

UNIVERSITE DE NICE-SOPHIA ANTIPOLIS - UFR SCIENCES  
Ecole Doctorale de Sciences Fondamentales et Appliquées

## T H E S E

pour obtenir le titre de  
**Docteur en Sciences**  
de l'UNIVERSITE de Nice-Sophia Antipolis  
Spécialité : "Physique Pluridisciplinaire"

présentée et soutenue par le  
*Jeffrey SATINOVER*

TITRE :  
**ILLUSORY AND GENUINE CONTROL IN OPTIMIZING  
GAMES AND FINANCIAL MARKETS**

Thèse dirigée par : *Didier SORNETTE*  
soutenue le 24 Juillet, 2008

### JURY:

Professeur	Pierre BERNHARD	Examineur
Professeur	Damien CHALLET	Rapporteur
Professeur	Taisei KAIZOJI	Rapporteur
Professeur	Victor PLANAS	Examineur
Professeur	Didier SORNETTE	Directeur de thèse



## Résumé

Les gens aiment croire qu'ils sont maîtres de leur destin. Mais on comprend mal les circonstances dans lesquelles nous pouvons surestimer ou perdre notre capacité à optimiser les résultats, surtout lorsqu'ils sont le résultat de l'agrégation de différents produits d'optimisation. En utilisant des simulations numériques et des chaînes de Markov, nous montrons que dans le jeu de minorité (MG) à horizon de temps fini (THMG) et dans les jeux de Parrondo (PG), des agents qui optimisent sur la base d'informations précédentes obtiennent des performances inférieures à celles d'agents aléatoires. Le résultat inverse est obtenu pour les jeux de la majorité et le "dollar game". De plus, les agents qui déploient leur stratégie la moins bonne dans le passé surperforment tous les autres agents dans les jeux de la minorité, et sous-performent dans les jeux de la majorité et le dollar-game. Ainsi, les stratégies à faible entropie sous-performent dans certaines circonstances, ce qui est paradoxal compte-tenu de l'utilisation de l'information existante. Nos résultats permettent de préciser les conditions sous lesquelles le phénomène d'illusion du contrôle peut apparaître, et permet de quantifier l'amplitude de l'effet, dans des contextes a priori définis pour que l'optimisation apporte un plus.

Pour aider à comprendre cette dynamique, nous démontrons que des séries chronologiques dans le TH- MG, -MAJG et -\$G possèdent des degrés caractéristiques de persistance – une mesure de la tendance à répéter des motifs.

Ces séries peuvent être décomposées en des superpositions spécifiques de cycles déterministes sur des graphes. Ces «taxonomies» conduisent naturellement à des questions connexes: Les séries financières peuvent-elles être caractérisées par leur cycle de décomposition? Est-il possible de générer des prédicteurs utilisant ces cycles de décomposition? Nous proposons des réponses en fournissant des exemples tirés des marchés financiers, où les actifs sont des actions, métaux précieux et taux de change.

## Abstract

People like to believe they are in control of their destiny. This ubiquitous trait seems to increase motivation and persistence, and is probably evolutionarily adaptive[1, 2]. But there is little understanding of when and under what circumstances we may over-estimate[3] or even lose our ability to optimize outcomes, especially when they are the result of aggregations of individual optimization processes.

Here, we demonstrate analytically using the theory of Markov chains and by numerical simulations in four classes of games, the Time-Horizon Minority Game (THMG) [4, 5], the Time-Horizon Majority Game (THMAJG), the Time-Horizon Dollar Game (TH\$G) and the Parrondo Games [6, 7], that agents who optimize their strategy based on past information sometimes perform worse than non-optimizing agents, and other times better. In other words, low-entropy (more informative) strategies may under certain circumstances under-perform high-entropy (or random) strategies. Comparing the different results in different games sharpens the definition of the “illusion of control”—and the distinction between genuine and illusory control—in set-ups a priori defined to emphasize the importance of optimization.

We demonstrate as well that if, in the Time-Horizon Minority Game, a subset of agents deploy what appears to be their worst-performing strategies instead of their best as do standard agents, these not only outperform the standard agents, they generate returns for themselves that can be absolutely positive. On the other hand, in the THMAJG and TH\$G this is not true—agents that choose their worst-performing strategies underperform those that choose their best. As part of the attempt to understand the peculiarities of these dynamics, we demonstrate that the THMG, -MAJG and -\$G may all be decomposed into a (finite-set) superposition of wholly deterministic cycles on graphs.

The analysis of Time-Horizon games in these terms leads naturally to related but distinct questions: Is it possible to characterize real-world financial series by their cycle decompositions? And is it possible to generate working predictors by decomposing real-life time-series into a cycle decomposition? We answer by providing toy examples from the domain of equities (the NASDAQ Composite Index), of precious metals markets (the Philadelphia Exchange Gold and Silver Index) and of high-frequency foreign exchange (the former USD/DM exchange rate).

## **Les mots cléf**

applications à la théorie des jeux mathématiques et économie, automates cellulaires, de nombreux corps théorie, econophysics, foule dynamique, graphes aléatoires, modèles des agent interactive, modèles des marchés financiers, non-linéaire dynamique, phénomènes critiques de systèmes socio-économiques, physique et de la société, processus stochastiques, réseaux de neurones, systèmes adaptatifs, systèmes complexes, systèmes désordonnés (théorie), théorie des jeux

## **Key words**

adaptive systems, applications to game theory and mathematical economics, cellular automata, complex systems, critical phenomena of socio-economic systems, crowd dynamics, disordered systems (theory), econophysics, game theory, interacting agent models, many body theory, models of financial markets, neuronal networks, nonlinear dynamics, physics and society, random graphs, stochastic processes

## **Institution**

Prepared at the Laboratoire de Physique de la Matière Condensée, CNRS U. M. R. 6622,  
Université de Nice - Sophia Antipolis, Parc Valrose, 06108 Nice Cedex 2, France

# Contents

Résumé.....	3
Abstract.....	4
Les mots cléf.....	5
Key words.....	5
Institution.....	6
Contents.....	7
Acknowledgements.....	10
Dedication.....	11
Chapter 1. Introduction.....	13
1.1 Plan of the Dissertation.....	14
1.2 Associated publications.....	14
Chapter 2. The “illusion of control” in historical context: from society and the market economy to the physics of complex adaptive systems.....	16
2.1 Spontaneous Order in Social and Economic Systems.....	16
2.2 The “law of unintended consequences”.....	19
2.3 The physics of complex systems.....	21
Chapter 3. The “illusion of control” in agent-based and Parrondo Games.....	26
3.1 “Illusion of control” in the Time-Horizon Minority Game (THMG).....	26
3.1.1 Definition and summary of main results for the Time-Horizon MG (THMG).....	26
3.1.2 Statement of our main results on the “illusion of control” in the THMG.....	28
3.1.3 Quantitative statement and tests.....	29
3.1.4 Generalizations.....	30
3.1.5 Illusion of control and the crowding-out mechanism.....	34
3.1.6 Robustness of the “illusion of control” phenomenon: THMG versus MG.....	35
3.1.7 First-entry games and symmetric evolutionary stable equilibria.....	37
3.2 Parrondo Games.....	38
3.2.1 Single-player capital-dependent Parrondo effect.....	38
3.2.2 Single-player history-dependent Parrondo effect.....	39
3.2.3 Multiple-player capital-dependent Parrondo effect and its reversal under optimization.....	42
3.2.4 Single-player capital-dependent Parrondo effect and its reversal under optimization.....	42
3.3 Conclusions.....	44
Chapter 4. Control in the Majority Game and Dollar Game.....	45
4.1 Optimization in the MAJG and \$G: No illusion.....	45
4.2 Minority versus Majority and \$ Games.....	45
4.2.1 Recap of the “illusion of control” in MG and Time-Horizon MG (THMG).....	45
4.2.2 Definition and overview of the Majority Game (MAJG).....	46
4.2.3 Definition and overview of the Dollar Game (\$G).....	47
4.3. Main results on the “illusion of control” in the THMG v. MAJG and \$G.....	47
4.3.1 Quantitative statement and tests.....	48
4.3.1.1 Analytic Calculation versus Numeric Simulation.....	48
4.3.1.2 Illusory versus Genuine Control for $\tau \ll \tau_{eq}$ .....	49
4.3.1.3 Illusory versus Genuine Control for $\tau < \tau_{eq}$ .....	51
4.4. Interpretations: crowding-out, anti-optimizing agents and persistence.....	53
4.4.2. Illusion of control and “anti-optimizing” agents.....	53

4.5 Conclusions .....	54
Chapter 5. Persistence and anti-persistence in the MG, MAJG and \$G .....	56
5.1 Persistence versus Anti-persistence in the THMG, THMAJG and TH\$G.....	57
5.2 Cycle structure of a binary series.....	60
5.3 Cycle distinctions among MG, AJG and \$G.....	60
5.4 Cycle analysis of the illusion of control .....	65
5.5 Cycle-based predictors for the MG, MAJG and \$G .....	66
5.6 Approximation of a Series as a Perturbed Dominant Cycle.....	67
Chapter 6. Cycle- and persistence-based real-world (toy) predictors .....	69
6.1 Cycle-based predictor for the NASDAQ composite.....	69
6.2 Cycle predictors for other time-series.....	71
Chapter 7. Supplementary findings in the MG .....	76
7.1 Choosing the worst strategy in the THMG/MG.....	76
7.1.1 Sample Results from the MG and THMG .....	77
7.1.2 Microscopic examination of agent performance with and without C agents in the THMG.....	80
7.2 Illusion of control in the THMG at various $\tau$ .....	88
7.3 Performance of optimizing (S) versus random-choice agents in the THMG.....	90
Chapter 8. Discussion.....	91
8.1 Status of the illusion of control; anti-persistence.....	91
8.2 Cycle decomposition .....	91
8.3 Limitations and future lines of research.....	92
8.3.1 Effects of $S > 2$ in MG, MAJG and \$G.....	92
8.3.2 Grand Canonical MG, MAJG and \$G.....	92
8.3.3 Agents with mixed memory and/or lookback .....	93
8.3.4 Other Games .....	93
Chapter A9. Appendix: mathematical methods.....	94
A9.1: Analytic Methods for the TH –MG, -MAJG and -\$G .....	94
A9.1.1 The Time Horizon Minority Game: Choosing the Best Strategy .....	94
A9.1.2 The Time Horizon Minority Game: Choosing the Worst Strategy .....	100
A9.1.3 Analytic methods for the THMAJG and TH\$G .....	106
A9.1.4 Analytic methods for the cycle decomposition of binary series.....	109
A9.1.4.1.Cycle Decomposition of TH games.....	109
A9.1.4.2.Extraction of Cycles from a Binary Series .....	110
A9.1.4.3 Analytic Form of a Cycle Decomposition.....	113
A9.1.4.4 Comparison to Numerical Simulation .....	115
A9.2 Analytic expression of the general Parrondo effect.....	116
Chapter A10. Appendix: anti-persistence using an equity-ranking predictor.....	118
A10.1 Introduction.....	118
A10.2 Background.....	118
A10.3 Review of the VL ranking system and its performance .....	121
A10.4 Methods .....	124
A10.4.1 Equities .....	124
A10.4.2 Input data and outputs .....	124
A10.4.3 Selection of trading period start .....	125
A10.4.4 Network Architecture.....	125
A10.4.5 Training.....	125
5 Results .....	126
A10.5.1 Hedged Returns.....	126
A10.5.1.1 Overall results.....	126
A10.5.1.2 Long Returns .....	127

