 805.
- Russell, J., Weiss, A., and Mendelsohn, G. (1989). Affect grid: A single-item scale of pleasure and arousal. *Journal of Personality and Social Psychology*, 57(3):502, 493.
- Russell, S. and Norvig, P. (2009). *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3 edition.
- Sadka, R. (2006). Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk. *Journal of Financial Economics*, 80(2):309–349.
- Samanidou, E., Zschischang, E., Stauffer, D., and Lux, T. (2007). Agent-based models of financial markets. *Reports on Progress in Physics*, 70(3):409.
- Sandholm, T. (2000). eMediator: a next generation electronic commerce server. In *Proceedings of the fourth international conference on Autonomous agents*, pages 341–348, Barcelona, Spain. ACM.

- Sanfey, A. G., Loewenstein, G., McClure, S. M., and Cohen, J. D. (2006). Neuroeconomics: cross-currents in research on decision-making. *Trends in Cognitive Sciences*, 10(3):108–116.
- Schacter, D. L. (2002). *The Seven Sins of Memory: How the Mind Forgets and Remembers*. Mariner Books, 1st edition.
- Schank, R. C. and Abelson, R. P. (1977). *Scripts, Plans, Goals, and Understanding: An Inquiry Into Human Knowledge Structures*. Psychology Press.
- Scherer, K. R. (2005). What are emotions? and how can they be measured? *Social Science Information*, 44(4):695–729.
- Scherer, K. R., Schorr, A., and Johnstone, T. (2001). *Appraisal Processes in Emotion: Theory, Methods, Research*. Oxford University Press, USA, 1 edition.
- Schinckus, C. (2014). Stylized facts and multiple realizability in econophysics. *Results in Physics*, 4:135–136.
- Schmidt, U. and Traub, S. (2002). An experimental test of loss aversion. *Journal of Risk and Uncertainty*, 25(3):233–249.
- Schmidt, U. and Zank, H. (2005). What is loss aversion? *Journal of Risk and Uncertainty*, 30(2):157–167.
- Schutzwohl, A. (1998). Surprise and schema strength. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(5):1182–1199.
- Schutzwohl, A. and Reisenzein, R. (1999). Children’s and adults’ reactions to a schema-discrepant event: A developmental analysis of surprise. *International Journal of Behavioral Development*, 23(1):37–62.
- Scotti, C. (2016). Surprise and Uncertainty Indexes: Real-time Aggregation of Real-Activity Macro Surprises. *Journal of Monetary Economics*.
- Serrano-Padial, R. (2012). Naive traders and mispricing in prediction markets. *Journal of Economic Theory*, 147(5):1882–1912.
- Sharpe, W. F. (1994). The Sharpe Ratio. *The Journal of Portfolio Management*, 21(1):49–58.

- Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, 40(3):777–790.
- Shefrin, H. and Statman, M. (2003). The Contributions of Daniel Kahneman and Amos Tversky. *Journal of Behavioral Finance*, 4(2):54–58.
- Sherry, D. F. and Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, 94(4):439–454.
- Shiller, R. (2002a). The Irrationality of Markets. *Journal of Psychology and Financial Markets*, 3(2):87–93.
- Shiller, R. J. (1987). Investor Behavior in the October 1987 Stock Market Crash: Survey Evidence. Working Paper 2446, National Bureau of Economic Research.
- Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press, Princeton.
- Shiller, R. J. (2002b). Bubbles, Human Judgment, and Expert Opinion. *Financial Analysts Journal*, 58(3):18–26.
- Shiller, R. J. (2003). From efficient market theory to behavioral finance. *Journal of Economic Perspectives*, 17(1):83–104.
- Shiller, R. J. (2013). Is Economics a Science?
- Shleifer, A. and Summers, L. H. (1990). The Noise Trader Approach to Finance. *Journal of Economic Perspectives*, 4(2):19–33.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.
- Skidelsky, R. J. A. (2010). *Keynes: The Return of the Master*. Penguin Books.
- Skinner, D. and Sloan, R. (2002). Earnings surprises, growth expectations, and stock returns or don't let an earnings torpedo sink your portfolio. *Review of Accounting Studies*, 7(2):289–312.
- Smith, V. L. (1991). Rational choice: The contrast between economics and psychology. *The Journal of Political Economy*, 99(4):877–897.
- Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica*, 56(5):1119–1151.