A10.5.1.3 Short Returns .....	128
A10.5.2 Risk-adjusted returns .....	129
A10.5.2.1 Sharpe ratios .....	129
A10.5.2.2 Jensen's alpha .....	130
A10.6 Antipersistence in predictor behavior .....	132
A10.6.1 Distribution of ANN predictions as a composite of success and failure .....	133
A10.6.2 Anti-persistence in predictor performance .....	137
A10.6.3 Effects of overall predictor performance on the top and bottom ends of the predictor rankings .....	137
A10.6.4 T/B Portfolios aggregated by global success or failure of the predictor(s) .....	140
A10.6.4.1 T/B 10 = H10 Portfolios .....	140
A10.6.4.2 T/B 100 = H100 Portfolios .....	149
A10.7 Discussion .....	150
A10.7.1 Data Cleaning and Errors .....	150
A10.7.2 Why Does the ANN Predictor Work? .....	150
A10.7.3 Relation between the ANN and the VL ranking system .....	151
References .....	152

## Acknowledgements

Without the inspiration, encouragement and assistance of Prof. Didier Sornette, this thesis would have been impossible. I am deeply grateful that he agreed to be my thesis director and mentor during these years in association with the University of Nice. I thank most sincerely the University, too, and the Condensed Matter Physics Laboratory/CNRS for providing me an intellectual home over the past three years.

R. Shankar, Professor and Chairman of Physics at Yale University, was a great source of encouragement, support and intellectual stimulus in the earlier years of my physics study. I am grateful to him (and to his family) for their friendship.

Jørgen Andersen at the Condensed Matter Physics Laboratory of the University of Nice/CNRS made some extremely helpful observations that initiated the direction of this thesis. I have benefited as well from observations made by Didier Sornette's colleagues at the ETH, Zurich. I am grateful to Riley Crane, Gilles Daniel and especially Yannick Mallevergne for helpful feedback on important sections of the manuscript.

Damien Challet initially posed the question as to whether an “illusion of control” discovered for Minority Games would be found in Majority Games as well.

Over the years, Anne Sornette has done more to help smooth my way for which I can not possibly express adequate thanks. She is truly an angel.

The research on artificial neural networks in Chapter (Appendix) 10 was carried out while the author was a graduate student in the Department of Physics at Yale University, as part of the Condensed Matter Theory Group, while concurrently serving as Director of Research and Quantitative Finance for Von Kohorn Research and Advisory, Westport Connecticut (VKRA). The initial phases of this research was conducted primarily using the facilities of VKRA. The author is especially grateful for the dedicated programming and data management assistance at that stage of Sarah Satinover.

Finally, I must thank my wife Julie and daughters Sarah, Anni and Jenny—for putting up with me during the years of intense study, research and writing that have led to this thesis. Words truly cannot adequately express how deeply I am indebted to their support and forbearance.

## **Dedication**

To RPF, and to my father, in memoriam

“Every step and every movement of the multitude, even in what are termed enlightened ages, are made with equal blindness to the future; and nations stumble upon establishments, which are indeed the result of human action, but not the execution of any human design.”

— Adam Ferguson, *An Essay on the History of Civil Society* (1767): Part Third. Section II, p. 122 of the Duncan Forbes edition, Edinburgh University Press, 1966.

## Chapter 1. Introduction

Human beings like to believe they are in control of their destiny. This ubiquitous trait seems to increase motivation and persistence, and is probably evolutionarily adaptive [1, 2]. But how good really is our ability to control? How successful is our track record in these areas? There is little understanding of when and under what circumstances we successfully optimize and control outcomes, and when we may over-estimate [3] or even lose that ability, especially when outcomes are the result of the aggregations of many individual optimization processes, i.e. within a complex adaptive system (CAS).

Using the theory of Markov chains and by numerical simulations in four classes of games, the Time-Horizon Minority Game (THMG) [4, 5], the Time-Horizon Majority Game (THMAJG), the Time-Horizon Dollar Game (TH\$G) and the Parrondo Games [6, 7], we study analytically how agents who optimize their strategy based on past information sometimes perform worse than non-optimizing agents, and other times better. In other words, low-entropy (more informative) strategies may under certain circumstances under-perform high-entropy (or random) strategies. Comparing the different results in different games sharpens the definition of the “illusion of control”—and the distinction between genuine and illusory control—in set-ups a priori defined to emphasize the importance of optimization.

We demonstrate as well that if, in the Time-Horizon Minority Game, a subset of agents deploy what appear to be their worst-performing strategies instead of their best, as do standard agents, these not only outperform the standard agents, they generate returns for themselves that can be absolutely positive. On the other hand, in the THMAJG and TH\$G this is not true—agents that choose their worst-performing strategies underperform those that choose their best. As part of the attempt to understand the peculiarities of these dynamics, we demonstrate that the THMG, -MAJG and -\$G may all be decomposed into a (finite-set) superposition of wholly deterministic cycles on graphs.

It is further the case that the time-series generated by the THMG are, for a range of parameters, *anti-persistent*; the time-series generated by the THMAJG and TH\$G are *persistent*, where “persistence” is a measure of the tendency of a binary sequence to contain repeating patterns.

The analysis of Time-Horizon games in these terms leads naturally to related but distinct questions: Is it possible to characterize real-world financial series by their persistence and/or cycle decompositions? And is it possible to generate working predictors by utilizing persistence and decomposing real-life time-series into a cycle decomposition? We answer by providing toy examples from the domain of equities (the NASDAQ Composite Index), of precious metals markets (the Philadelphia Exchange Gold and Silver Index) and of high-frequency foreign exchange (the former USD/DM exchange rate).

The work presented here explores a very particular set of games that for the most part exemplify “complex adaptive systems” (CAS; not so Parrondo Games). And in particular, we draw attention to the fact that the MG and THMG are characterized by anti-persistence in the generated time-series, whereas series in both the MAJG/THMAJG and \$G/TH\$G are characterized by persistence. Specific as the examples may be, the exten-

sion to financial markets suggests the more general importance of these games and their peculiarities in understanding complex systems both theoretically and in the real world.

## **1.1 Plan of the Dissertation**

In Chapter 2. we briefly discuss the question of illusory versus genuine control in complex systems in its historical context. We take from this discussion the motivation for and importance of attempting to understand quantitatively and where possible analytically, the circumstances under which control may be genuine and when it is likely to be illusory.

In Chapter 3. we survey and summarize descriptively the main results of the study in the market-modeling MG and in the PG. We discuss the important differences, both theoretical and applied, between the standard MG and its “time-horizon” variant, the THMG.

In Chapter 4. we survey and summarize descriptively the main results of the study in the two other classes of market-modeling games the MAJG and the \$G. We discuss the important differences, both theoretical and applied, between the standard MAJG and \$G their “time-horizon” variants. We discuss the different results with respect to the illusion of control in the MG/THMG versus the MAJG/THMAJG//\$G/TH\$G.

In Chapter 5. we present the differences among the main time-series generated by these games in terms of persistence/anti-persistence and in terms of their decomposition into cycles. The cycle decomposition method is used to analyze the illusion of control in the THMG and to illustrate how these games may be understood as perturbed deterministic processes.

In Chapter 6. we present the results of generating real-world predictors based on the cycle decomposition method as well as persistence/anti-persistence.

In Chapter 7. we illustrate certain additional findings and subtleties in these games corollary to the main findings already presented.

In Chapter 8. we discuss limitations in the present research and suggest further avenues of research opened up by the present findings.

Chapter 9. is a mathematical methods appendix. We present formal analytic expressions for all four games. We also present methods for determining the persistence/anti-persistence of a binary series, for decomposing a given binary series into a superposition of cycles on graphs and for obtaining analytically the cycle-decomposition of a known Markov transition matrix.

Chapter 10. is an appendix that describes an alternate prediction method based on neural networks. We discuss the results in terms of what they reveal about changes in regime (i.e., phase transitions) in real world complex systems. The method illustrates a certain degree of predictive control along with many instances of the “illusion of control”.

## **1.2 Associated publications**

The contents of five published, submitted and working papers are incorporated in this thesis. They are:

J.B. Satinover and D. Sornette,  
Illusion of Control in a Brownian Game,  
Physica A 386, 339-344 (2007)  
(<http://arxiv.org/abs/physics/0703048>)

J.B. Satinover and D. Sornette,  
“Illusion of Control” in Minority and Parrondo Games,  
Eur. Phys. J. B 60, 369-384 (2007)  
(<http://arxiv.org/abs/0704.1120>)

J.B. Satinover and D. Sornette,  
Illusory versus Genuine Control in Agent-Based Games,  
submitted to Eur. Phys. J. B. (2008)  
(<http://arXiv.org/abs/0802.4165>)

J.B. Satinover and D. Sornette,  
Cycles, determinism and persistence in agent-based games and financial time-series,  
submitted to Quantitative Finance (2008)  
(<http://arXiv.org/abs/0805.0428>)

J.B. Satinover and D. Sornette  
Anti-persistence in an Equity-Ranking Predictor,  
Swiss Finance Institute Research Paper No. 08-15 (2008)  
<http://ssrn.com/abstracts=1155427>

## Chapter 2. The “illusion of control” in historical context: from society and the market economy to the physics of complex adaptive systems

### 2.1 Spontaneous Order in Social and Economic Systems

In a famous passage from Chapter 18 of *Economic Sophisms*, Friedrich Bastiat wonders why no one ever fears there will be no bread the next morning in Paris[8]:

Here are a million human beings who would all die in a few days if supplies of all sorts did not flow into this great metropolis. . . . What, then, is the resourceful and secret power that governs the amazing regularity of such complicated movements...? [T]he principle of free exchange.

What exactly is meant by “free exchange?” From a physicist’s perspective the answer is this: Quantifiable interactions (“exchanges”) that are relatively free (or “free-enough”) from globally-imposed (“top down”) rules.

This is the essential claim of the Austrian School of Economics whose dominant voice was Friedrich Hayek: “For Hayek, market institutions are epistemic devices—means whereby information that is scattered about society and known in its totality by no one can be used by all by being embodied in prices”[9]. “Society and the Market are spontaneous orders—results of human action but not of human design” [10].

The essence of society and the market economy is *cooperativity*, which we distinguish from *cooperation*. Cooperativity amongst agents in an economy or citizens in a society arises tacitly without the deliberate seeking and coordination of action that constitutes cooperation proper. This unseen cooperativity is given literary vivid expression in Leonard Read's “I, Pencil”[11]:

[M]illions of human beings have had a hand in my creation.... There isn't a single person in all these millions, including the president of the pencil company, who contributes no more than a tiny, infinitesimal bit of know-how.... Neither... the miner of graphite in Ceylon [nor] the logger in Oregon can be dispensed with, any more than can the chemist at the factory or the worker in the oil field... There is a fact still more astounding: the absence of a master mind, of anyone dictating or forcibly directing these countless actions which bring me into being.

Hayek explained this process from an socioeconomic perspective in a 1945 article in the *American Economic Review*, “The Use of Knowledge in Society”[10], at roughly the same time that John Von Neumann was beginning to develop the earliest notions of self-organizing computational systems[12]. These, von Neumann argued, are the foundational model for biology and evolution: so called “Cellular Automata” governed exclusively by nearest-neighbor interactions, yet which demonstrated development, global self-organization and self-propagation. Hayek wrote [10]:

[I]n a system in which the knowledge of the relevant facts is dispersed among many people, prices can act to coordinate the separate actions of different people...The whole acts as one market... [and] brings about the solution which ... might have been arrived at by one single mind possessing all the information which is in fact dispersed among all the people involved in the process.

The mathematical formalization of self-organizing cooperativity began with the founding of what is now known as computational neuroscience in particular and the more general study of so-called “complex adaptive systems” (CAS).

Although “complex systems” (CS) and “complex adaptive systems” have become major areas of study in the physical, biological and social sciences, no fixed definition of these has ever been agreed upon: The essence of what makes a system “complex” (not to be confused with merely “complicated”) has proven notoriously elusive. In his famous concurrence to the court’s opinion in the obscenity case of *Jacobellis v. Ohio* (1964), United States Supreme Court Justice Potter Stewart wrote of pornography that it was “hard to define,” but nonetheless, “I know it when I see it.” So, too, with CS/CAS.

Nonetheless, certain characteristic features of CS/CAS can be identified with fairly widespread agreement. A system may be considered *complex*[13] and is therefore a CS if:

- It is composed of many interacting units (“agents”) such that...
- ...it demonstrates “emergent” traits, i.e., traits attributable to the system-as-a-whole that are not characteristic of the constituent units and that do not arise from the simple statistical averaging over agents’ quantified traits.


A CS is *adaptive* and is therefore a CAS if further:

- Its agents they respond differentially in systematic fashion to changes in some condition of the whole (global state) or part of the whole (i.e., local state which need not necessarily be a neighborhood but often is).
- The differential response of its agents is *goal-directed*, i.e., there is either an implicit or explicit fitness function that dictates agent response. The fitness function may be defined locally, globally or both.

By these definitions, CS/CAS may be demonstrated in naturally-arising physical systems, in living systems, in social systems and in abstract (usually computer-based) models [14]. A living terrestrial nervous system composed of neurons is a CAS as is likewise an artificial neural network whether its processing elements are constructed out of hardware units or are software simulations.

An example of a computer-generated CAS that arises naturally in both an inorganic physical system and in a biological system is the so-called “compete-cooperate” cellular automaton (CA). **Figure 1** shows all three: On the left the CA, in the middle the Belousov-Zhabotinsky chemical reaction in a thin film of liquid (i.e., between two plates), on the right a colony on agar of the amoeba species *dictyostelium discoideum* [15]. All three develop the same time-varying but stable pattern: spirals of contrasting value (color) that rotate and propagate outward.

In these instances, how an “agent” (abstract, molecular, single cell organism) responds is unrelated to any global measure of the system as a whole. Instead, it responds to the state of its nearest neighbors, adopting a like value when the sum of neighborhood values is below a certain threshold and a different value when it is above. Thus it incorporates both “cooperation” (i.e., imitation) and “competition” (anti-imitation).



**Figure 1:** Spiral wave formation in from left to right: Artificial (software) compete-cooperate cellular automaton, thin-film implementation of the Belousov-Zhabotinsky oscillating chemical reaction and colony formation on agar of the amoeba *dictyostelium discoideum*. (With permission [15])

In general computational science, the term “Artificial Intelligence” has come exclusively to refer to “expert systems” in which a set of predetermined algorithms are preprogrammed into a machine. No “adaptive learning” from experience takes place. Interestingly, this definition has become embedded largely under the influence of Marvin Minsky at M.I.T., who nonetheless began his career seeking to exploit early discoveries of the opposite sort—that is, of what was then called “Perceptrons,” simple versions of artificial “neural networks”. Neural networks are machine learning devices modeled on biological nervous system tissue that employ trial and error learning to converge bottom-up toward solutions on their own. Like society and the market economy, and like John Von Neumann’s “cellular automata”, they are networks of interacting elements that spontaneously self-organize. Von Neumann—the “father of the modern computer”—considered this biologically-based architecture, *not* the top-down architecture with which we are so familiar, to be the vastly superior kind, as it was the one employed to such self-evident astounding versatility by nature.

Minsky, however, gave up on neural networks in the 1960’s. But in the early 1980’s, a Cal Tech physicist named John Hopfield discovered that the mathematics of certain kind of combined magnetic and anti-magnetic substances found in nature—so-called “spin-glasses” exactly mimicked that of a distinct kind of neural network remarkably similar to certain biological neural networks. Both provided a model for the spontaneous formation and storage of *memories*, the foundational capacity of intelligence, alternately conceptualized as the global organization of the network as a whole. With his 1982 presentation on the subject to National Academy of Sciences[16], computational neuroscience was born as a subject of immense challenge and interest. Neural networks quickly became a field of massive academic research (initially dominated by physicists with a keen interest in spin-glass phenomena) and eventually industrial investment[17].

The “Hopfield net” is a certain kind of fully connected neural network mathematically identical to a network of spins in a state of matter (“spin glasses”) characterized by competition, “frustration” and multiple equilibria. Both the artificial Hopfield net and the naturally-occurring spin glass are capable of evolving spontaneous order characterized by the spontaneous (“bottom-up”) formation of associative memories (that is “content addressable memories”) of the kind the human brain forms during the process of learning—distinct from the deliberate (“top-down”) embedding by a “teacher” of addressable memory storage in standard computers.

Networks of artificial processing elements, either hardware or simulated in software in artificial neural networks, networks of biological neurons, networks of interacting spins

in physical “spin-glasses” and networks of individuals forming a market economy all process information and store memories in related ways. (Related, but not necessarily identical as in the example shown in **Figure 1**) For example, the neurons in brain tissue are networked differently than in Hopfield nets: They are far less densely interconnected. For a long time it was not understood how they were able to store information as densely as a fully-interconnected Hopfield net is (whose memory capacity per neuron is actually relatively small compared to other ANN architectures of a kind not known in nature), and have essentially the same mathematical behavior. Recently it has been found that a network of oscillating processors, weakly and sparsely interconnected, can display essentially the same characteristics as a fully connected Hopfield net when driven by an external pulse generator [18-20]. In the human brain, this pulse generator is the thalamus, which serves as a central clearing house for emotions. Thus, a central role of emotion from an evolutionary perspective is likely as a means for efficiently inducing different kinds of memory formation) They do so not because of any superimposed external rule (as when a memory is imposed or “written” into an address in standard “Random Access Memory”), but in consequence of a spontaneous order that arises as each element interacts with any other element to which it is connected or related (or, as the case may be, to some common data input such as a recent string of thalamic impulses or a recent history of prices for a commodity).

The “knowledge” or “intelligence” or “memories” are not stored in any one location or set of discrete locations but rather are tacitly embedded and distributed in the whole—in the structure of relations and interactions among the elements. In the discipline of computational science, such information processing (or intelligence) is called “distributed,” as in “distributed processing,” or sometimes, “massively parallel processing” and is characterized by tacit “cooperativity” rather than explicit “cooperation”. Unsurprisingly, just as distributed processing can occur within a set of miniature processing elements inside a single desktop computer, distributed processing can (and does now) take place among worldwide networks of computers, with each individual computer serving as a single element within the global network—even if each individual computer is itself a massively parallel network of elements at a smaller scale. Likewise, if a human single cortex is understood as a sheet of neurons engaged in self-organized parallel computation, then a group of individuals can plausibly be understood as similarly engaged in massively parallel computation—that is, in a process of problem solving, or of self-organization in which solutions to problems arise as a global order which while highly intelligent is the product of neither intention nor design, but the inevitable consequence, rather of completely local (i.e., individual agent) adaptation and action.

## **2.2 The “law of unintended consequences”**

Even if it is now an accepted principle among researchers into “complex systems”, the ubiquity of cooperativity in the absence of deliberate coordination is a reality that is not merely alien but offensive to “common sense”. We have noted that human beings prefer to believe they are in control of their destiny. The commonplace inability to optimize while intending to; often to obtain worse results from the attempt; and indeed even to attain the best results from doing that which appears the most counter-intuitive (“anti-optimization”), is known in sociology and politics as “the law of unintended consequences”. The law is widely described in academic sociology and economic theory (the

first formal analysis of this law appeared in an extremely influential 1936 article by sociologist Robert K. Merton: “The Unanticipated Consequences of Purposive Social Action”[21].); in politics and practical finance it is as widely ignored.

The phenomenon of unintended consequences is documented primarily empirically through studies that retrospectively show how the social and/or economic consequences of a deliberate policy intervention fail to match intentions or even, perversely, worsen a condition the intervention was intended to ameliorate. An all-too common example in finance is the “whipsaw” phenomenon: During periods of high volatility, traders act emotionally in accord with the immediately preceding market move only to fall victim to its reversal at the moment they act.