- Snowberg, E., Wolfers, J., and Zitzewitz, E. (2012). Prediction Markets for Economic Forecasting. Working Paper 18222, National Bureau of Economic Research.
- Soffer, L. C., Thiagarajan, S. R., and Walther, B. R. (2000). Earnings preannouncement strategies. *Review of Accounting Studies*, 5(1):5–26.
- Sommer, J. (2013). Robert shiller: A skeptic and a nobel winner. *The New York Times*.
- Sorkin, A. R. (2010). *Too Big to Fail: Inside the Battle to Save Wall Street*. Penguin, London.
- Sornette, D. (2014). Physics and financial economics (1776–2014): puzzles, Ising and agent-based models. *Reports on Progress in Physics*, 77(6):062001.
- Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, 82(3):171–177.
- Squire, L. R. (2009). Memory and brain systems: 1969-2009. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 29(41):12711–12716. PMID: 19828780.
- Sterman, J. D. (1994). Learning in and about complex systems. *System Dynamics Review*, 10(2-3):291–330.
- Stock, J. H. and Watson, M. W. (2003). Forecasting Output and Inflation: The Role of Asset Prices. *Journal of Economic Literature*, 41(3):788–829.
- Struzik, Z. (2003). Taking the pulse of the economy. *Quantitative Finance*, 3(4):C78–C82.
- Studer, B., Limbrick-Oldfield, E. H., and Clark, L. (2015). ‘Put Your Money Where Your Mouth Is!’: Effects of Streaks on Confidence and Betting in a Binary Choice Task. *Journal of Behavioral Decision Making*, 28(3):239–249.
- Summerfield, C. and Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19(1):27–34.
- Summers, L. H. (1986). Does the stock market rationally reflect fundamental values? *Journal of Finance*, 41(3):591–601.

- Takahashi, H. and Terano, T. (2003). Agent-Based approach to investors' behavior and asset price fluctuation in financial markets. *Journal of Artificial Societies and Social Simulation*, 6.
- Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. Random House, 1 edition.
- Taleb, N. N. (2008). *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*. Random House, 2 updated edition.
- Taleb, N. N. (2014). *Antifragile: Things That Gain from Disorder*. Random House Trade Paperbacks, New York, reprint edition.
- Teigen, K. H. and Keren, G. (2003). Surprises: low probabilities or high contrasts? *Cognition*, 87(2):55–71.
- Tesfatsion, L. (2002). Agent-Based computational economics: Growing economies from the bottom up. *Artificial Life*, 8:55–82.
- Thaler, R. H. (2016). Behavioral Economics: Past, Present, and Future. *American Economic Review*, 106(7):1577–1600.
- Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press, New York, NY.
- The Nobel Prize Foundation (2002). The Prize in Economic Sciences 2002 - Press Release.
- The Nobel Prize Foundation (2013). The Prize in Economic Sciences 2013 - Press Release.
- Thomsen, D. K. (2009). There is more to life stories than memories. *Memory*, 17(4):445.
- Thomson, D. M. and Tulving, E. (1970). Associative encoding and retrieval: Weak and strong cues. *Journal of Experimental Psychology*, 86(2):255–262.
- Timmermann, A. and Granger, C. W. J. (2004). Efficient market hypothesis and forecasting. *International Journal of Forecasting*, 20(1):15–27.
- Tom, S. M., Fox, C. R., Trepel, C., and Poldrack, R. A. (2007). The Neural Basis of Loss Aversion in Decision-Making Under Risk. *Science*, 315(5811):515–518.