Another more complex (and quantitatively analyzed) historical example arose during the so-called “Negative Income Tax Experiment (NIT) conducted in the United States between 1968 and 1980 (at which time a second ten-year study sample was cancelled two years in because of the unanticipated pernicious results confirming what was found in the first ten-year sample). This real-world controlled experiment in social policy studied 8,700 individuals. By its termination in 1980 it had generated more than one hundred publications. The costs of research and administration were in the many millions at least; accurate tallies ultimately were lost track of.

In brief, the NIT provided payments to individuals whose income fell below a given threshold. It was intended to mitigate the detrimental effects of poverty on labor force participation, marital stability, family formation, fertility, migration and many other conditions. Numerous sites around the U.S. were selected for experimentation and in each region the subject population was randomly split into an experimental group that received a NIT and a control group that did not, both groups being matched in advance demographically. By 1980 it had become clear that its effects in every area were essentially opposite to what was intended: By almost every measure, the NIT-receiving group was both worse-off and less-contributory to society than the control group.[22-25]. Furthermore, the groups with the most detrimental results (e.g., young, male non-heads of families) were those with the greatest potential to effect long-term changes on the poverty rate — either for better or for worse[26].

On the other hand, interventions that will surely seem counter-intuitive, if not wholly irrational, can unexpectedly produce a contrary positive result—in the spirit of “reculer pour mieux sauter.” The long-standing political argument over whether reducing marginal tax rates may actually yield higher tax revenues to the taxing entity reflects the belief of a substantial part of the economics profession (albeit far from all of it) that precisely such a counter-intuitive method yields optimal results. Another example in finance—not at the level of policy-making and “social engineering”, but at the practical level of day to day accounting—is presented in ref.s [27-30]: Accounting tools are meant to increase managers’ accuracy in forecasting and budgeting. But the use of these tools generates an excessive degree of optimism which is often more than enough to offset the tools’ benefits.

## **2.3 The physics of complex systems**

But many such debates in politics, sociology and even economics take place qualitatively using history (or anecdote) and arguments from political philosophy. When quantitatively, as in economics and sometimes sociology, empirical data must be used that is rarely clean and free from confounding variables. (The NIT example above is a rare instance of a controlled—if scarcely double-blind!—socio-economic experiment carried out on populations in the real-world.) Ideally, one would like to be able to complement a discussion of genuine versus illusory optimization using mathematically definable agent-based models that map pertinent features of complex systems arising both in nature and in society.

The MG has elicited widespread interest precisely because it has this characteristic and yet its rules are remarkably simple. (See the computer pseudo-code in ref. [5].) By contrast, many agent-based simulations that have arisen in the sociology literature deploy agents that are extremely complex on the assumption that agents able to model the salient behavioral features of human beings must necessarily be intrinsically complicated (See the computer code in ref. [31].) Systems of such complicated agents are susceptible of numerical simulation and statistical review but of very little in the way of analytical expression. To be fair, it must be said—and the issue will arise in this work—that many real-world situations involve CAS where even if an equilibrium state is susceptible of full analytic expression, in the real world equilibrium is not attainable—these real-world situations involve systems that are always far from equilibrium. In other instances, the dynamic of interest is effectively always non-equilibrated. For example, an argument can be made that because of the constant introduction of endogenous factors such as innovation in both product per se and trading method; and also because of an ever-shifting array (of a very large number) of external influences such as interest-rates, foreign exchange, politics and climate, no time-series generated by a real-world market is capable of attaining a steady-state, regardless of the number of time-steps.

And here we dealing only with a “simple” one-dimensional time-series (price or price-change). The best work on agent-based sociology is thus inevitably driven to deploy agents so complex in their behavior that they are themselves best considered CAS, and the CAS arising out of their interactions can be mathematically analyzed only in approximation. A classic example is that of the relation between ANN and the neural structure of the brain (i.e. between computational neuroscience and neuroscience proper). The rules embedded in individual processing elements in many ANN yield fully analytic expressions for the nets as a whole. But real world-neurons have a plasticity in function and even physical form that make it evident that each is itself a highly complicated CAS whose behavior arises out of lower-level biological structures and biochemical signaling.

The MG thus models the behavior of agents in a market to a remarkable degree considering its simplicity. In particular it models the phenomenon of agents competing for a limited resource. This kind of competition arises not only in certain kinds or aspects of financial markets, it also models other kinds of human behavior, i.e., competing for the fastest lane on a multi-lane road—as soon as a majority of drivers choose the presently fastest lane, it becomes instead the slowest. Assuming that all drivers are equally clever in the switching strategies, the smartest move would be to simply stay put in any randomly selected lane [32]. Other far more seemingly complex human social behaviors

might be well modeled by the MG, and the illusion of control: for example the quest for prestige, as exemplified by Groucho Marx: “I refuse to join any club that would have me as a member”.

In that spirit, we demonstrate that if, in the Time-Horizon Minority Game, a subset of agents deploy what appears to be their worst-performing strategies instead of their best, as do standard optimizing agents, these not only outperform the standard agents, they generate returns for themselves that can be absolutely positive. On the other hand, if all agents adopt the counter-intuitive method this effect disappears.

In 1936, when the law of unintended consequences was first formally denoted as such and studied, attempts at deliberate social engineering were in their infancy. Nowadays it is common, if not universal, for impact studies of social policies to be required prior to implementation. (Arguendo: Most of these studies leave much to be desired with respect to their accuracy and predictive capacity both because of a-priori political distortions and because such accuracy may be simply impossible to attain. When this is so, these studies have little value from a scientific perspective and serve primarily to reinforce the political considerations.) In the economic domain, complex as it is, a measure of success may be found: For example, in the very mixed “free” (private) cum “socialized” (Medicare, Medicaid) market that characterizes medicine in the United States, arguments on behalf of changes in Medicare payments to physicians are now routinely accompanied by analyses of the so-called “behavioral” effect of a change [33]: Decreasing physician compensation for a certain procedure will to the first order reduce the cost to government of that procedure. But it is presumed that if the cost is lower the demand for it will increase among patients (elementary pricing theory) *and* that physicians will both look for plausible opportunities to increase the procedures use and to substitute alternate procedures to protect their income. If the combination of elementary price elasticity plus physician behavioral changes result in no net change to physician income, the intended savings will be wholly lost. Worse than that, physicians may grow to perceive their Medicare income as being routinely subject to arbitrary seeming diktats. They may thus adapt (as a bacterium does to ever-more powerful antibiotics) with ever more sophisticated methods for eluding the intended cost-reductions. In the practice of medicine, the balance of time and effort devoted by the physician to fiduciary professionalism on the one hand and protective self-interest on the other may be adversely shifted toward the latter, damaging the profession’s collective benefit to society.

The MG has provided a simple analytic model for studying the incorporation of impact on a self-organizing system. If in making their decisions a small subset of agents in the THMG take into account their own impact on the system as a whole, they perform significantly better than agents that do not. If all agents take into account their impact, the system as a whole performs better, but with a caveat: The frustration that makes the MG so useful a model for markets (and by extension other social systems) disappears. The system settles into a deterministic Nash equilibrium which cannot be said to model typical complex systems’ behavior.

Furthermore, the outperformance of a minority of agents that account for their impact is not as great as the outperformance of a same-sized minority that simply “anti-optimizes”, as we will discuss.

Fiske notes that the illusion of control is especially strong under certain circumstances [34]:

- When the choice of action is carefully considered
- When the choices of action are known
- When the individual's actions contribute to the overall state ("involved in the situation", rather than an observer)
- When the overall state arises from (or is perceived as arising from) a competition with other individuals

He observes that "this perfectly describes the situation of a person deciding whether they can get away with a social transgression." We note that it rather more accurately describes the situation of an agent in the MG or THMG. These, of course, are "toy markets" and the four descriptors above equally-well describe the situation of a trader. Perhaps in the reference to "social transgression" we see an example of a common stereotype prevalent among majorities that the winning minority in a market system are e.g., "robber barons" and that therefore "property is theft".

We defer further political philosophizing to emphasize that even simple collective games such as the MG can be studied as models of not just markets but social systems and that the illusion of control is a ubiquitous phenomenon deserving of quantitative and analytic study. To the extent that there genuinely exists an "illusion of control", both losers' negative assessments of why winners win (theft) and winners' self-important assessments of why they do (skill) are inaccurate. (As an example of such self-importance, consider Leo Hunderly, Jr. on his net worth of 150 million USD: "...there are people, including myself...who because of their uniqueness warrant whatever the market will bear." [35]. By contrast, a much earlier folk source asserts that "...the race is not to the swift, ... nor yet riches to men of understanding, nor yet favour to men of skill; but time and chance happeneth to them all." [36]. The same points is mathematically formalized in a recent model explaining the Zipf law of (inter alia) wealth distribution and firm size: Ref. [37] demonstrates that the law is attributable more to the random part of growth processes than to systematic drift—i.e., though both are as a rule present, chance dominates talent. A similar point is made in appendix chapter A10, where simulations demonstrate that the rank changes in an equity ranking system widely used by both individual and professional investors are attributable more to underlying stochastic processes than to intelligent assessment of rank.

Beliefs to the contrary are in accord with the extensively documented fact that people are "unrealistically optimistic especially when they extrapolate from their own past experience" [34], consistent with the illusion demonstrated in the MG and THMG.

However, markets (or phases of a market) are not exclusively characterized by the minority mechanism. For example, market bubbles are characterized by the herd mentality (to wit: "The Madness of Crowds"[38]), i.e., the tendency of people to join a *majority*, and therefore for a time at least, in a self-fulfilling prophesy (positive feedback) drive up the price of a commodity. In a real world market, the underlying commodity is ultimately something tangible and it, of course, is in limited supply, otherwise no amount of chasing after it could drive the price up. But in an artificial market such as the MG, it is, as it

were, the price itself that agents chase after. And in real-world bubbles, this is what a majority of traders appear to do as well. The so-called “momentum” model of trading is based on the knowledge that regardless of underlying fundamentals, price sequences trend in one direction or the other more frequently than would be expected of a random walk. This leads to a greater frequency of extended rises. Thus many traders and certain phases of markets may be better characterized by a “Majority Game” (MAJG), where agents attempt not to be in the minority (and lose when they are) but the majority.

However, in order to realize a paper gain of any sort, especially during a bubble, a trader must sell his holdings, i.e., leave the majority. But if a majority of traders do this at the same time, the bubble breaks and a large market drawdown ensues—the same mechanism as the bubble but in reverse. In this case, a trader scarcely wants to be in the majority, and a MAJG would seem intuitively to be incapable of capturing this dynamic.

In fact, the ideal trader wants to be in the majority just as the market begins a rise, but in the minority just as it begins a decline. The “\$-Game” (\$G) is structured identically to the MG and MAJG but the function describing the payoff to agents is altered to capture this shifting objective [39]. All three games in their “time-horizon” variants may be expressed analytically for a finite number of agents in terms of Markov transition matrices. The relative simplicity of expression highlights the different circumstances under which the “illusion of control” emerges in differing CAS.

It is important to ask whether the illusion of control—to the extent that it is present—is somehow unique to the THMG and/or MG, or whether it may be found in other settings. As an adjunct to the primary discussion we demonstrate that a similar illusion may be found in the most natural setting of a very different kind of game, the single-player Parrondo game. A Parrondo game is a game-theoretic formulation of a physical phenomenon, the asymmetric drift of a charged particle in a so-called flashing-ratchet, i.e., an on-and-off sawtooth-shaped potential: A particle executing 1D random motion in such a constant potential will demonstrate no net drift from its original position. But if the potential is flashed altogether off and on, either at random times or cyclically, the particle will drift in the direction of the shallower of the two slopes of the “teeth”. Note that the mean potential is simply a sawtooth ratchet of half the maximum excursion in which a particle would show no drift. Thus a seeming “paradox” arises if one incorrectly presumes that the flashing ratchet potential (say  $\frac{1}{2}$  the time on,  $\frac{1}{2}$  off) is equivalent to a stable ratchet potential of  $\frac{1}{2}$  height: One is tempted to think, at first, that the flashing potential must be equal to the mean of both since the instantaneous change in the potential off and on exerts no lateral force on the particle.

This phenomenon can be analogized to a player who alternates among three negatively biased binary games (unfair coin-tosses). Under certain constraints the player will on average nonetheless win—the so-called “Parrondo Paradox” or “Parrondo Effect” (PE) [6, 7]. At first glance, this seems as though it were like a gambler winning with statistical certainty against a house in Las Vegas simply by shifting at random between say Roulette and Craps—an evident impossibility. In fact, the Parrondo Effect is unlike this and cannot be used to devise winning gambling strategies. Using the theory of Markov chains it has been shown that there is no real paradox in Parrondo games, as there is none in the flashing ratchet. The genuine phenomenon leads naturally to the question, “Is it possible to optimize the gains obtained from playing losing games?”

Furthermore, it turns out that if one sets up a certain (peculiar) optimization scheme involving *many* players, an attempt to optimize the outcome yields the reverse result: Instead of doing better, the players do worse, obtaining negative results even against the positive “current” (drift) induced by the Parrondo Effect. We demonstrate using a setting more natural for the Parrondo game (which, unlike the MG, is inherently single-player) that the attempt to optimize using prior history yields a similarly strong reversal of outcome, against the “current” of the Parrondo Effect proper. Attempts have been made to apply the Parrondo Effect to trading. These have unsurprisingly been unsuccessful. Our results demonstrate that the attempt to achieve success by optimizing is likely not only to fail but to yield even worse results[40].

We turn now to Chapter 3. wherein we summarize the main findings of this research.

## Chapter 3. The “illusion of control” in agent-based and Parrondo Games

In this chapter we survey and summarize the main results of our study of three classes of market-modeling CAS games (MG, MAJG, \$G) and a fourth non-CAS game (PG). We concentrate first on the MG and THMG, using it and the pertinent results as the prototype. We then extend our observations to the MAJG and \$G. We discuss the important differences, both theoretical and applied, between the standard MG (and the MAJG and \$G modeled on it) and the so-called “time-horizon” variants of these games. We describe our results quantitatively and graphically, and defer an in-depth presentation of mathematical methods for the Appendix.

### 3.1 “Illusion of control” in the Time-Horizon Minority Game (THMG)

The success of science and technology, with the development of ever more sophisticated computerized integrated sensors in the biological, environmental and social sciences, illustrate the quest for control as a universal endeavor. The exercise of governmental authority, the managing of the economy, the regulation of financial markets, the management of corporations, and the attempt to master natural resources, control natural forces and influence environmental factors all arise from this quest. Langer’s phrase, “illusion of control” [3] describes the fact that individuals appear hard-wired to over-attribute success to skill, and to underestimate the role of chance, when both are in fact present. Whether control is genuine or merely perceived is a prevalent question in psychological theories. The following presents two rigorously controlled mathematical set-ups demonstrating generic circumstances in which optimizing agents perform worse than their non-optimized strategies, or than non-optimizing (fixed- or random-choice) agents.

#### 3.1.1 Definition and summary of main results for the Time-Horizon MG (THMG)

We first study a variant of Minority games (MGs), which constitute a sub-class of market-entry games. MGs exemplify situations in which the “rational expectations” mechanism of standard economic theory fails. This mechanism in effect asks, “what expectation model would lead to collective actions that would on average validate the model, assuming everyone adopted it?”[41]. In minority games, a large number of interacting decision-making agents, each aiming for personal gain in an artificial universe with scarce resources, try to anticipate the actions of others on the basis of incomplete information. Those who subsequently find themselves in the minority group gain. Therefore, expectations that are held in common negate themselves, leading to anti-persistent behavior both in the aggregate and for individuals. Minority games have been much studied as repeated games with expectation indeterminacy, multiple equilibria and inductive optimization behavior.

Consider the Time-Horizon MG (THMG), where  $N$  players have to choose one of two alternatives at each time step  $t$  based on information represented as a binary time series  $A(t)$ . The action chosen by agent  $i$  is denoted  $a_i(t)$  and the binary time series by

$A(t) = \sum_{i=1}^N a_i(t)$ . Those whose choice is the one chosen by the minority win. Each agent is endowed with  $S$  strategies. Each strategy gives a prediction for the next outcome  $A(t)$  based on the history of the last  $m$  realizations  $A(t-1), \dots, A(t-m)$  simplified into binary form—0 replaces  $A(t)$  when one state is in the minority, i.e.,  $A(t) < 0$ , 1 when the other is. ( $m$  is called the memory size [length of the collective history] used by the agents). Each agent holds the same number  $S$  of (in general different) strategies among the  $2^{2^m}$  total number of possible strategies. (A strategy associates every possible  $m$ -bit sequence  $\mu(t)$  with a predicted next binary bit.) The  $S$  strategies of each agent are chosen at random, with replacement, once and for all at the beginning of the game. At each time  $t$ , in the absence of better information, in order to decide between the two alternatives for  $A(t)$ , each agent uses her most successful strategy in terms of payoff accumulated in a rolling window of finite length  $\tau$  up to the last information available at the present time  $t$  (the case of a limitlessly growing  $\tau$  corresponds to the standard MG; the term “Time Horizon MG” refers to the case of a fixed and finite  $\tau$ ).

This is the key optimization step. If her best strategy predicts  $Sgn[A(t) = +1]$  (resp.  $-1$ ), she will take the action  $a_i(t) = -1$  (resp.  $+1$ ).  $a_i(t)$  is then added to the information set available for the next iteration at time  $t+1$ , along with its associated payoff. In the simplest instance of the minority rule, the corresponding instantaneous payoff of agent  $i$  is given by

$$g_i(t) = -Sgn[a_i(t)A(t)] \quad (1)$$

(and similarly for each strategy for which it is added to the  $\tau - 1$  previous payoffs). As the name of the game indicates, and as the minus sign in eqn. (1) indicates, if a strategy is in the minority ( $a_i(t)A(t) < 0$ ), it is rewarded. In other words, agents in the THMG try to be anti-imitative. The richness and complexity of minority games stem from the fact that agents strive to be different. Previous investigations have shown the existence of a phase transition marked by agent cooperation and efficiency between an inefficient regime and a random-like regime as the control parameter  $\alpha \equiv 2^m/N$  is increased: In the vicinity of the phase transition at  $\alpha_c = 2^{m_c}/N \approx 0.34$  (for both the THMG and MG proper), the size of the fluctuations of  $A(t)$  (as measured by its normalized variance  $\sigma^2/N$ ) falls below the random coin-toss limit for large  $m$ 's (assuming fixed  $N$ ) when agents always use their highest scoring strategy [4]. In other words, for a range of  $m$  (given  $N, S$ ), agent performance is better than what strategy performance would be in a game with no agents optimizing. The phenomenon discussed here is that when optimizing, and averaged over all actual agents and strategies in a given realization, agents in the TH variant of the MG nonetheless generally underperform the mean of their own measured strategy performance and do so in all phases for reasonable lengths of  $\tau$  (as also the mean over all strategies in a given realization; all this under the restriction that agents do not account for their own impact).

For any given realization, however, a minority of agents outperform the mean of their strategies and the majority of other agents. Some may also achieve net positive gain, if rarely.). In the MG proper, however,  $\tau$  is unbounded and a stationary state is reached at some very large  $\tau_{eq} \geq 2^m \times 200$  where a subset of agents “freeze” their choice of strategy:

One virtual strategy score attains a permanently higher value than any other. These frozen agents in general do outperform the mean of all strategies in a given realization as well as the mean of their own  $S$  original strategies: They perform precisely as well as their best. We focus on results in the THMG with an eye towards real-world markets in which because the time series being predicted are non-stationary, trading strategies are weakened if they incorporate an unbounded (and uniformly-weighted) history of prior strategic success or failure: Remote history is less important than recent history and beyond a certain point is meaningless. Unless specifically stated otherwise, throughout this paper, whenever we compare agent to strategy performance, we always mean the performance of agents’ strategies as measured by the accumulation of hypothetical points averaged over *all* agents in the system and the set of *all* of their strategies. Furthermore, in selecting a strategy the agents do not take account of the impact of their choice on the probable minority state—that is, they do not consider that their own selection of action reduces the probability that this action will be the minority one. (We refer to such agents as “standard”.)