- Treleaven, P., Galas, M., and Lalchand, V. (2013). Algorithmic trading review. *Commun. ACM*, 56(11):76–85.
- Trueman, B. (1994). Analyst Forecasts and Herding Behavior. *Review of Financial Studies*, 7(1):97–124.
- Tsang, E. (2009). Forecasting - where computational intelligence meets the stock market. *Frontiers of Computer Science in China*, 3(1):53–63.
- Tulving, E. (1972). Episodic and semantic memory. In *Organization of Memory*, pages 382–402. Academic Press, Inc, New York, NY.
- Tulving, E. (1983). *Elements of Episodic Memory*. Oxford University Press, USA.
- Tulving, E. (2002). Episodic memory: from mind to brain. *Annual Review of Psychology*, 53:1–25. PMID: 11752477.
- Tulving, E. and Osler, S. (1968). Effectiveness of retrieval cues in memory for words. *Journal of experimental psychology*, 77(4):601, 593.
- Tulving, E. and Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 80(5):352–373.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.
- Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458.
- Tversky, A. and Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*.
- U.S. Commodity Futures Trading Commission and U.S. Securities & Exchange Commission (2010). Findings regarding the market events of May 6, 2010. Technical report.
- Varian, H. R. (2014). Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*, 28(2):3–28.
- Vega, C. (2006). Stock price reaction to public and private information. *Journal of Financial Economics*, 82(1):103–133.

- Verrecchia, R. E. (1982). Information Acquisition in a Noisy Rational Expectations Economy. *Econometrica*, 50(6):1415–1430.
- Wan, H. A., Hunter, A., and Dunne, P. (2002). Autonomous agent models of stock markets. *Artificial Intelligence Review*, 17(2):87–128.
- Warneryd, K. (2001). *Stock-Market Psychology: How People Value and Trade Stocks*. Edward Elgar Publishing, illustrated edition edition.
- Wasserman, E. A. and Castro, L. (2005). Surprise and change: Variations in the strength of present and absent cues in causal learning. *Learning & Behavior*, 33(2):131–146.
- Weiss, G. (1999). *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. The MIT Press.
- Werner, N. S., Duschek, S., and Schandry, R. (2009). Relationships between affective states and decision-making. *International Journal of Psychophysiology*, 74(3):259–265.
- Whalen, P. J. (2007). The uncertainty of it all. *Trends in Cognitive Sciences*, 11(12):499–500.
- Whaley, R. E. (1993). Derivatives on Market Volatility. *The Journal of Derivatives*, 1(1):71–84.
- Whaley, R. E. (2000). The Investor Fear Gauge. *The Journal of Portfolio Management*, 26(3):12–17.
- Whaley, R. E. (2009). Understanding the VIX. *The Journal of Portfolio Management*, 35(3):98–105.
- White, A. M. (2003). What happened? alcohol, memory blackouts, and the brain. *Alcohol Research & Health: The Journal of the National Institute on Alcohol Abuse and Alcoholism*, 27(2):186–196. PMID: 15303630.
- Windrum, P., Fagiolo, G., and Moneta, A. (2007). Empirical Validation of Agent-Based Models: Alternatives and Prospects. *Journal of Artificial Societies and Social Simulation*, 10(2).
- Wixted, J. T. (2004). The Psychology and Neuroscience of Forgetting. *Annual Review of Psychology*, 55(1):235–269.

- Wolford, G., Newman, S. E., Miller, M. B., and Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology = Revue Canadienne De Psychologie Expérimentale*, 58(4):221–228.
- Wooldridge, M. (2002). *An Introduction to MultiAgent Systems*. John Wiley & Sons, 1st edition.
- Wooldridge, M. and Jennings, N. R. (1995). Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10(02):115–152.
- Wurgler, J. (2000). Financial markets and the allocation of capital. *Journal of Financial Economics*, 58(1-2):187–214.
- Wurman, P. R., Wellman, M. P., and Walsh, W. E. (1998). The michigan internet AuctionBot: a configurable auction server for human and software agents. In *Proceedings of the second international conference on Autonomous agents*, pages 301–308, Minneapolis, Minnesota, United States. ACM.
- Yang, C., Hu, S., and Xia, B. (2012). The endogenous dynamics of financial markets: Interaction and information dissemination. *Physica A: Statistical Mechanics and its Applications*, 391(12):3513–3525.
- Yeung, C. H. and Zhang, Y. C. (2008). Minority games.
- Yonelinas, A. P. and Ritchey, M. (2015). The slow forgetting of emotional episodic memories: an emotional binding account. *Trends in Cognitive Sciences*, 19(5):259–267.
- Zhang, R., Brennan, T. J., and Lo, A. W. (2014). The origin of risk aversion. *Proceedings of the National Academy of Sciences*, 111(50):17777–17782.
- Zhang, W. and Semmler, W. (2009). Prospect theory for stock markets: Empirical evidence with time-series data. *Journal of Economic Behavior & Organization*, 72(3):835–849.