### 3.1.2 Statement of our main results on the “illusion of control” in the THMG

Our main result may be stated concisely from the perspective of utility theory: Throughout the space of parameters ( $N, m, S, \tau \ll \tau_{eq}$ ), the mean payoff of agents’ strategies (as calculated by each agent averaged over all strategies and agents in a realization) not only surpasses the mean payoff of supposedly-optimizing agents (averaged over all given agents), but the respective cumulative distribution functions (CDF) of payoffs show a first-order stochastic dominance of strategies over agents. Thus, were the option available to them, agents would behave in a risk-averse fashion (concave utility function) by switching randomly between strategies rather than optimizing. This result generalizes when comparing optimizing agents with  $S > 1$  strategies with agents having only one strategy (or equivalently  $S$  identical strategies), when the single strategies are actually implemented. (This takes into account any difference in strategy performance that may arise from the simple fact of a strategy actually being deployed). The same result is also found when comparing optimizing agents with agents flipping randomly among their  $S$  strategies. Agents are supposed to enhance their performance by choosing adaptively between their available strategies. In fact, the opposite is true: By our metric, the optimization method would appear to agents as strictly a method for worsening performance. (In the MG proper the situation is more complex. As detailed in [4], agents with two identical strategies—equivalent to having only one [and called therein “producers”]—always have net gain  $\leq 0$ . But this gain may be on average either greater or less than for agents that optimize among more than one strategy [called “speculators”]. Which is true depends inter alia on the proportion of producers to speculators: A very small proportion of producers will outperform speculators. There is an expected proportion of producers that arises from the average over many different random possible initial allocations of strategies among agents [i.e., quenched disorder]. Given this expected mean proportion, and averaging over *all* the agents in each initial allocation, the mean performance of strategies

is better than that of agents. This is true, however, only when agents do not choose their strategies at each step taking into account the impact of that selection.)

Let us restate our result for the THMG in the language of a financial market with traders trying to outperform the overall market. We argue that in using the THMG as a model for traders' actions, the following is the case: Every trader attempting to optimize by selecting his "best performing strategy" measures that performance virtually, not by contrast to an imagined setting where all traders select fixed strategies at random (to whose results he would have no access anyway). Even though the virtual performances of each of his basket of strategies might never have been implemented in reality, if he found that his real performance under a selection process was worse than the virtual performance of the strategies he had been selecting among, he would abandon the selection process. This would be true for most agents and not true only for a small minority. (If *every* trader were to do the same, of course, then one would end up with the random or fixed choice game as discussed below. This forms the usual standard of comparison for strategy performance in the MG literature. In reality, many traders, especially hedge funds and other large investment pools do attempt to account for their own impact in trading, but it is arguable that these assessments are accurate when averaged over all traders.)

The above argument echoes the finding of Doran and Wright [42], who report that two-thirds of all finance professors at accredited, four-year universities and colleges in the U.S. (arguably among the most sophisticated and informed financial investors) are passive investors who think that the traditional valuation techniques are all *unimportant* in the decision of whether to buy or sell a specific stock (in particular, the CAPM, APT and Fama and French and Carhart models).

### 3.1.3 Quantitative statement and tests

In the THMG, the "illusion of control" effect is observed for all  $N, m, S$  and  $\tau \ll \tau_{eq}$ . We use the Markov chain formalism for the THMG [9,10] to obtain the following theoretical prediction for the gains,  $\Delta W_{agent}$  averaged over all agents and  $\Delta W_{strategy}$  averaged over all strategies respectively, of agents and of all strategies in a given realization [43]:

$$\langle \Delta W_{Agent} \rangle = \frac{1}{N} \bar{A}_D \cdot \bar{\mu} \quad (2)$$

$$\langle \Delta W_{Strategy} \rangle = \frac{1}{2N} (\hat{\mathbf{s}}_{\mu} \cdot \bar{\mathbf{k}}) \cdot \bar{\mu} \quad (3)$$

Brackets denote a time average.  $\mu$  is a  $(m + \tau)$ -bit "path history" [9] (sequence of 1-bit states);  $\bar{\mu}$  is the normalized steady-state probability vector for the history-dependent  $(m + \tau) \times (m + \tau)$  transition matrix  $\hat{\mathbf{T}}$ , where a given element  $T_{\mu_t, \mu_{t-1}}$  represents the transition probability that  $\mu_{t-1}$  will be followed by  $\mu_t$ ;  $\bar{A}_D$  is a  $2^{(m+\tau)}$ -element vector listing the particular sum of decided values of  $A(t)$  associated with each path-history;  $\hat{\mathbf{s}}_{\mu}$  is the table of points accumulated by each strategy for each path-history;  $\bar{\mathbf{k}}$  is a  $2^{(m+\tau)}$ -element vector listing the total number of times each strategy is represented in the collection of  $N$  agents. As shown in the supplementary material,  $\hat{\mathbf{T}}$  may be derived from  $\bar{A}_D$ ,  $\hat{\mathbf{s}}_{\mu}$  and

$\bar{N}_U$ , the number of undecided agents associated with each path history. Thus agents' mean gain is determined by the non-stochastic contribution to  $A(t)$  weighted by the probability of the possible path histories. This is because the stochastic contribution for each path history is binomially distributed about the determined contribution. Strategies' mean gain is determined by the change in points associated with each strategy over each path-history weighted by the probability of that path.

We find excellent agreement between the numerical simulations and the analytical predictions (2) and (3) for the THMG. For instance, for  $m = 2$ ,  $S = 2$ ,  $\tau = 1$  and  $N = 31$ ,  $\langle \Delta W_{Agent} \rangle = -0.22$  for both analytic and numerical methods (payoff per time step averaged over time and over all optimizing agents) compared with  $\langle \Delta W_{Strategy} \rangle = -0.06$  also (similarly averaged over all strategies) for both analytic and numerical methods. In this numerical example, the average payoff of individual strategies is larger than for optimizing agents by 0.16 units per time step. The numerical values of the predictions (2) and (3) are obtained by implementing each agent individually as a coded object.

In the THMG, the mean per-agent per-step payoff  $\langle \Delta W_{Non-Opt} \rangle$  accrued by non-optimizing agents (they have only one fixed strategy, or equivalently their  $S$  strategies are identical; a.k.a. “producers”) is larger than the payoff  $\langle \Delta W_{Agent} \rangle$  of optimizing agents (a.k.a. “speculators”). In general, this comparative advantage decreases with their proportion but much less rapidly than in the MG proper [6]. For example, with  $m = 2$ ,  $S = 2$ ,  $\tau = 1$  and  $N = 31$ , and 2500 random initializations and  $N_{opt}$  optimizing agents,  $\langle \Delta W_{Non-Opt} \rangle - \langle \Delta W_{Agent} \rangle = (2.380, 2.270, 2.289, 2.275, 2.145, 2.060, 2.039, 1.994, 1.836, 1.964) \times 10^{-3}$  for  $N_{opt} = 1, 2, \dots, 10$ . More generally, the following ordering holds: payoff of individual strategies  $>$  payoff of non-optimizing agents  $>$  payoff of optimizing agents. The first inequality is due to the fact that not all individual strategies are implemented and the theoretical payoff of the non-implemented strategies does not take into account what their effect would have been (had they been implemented). Implementation of a strategy tends to decrease its performance (this is similar to the market impact of trading strategies in financial markets associated with slippage and market friction). Non-optimizing agents by definition always implement their strategies. However, the higher payoff of non-optimizing compared with optimizing agents shows that the illusion-of-control effect is not due to their actually being deployed, but is a genuine observable effect.

### 3.1.4 Generalizations

The amplitude of the illusion-of-control effect in the THMG highlights important differences between the MG proper, in which  $\tau$  is sufficiently large so as to allow the system to attain equilibrium with many “frozen” agents ( $\sim 10^4 - 10^6$  time steps; though some agents will freeze relatively quickly) and the THMG in which  $\tau$  is arguably of a length comparable to real-world investment “lookbacks”. (A reminder: A strategy in both the MG/THMG and in many real-world “technical” trading methods, relates  $m$  relatively recent actual changes in price to a future predicted change. The hypothesized utility of different strategies is “back-tested” over different historical “lookback.”) The effect also highlights the distinction between optimizing agents with  $S$  maximally distinct strategies

(in the sense of Hamming distance) and non-optimizing agents with  $S$  identical strategies—a distinction with different characteristics in the THMG than in the MG proper.


It is helpful to generalize the latter distinction by characterizing the degree of similarity between the  $S$  strategies of a given agent using the Hamming distance  $d_H$  between strategies. (The Hamming distance between two binary strings of equal length is the number of positions for which the corresponding symbols are different, normalized on the unit interval.) Non-optimizing agents with  $S$  identical strategies correspond to  $d_H = 0$ . By contrast, optimizing agents with  $S$  maximally distinct strategies correspond to  $d_H = 1$ . Since on average agents in the THMG with  $d_H = 0$  out-perform agents with  $d_H > 0$ , it is natural to ask whether the ranking of  $d_H$  could be predictive of the ordering of agents' payoffs. But first it is important to clarify differences with respect to  $d_H$  in the MG versus the THMG.

The first difference to emphasize is that in the MG, where the system runs to equilibrium prior to making performance assessments, one of the chief features of the stationary state is that some (and sometimes many, i.e., max 39%) of the optimizing agents i.e., speculators" which by definition must have strategy  $d_H > 0$ ) "freeze". That is, the effective  $\tau$  is long enough so that one of the  $S$  strategies for some agents will attain permanently the largest number of "virtual" points. It will then always be deployed, somewhat similar it might seem (but only at first glance), to an agent with  $S$  identical strategies from the start.

The second difference is that at equilibrium in the MG, the relation between agent performance and  $d_H$  inverts at the critical point  $\alpha_c$  [44]: On average, for  $\alpha > \alpha_c$ , agents with larger  $d_H$  outperform those with smaller  $d_H$ —and outperform the mean over the selected strategies. This reversal is due to the freezing of a subset of agents which is *not* the same as an agent having been assigned two identical strategies from the start. Over the very long run-up to equilibrium, frozen agents have the opportunity to choose what is in fact a better strategy, but so little better in general that it requires a great many time-steps for the difference to emerge and become embedded in the system. It is unsurprising that a larger Hamming distance between strategies offers more opportunity for such a differentiation to occur. Conversely, for agents with  $d_H = 0$ , such selection is by definition impossible.

Note that for extremely short  $\tau$  (e.g., 1, 10), the phase-transition appears attenuated to the point of being undetectable by this metric: Rather, mean agent performance increases monotonically and approaches asymptotically that of mean strategy performance. As  $\tau$  increases a number of things happen. First, the phase transition at  $\alpha_c$  appears and grows increasingly sharp. Second, the overall scale of agent return (comparably, volatility of  $A(t)$ , i.e.,  $-\sigma^2/N$ ) as a function of  $m$  oscillates (rather dramatically, in fact) with a period equal to  $2 \times 2^m$  for real histories [43]—but does not vary for random ones. (A major point of discussion in the MG literature, addressed only tangentially in this thesis, is the interesting fact that the game's hallmark phase-transition appears even if the "histories" to which agents respond are not genuine but mere random sequences freshly generated at every  $t$ .. The phase transition under random histories further stands in the same relation to

$m$  as with real histories. The critical feature that preserves the transition in its sharpest form—with its implied maximum cooperativity—and its relation to  $m$  is that all agents must respond to the same history of length  $m$  at each  $t$ . However, there are subtle differences in the volatility of  $A(t)$  and in agent performance as a function of phase, between games with random and real histories. Some of these differences are referenced with respect to the “illusion of control” in Chapter 7. See ref.’s [45-49] for extended discussions of this point.) Third, so long as  $\tau$  remains “reasonable”, a reversal of the relation between  $d_H$  and agent performance does not occur (i.e., the larger the  $d_H$ , the smaller the agent gain in wealth, for all  $\alpha$ ). “Reasonable” lengths for  $\tau$  in the THMG, from the perspective of relative stationarity in real-world financial time-series, cannot be denoted without taking into account the regime: For a modest number of agents at small  $m$  (e.g., for  $N = 31$ ,  $m < 4$ ),  $200 \times 2^m$  time-steps is sufficient to reach equilibrium. But near the phase transition (for  $N = 31, m_c = 4$ ), many more steps are required, on the order of  $5000 \times 2^m \geq 80,000$ , equivalent to 320 years of daily price data, assuming that a time step equals one trading day. For  $\tau$  on the order of  $10^3$ , no reversal of the relation between small  $d_H$  and better agent performance occurs. **Figure 2** provides a typical example of the distribution of agent returns by  $d_H$ . A similar distribution with negative linear slope occurs for all reasonable values of  $m$  and  $\tau$  short of  $\tau_{eq}$ .



**Figure 2:** Typical distribution of agent returns by Hamming distance between component strategies. A simple linear fit with slope  $\theta(t) = -.061$  demonstrates decreasing mean agent performance with increasing  $d_H$ ; the time-horizon  $\tau = 2570$  is even longer than “reasonable”; it is more than long enough for a sharp phase transition to be present at  $3 \leq m \leq 4$  (but still well short of equilibrium, i.e.,  $\tau < \tau_{eq} \approx 12,800$ ); the memory  $m = 6$  is well past the phase transition after which in the MG proper the relation between  $d_H$  and  $\Delta W_{agent}$  inverts.

A non-zero  $d_H$  implies that there are at least two strategies among the  $S$  strategies of the agent which are different. But, if  $d_H$  is small, the small difference between the  $S$  strategies makes the optimization only faintly relevant ( $d_H$  small implies that the same action

will be taken most of the time regardless of which strategy is selected) and one can expect to observe a payoff similar to that of non-optimizing agents, therefore larger than for optimizing agents with large  $d_H$ 's. This intuition is indeed confirmed by our calculations: the payoff per time step averaged over all agents is a decreasing function of  $d_H$ , as originally discussed in [44, 50] for the MG at equilibrium (and for  $\alpha < \alpha_c$ ).

The illusion of control effect suggests that the initial set-up of the THMG in terms of  $S$  fixed strategies per agent is evolutionarily unstable (when agents do not select strategies taking account of their impact). It is thus important to ask what happens when agents are allowed to replace strategies over time based on performance. A number of authors have investigated this issue in the MG, adding a variety of longer-term learning mechanisms on top of the short-term adaptation that constitutes the basic MG [44, 50-55]. Inter alia, ref. [51] demonstrates that if agents are allowed to replace strategies over time based on performance, they do so by ridding themselves of those composed of the more widely Hamming-distant tuples. Agents that start out composed of identical strategies do not change at all; those composed of strategies close in Hamming space change little. Similarly, the authors of [50] explicitly fixed agents with tuples of identical strategies and found they performed best. Another important finding in [50] is that the best performance attainable is equivalent to that obtained by agents choosing their strategies at random. Note that learning only confers a *relative* advantage. In general, agents that learn outperform agents that don't. This is certainly true for this privileged subset of agents among standard ones. But the performance of learning agents approaches a maximum most closely attained by agents where the Hamming distance between strategies is 0. These agents neither adapt (optimize) nor learn. One might say that when learning is introduced, the system learns to rid itself of the illusory optimization method that has been hampering it. (Note that if one compares optimizing agents' performance to the performance of a separate system composed entirely of non-optimizing agents, there are regimes in  $m$  for which the optimizing agents do better: The standard metric of comparison in the MG literature. This could arise "in reality" only if traders deliberately ignored the evidence most perceive, namely, that the mean of their own strategies appear to be outperforming the optimization process that chooses deliberately among them. We emphasize "most" here, because a smaller proportion of traders' would in fact see their optimization process succeeding. Again, mean agent performance underperforms mean strategy performance when averaged over all agents and all of the strategies represented in a given quenched disorder.)

There are exceptions, of course. Carefully designed privileges and certain kinds of learning can yield superior results for a subset of agents, and occasionally for all agents. But the routine outcome is that both optimization and straightforward learning cannot improve on simple chance, as measured by agents' own assessment of their strategies' respective virtual performance. The fact that the optimization method employed in the THMG yields the opposite of the intended consequence, and that learning eliminates the method, leads to an important question. We pose it carefully so as to avoid introducing either privileged agents or learning: Is the illusion-of-control so powerful in this instance that inverting the optimization rule could yield equally unanticipated and opposite results?

The answer is yes: If the fundamental optimization rule of the MG is symmetrically inverted for a limited subset of agents who choose their worst-performing strategy instead of their best, those agents systematically outperform both their strategies and other agents. They also can attain positive gain. (Details of this phenomenon are presented and discussed in chapter 7.) Thus, the intuitively self-evident control over outcome proffered by the THMG “optimization” strategy is most strikingly shown to be an illusion. Even learning and evolutionary strategies generally at best rid the system of any optimization method altogether. They do not attain the kind of results obtained simply by allowing some agents to reverse the method altogether. We emphasize the fact that extensive numerical studies confirm that the phenomenon here indicated persists over a very wide range of parameters in the MG and over all parameter values in the THMG. Hence, having a portfolio of  $S$  strategies to choose from is in the THMG always counter-productive, and the MG often so: (diversification + optimization) performs on average worse than a single fixed strategy.

Let us also mention briefly a related work by Menche and de Almeida. In the standard MG, the only public information are recent first places, while Menche and de Almeida [56] introduce a history of second places in the agents’ set of strategies, thus providing more information to the agents about the state of the game and about the quality of their strategies. They find that the resulting performance of the system becomes significantly better and the phase transition into the uncorrelated phase is strongly suppressed. Note that this variation grants agents greater computational capacity than agents in the standard MG. For  $S = 2$  it comes close to the simple variation we explore, where some agents choose their worst strategy instead of their best without changing the computational complexity of the game nor of individual agents.

### 3.1.5 Illusion of control and the crowding-out mechanism

Intuitively, the illusion-of-control effect in the MG results from the fact that a strategy that has performed well in the past becomes crowded out in the future due to the minority mechanism: Performing well in the recent past, there is a larger probability for a strategy to be chosen by an increasing number of agents, which inevitably leads to its demise. This argument in fact also applies to all the strategies that belong to the same reduced set; their number is  $2^{2^m}/2^{m+1}$ , equal to the ratio of the cardinality of the set of all strategies to the cardinality of the set of reduced strategies. (A reduced set of strategies—referred to in the literature as a “reduced strategy space”, is a subset of all  $2^{2^m}$  possible strategies that spans the full subset. This is possible because strategies that are Hamming very close differ little in their choice of action and point accumulation.) Thus, the crowding mechanism operates from the fact that a significant number of agents have at least one strategy in the same reduced subset among the  $2^m$  reduced strategy subsets. Optimizing agents tend on average to adapt to the past but not the present. They choose an action  $a_i(t)$  which is on average out-of-phase with the collective action  $A(t)$ . In contrast, non-optimizing agents average over all regimes for which their strategy may be good and bad, and do not face the crowding-out effect. The crowding-out effect also explains simply why anti-optimizing agents over-perform: choosing their worst strategy ensures that it will be the least used by other agents in the next time step, which implies that they will be in the mi-

nority. The crowding mechanism also predicts that the smaller the parameter  $2^m/N$ , the larger the illusion-of-control effect. Indeed, as one considers larger and larger values of  $2^m/N$ , it becomes more and more probable that agents have their strategies in different reduced strategy classes, so that a strategy which is best for an agent tells nothing about the strategies used by the other agents, and the crowding out mechanism does not operate. Thus, regions of successful optimization, if they occur at all, are more likely at higher values of  $2^m/N$ . (See Appendix A for further details.)

### 3.1.6 Robustness of the “illusion of control” phenomenon: THMG versus MG

It could be argued that this example of “illusion of control” is very specific because we consider the THMG, and not the standard MG. In the standard MG, some agents are frozen as a result of lengthy optimization, and some of these agents are able to win more than half of the time. This appears similar to the references we discuss wherein learning takes place and agents learn to “rid” themselves of strategy choice. But in the MG proper the “ridding” takes place at the more fundamental level of the basic optimization procedure and reflects a genuine (non-illusory) control that appears along with the phase transition (requiring an especially lengthy run-up to equilibrium). That some agents are able to win more than half the time when they are frozen is only in part analogous to when there exists a subset of select agents—for example ones that take into account their impact; or agents with two identical strategies (equivalent therefore to being frozen from the start, albeit without any preceding selection process); or agents that choose their worst strategy. From one perspective, the existence of genuine control in certain phases of the MG proper is an artifact of an “unreasonably” long equilibration process (and an equilibrium state arguably not found in real-world markets). From the opposite perspective, the illusion of control in the THMG is an artifact of “transients” relatively early in the system’s equilibration process.

In any event, we observe that in the THMG, agents with two identical strategies on average outperform those selecting among strategies, but do not do better than a 0.5 win rate—again, averaged over all those that do in fact do better than 0.5 and those that don’t. We also observe that agents that always choose their worst strategy (when they are a subset among a majority that choose their best as usual) have a better than 0.5 win rate on average for a number of parameter values. As detailed in [57, 58], when all agents take into account their impact, the agents do now outperform their strategies. However, the game settles into a Nash equilibrium [59] which is arguably an entirely different situation, one in which the “illusion” of control is no longer pertinent as the dynamics are in this case deterministic. (In game theory, a Nash equilibrium arises when every player’s action is optimal given the actions of the other players. A more realistic situation occurs when, for example, only some traders account correctly for impact, or when some or all account for impact only imperfectly. Depending on the extent of impact-accounting, the system may remain frustrated but the illusion of control may still disappear.)

Any agent accounting for impact looks back at the prior vote imbalance and determines what the imbalance would have been had she used each given strategy (not what its score would have been using just the strategy she actually did use). Similar methods are used by real-world traders taking positions large enough to have an impact on price. A consequence of taking into account impact is that, with a certain probability usually smaller

than one, at any time-step, an agent will select some strategy other than the “best” (i.e., as computed in the standard way sans impact). From this perspective, accounting for impact is similar (but not identical) to the computationally simpler act of agents in fact choosing other than the best strategy, always. For  $S = 2$ , this is the same as choosing the worst. Although we report in detail only on the effect of choosing the worst strategy for  $S = 2$ , the same principle holds for  $S > 2$ : a subset of agents choosing their worst strategy outperforms, on average, those that choose their next-to-worst, etc. When a small subset of agents take into account their impact, on average these perform better than those that do not. However, they do not perform as well as a similarly sized subset of agents choosing their worst strategy. The dynamics of a game composed entirely of agents choosing their worst strategy is not similar to the Nash-equilibrated structure of a game with all agents accounting for impact.

It is true in particular that there is a complex relation between maximum/minimum typical system fluctuation/cooperation and length of strategy score memory ( $\tau$ ), which we do not discuss in detail:  $\sigma^2/N$  is periodic in  $2 \cdot 2^m$  [43]. Nonetheless, as many simulations that we have performed illustrate for both real and random histories (see Chapter 7), the underperformance of standard agents vis-à-vis the mean of all strategies represented in a given  $\hat{\Omega}$ , and over many  $\hat{\Omega}$ , is found for all  $\tau$  up to (and greater than) the equilibrium number of steps at which point the THMG becomes equivalent to the MG. The difference between strategy performance (as we define it) and agent performance declines roughly exponentially with  $\tau$  but remains positive. When a critical point is present (sufficiently long  $\tau$  but still well short of  $\tau_c$ ) it reaches its positive minimum at the critical point. We find that strategy out-performance is greater for random histories at this point than for real histories. The phase transition central to the MG is most evident for  $\tau$  long and is attenuated for  $\tau$  short.

Even restricted to the THMG, the phenomena we are most interested in are essentially as prominent for  $\tau \approx 1000$ , say, as for  $\tau \approx 1$ . The rule of thumb for reaching the stationary state in the standard MG is to iterate for about  $200 \times 2^m$  time steps (it takes even more time close to the critical point  $\alpha_c$ ). Thus, for our simulations with  $m = 2$ , values of  $\tau \approx 1000$  and above probe the stationary regime of the standard MG and confirm the robustness of the illusion-of-control effect. It is reasonable to argue that for real-life trading situations, which are generally non-stationary, a 100 or 1000 time-unit “look back” is of significant interest (Real world  $\tau$ 's  $\leq 100$  are not unusual). “Look backs” long enough to achieve equilibrium, even for only tens of agents in the MG are (at least arguably) less likely to happen in reality. (On the other hand, studies that employ tick data—there are  $\sim 6400$  ticks in a day—may arguably require lookbacks on the order of a MG  $\tau_c$  and may be treated as at equilibrium.) If the subset of such agents taking into account their impact is large enough to be meaningful ( $\sim 1/3$  or more), one can see that the performance of these agents distributes itself as symmetrically as possible around zero. The remaining agents “perturb” this equilibrium. When the proportion of agents accounting for impact is large enough (depending on the other system parameters), the system as a whole settles into a deterministic equilibrium and there is no longer a phase transition at critical  $\alpha_c$ . This equilibrium is achieved more readily when  $\tau$  is large yet need not be too large.

However, in a THMG composed entirely of impact-accounting agents, with  $N = 31$ ,  $S = 2$  and  $m = 2$ , a near equilibrium state is attained for  $10 > \tau > 100$ . Also, for  $\tau = 1$  or  $\tau = 10$ , strategies outperform their agents as we have described. **For** , **the reverse is true.** In the MG (with standard agents that do not account for their impact), whenever the fluctuations of the global choice are better than with random ones, the agents perform globally better than in a game composed entirely of non-optimizing standard ones. As discussed in [57], which agents perform best during which phase is highly sensitive to the precise ratio of “producers” (agents with  $S$  identical strategies, hence non-optimizing) to “speculators” (agents with at least two different strategies), and to the degree to which agents have correlated actions as averaged over all histories. Frozen speculators in general perform best of all. We stress that this is not inconsistent with our observation that in the aggregate—not examining these “microscopic” differences among types and proportions of agents—standard agents at all  $\tau$  nonetheless under-perform the mean of all strategies in a given quenched disorder averaged over many different such initial configurations.

### 3.1.7 First-entry games and symmetric evolutionary stable equilibria

The above discussion leads to the conclusion that there is often a profound clash between optimization on the one hand and minority payoff on the other hand: an agent who optimizes identifies her best strategy, but in so doing by her “introspection”, she somehow knows the fate of the other agents, that it is probable that the other agents are also going to choose similar strategies, ... which leads to their underperformance since most of them will then be in the majority. It follows then that an optimizing agent playing a standard minority game should optimize at a second order of recursion in order to win: Her best strategy allows her to identify the class of best strategies of others, which she thus must avoid absolutely to be in the minority and to win (given that other players are just optimizing at the first order as in the standard MG). Generalization to ever more complex optimizing set-ups, in which agents are aware of prior-level effects up to some finite recursive level, can in principle be iterated ad infinitum.

Actually, the game theory literature on first-entry games shows that the resulting equilibria depend on how agents learn [60]: With reinforcement learning, pure equilibria involve considerable coordination on asymmetric outcomes where some agents enter and some stay out; learning with stochastic fictitious plays leads to symmetric equilibria in which agents randomize over the entry decisions. There may even exist asymmetric mixed equilibria, where some agents adopt pure strategies while others play mixed strategies. We consider the situation where agents use a boundless recursion scheme to learn and optimize their strategy so that the equilibrium corresponds to the fully symmetric mixed strategies where agents randomize their choice at each time step with unbiased coin tosses. Consider a THMG game with  $N$  agents total,  $N_R$  of which employ such a fully random symmetric choice. The remaining  $N_S = N - N_R$  “special” agents (with  $N_R \gg N_S$ ) will all be one of three possible types: agents with  $S$  fixed strategies that choose their best (respectively worst, referred to above as anti-optimizing) performing strategy to make the decision at the next step and agents with a single fixed strategy. Our simulations confirm that these three types of agents indeed under-perform on average the optimal fully symmetric purely random mixed strategies of the  $N_R$  agents (see **Figure 26** of Chapter 8). Here, pure random strategies are obtained as optimal, given the fully rational fully in-

formed nature of the competing agents. The particular results are sensitive to which strategies are available to the special agents and to their proportion. Their underperformance in general requires averaging over all possible strategies and  $S$ -tuples of strategies. (In appendix Chapter A9, we show sample numerical results for  $N_S = 1$ ).

### 3.2 Parrondo Games

We now turn to an entirely different kind of game of the Parrondo type (PG). In the Parrondo effect (PE) [61], individually fair or losing games are alternated in succession either periodically or randomly by a single player to yield a winning game. That random alternation of losing games can yield an overall net gain seems especially counterintuitive. The PE was originally conceptualized as the game-theoretic equivalent of the “flashing ratchet” effect: A charged particle that executes symmetric Brownian motion in a ratchet-shaped potential drifts unidirectionally if the potential is flashed on and off either at random or periodically [62, 63]. It has been proposed as a potential explanation for aspects of random-walk diffusion [64], diffusion-mediated transport [65], spin systems [66], enzyme synthesis and gene recombination [67] and to be applied in investment strategies and portfolio optimization [68-70].

#### 3.2.1 Single-player capital-dependent Parrondo effect

The original Parrondo Effect (PE) combines two “capital-dependent” games which together constitute the PG. A single player has (discrete)-time-dependent capital  $X(t), t = 0, 1, 2, \dots$ . The time evolution of  $X(t)$  is determined by tossing biased coins. If game  $A$  is played, the player’s capital changes by  $+1$  (“win”) with probability  $p$  and by  $-1$  (“loss”) with probability  $1 - p$ . If game  $B$  is played, the changes are determined by:

	Prob. of win	Prob. of loss
$X(t)/3 \in \mathbb{Z}$	$p_1$	$1 - p_1$
$X(t)/3 \notin \mathbb{Z}$	$p_2$	$1 - p_2$

For  $p = \frac{1}{2} - \varepsilon$ ,  $p_1 = \frac{1}{10} - \varepsilon$  and  $p_2 = \frac{3}{4} - \varepsilon$  ( $\varepsilon > 0$ ). If either game  $A$  or game  $B$  is played exclusively, they both lose. In other words,  $\langle X(t) \rangle$  decreases with  $t$ . But if the games are alternated at random  $\langle X(t) \rangle$  increases. This is because the capital dependent parameter  $\text{Mod}[X(t), M]$  (here with  $M = 3$ ) can drive the system into a sufficient frequency of the winning component of game  $B$  (e.g.,  $B_2$ ) to cause the PE. Winning by playing losing games is only a seeming paradox as the possible values of  $X(t)$  when both games are played are not equiprobable. Instead, they take on values that, for a range of biases in the coins, are favorable to the player, given the peculiar rules of game  $B$ . (One may also devise probabilities such that both games are winning, yet the combined game is losing, and so on. A more general definition of the PE includes such “negative” effects as well. This is more fully explicated in appendix Chapter A9 and in ref. [71].)

### 3.2.2 Single-player history-dependent Parrondo effect

Ref. [72] extends the basic PE. Game  $A$  remains as described above (a simple biased coin toss). Game  $B$  is replaced with a history- as opposed to capital-dependent coin (game) defined by the respective winning/losing probabilities of four biased coins. A specific bias is associated with each of the four possible two-step binary histories (00,01,10,11) of the player's wins (1) or losses (0). The choice of coin follows the history dependent rule:

Before last $t-2$	Last $t-1$	History	Coin (Game) at $t$	Prob. of win at $t$	Prob. of loss at $t$
Loss	Loss	00	$B_1$	$q_1$	$1-q_1$
Loss	Win	01	$B_2$	$q_2$	$1-q_2$
Win	Loss	10	$B_3$	$q_3$	$1-q_3$
Win	Win	11	$B_4$	$q_4$	$1-q_4$

Both games  $A$  and history-dependent games of type  $B$  can be expressed as Markov transition matrices. But in this case  $X(t)$ , the evolution of the capital, is non-Markovian. To relate the capital to history one may therefore define the Markov chain

$$\bar{Y}(t) = \begin{pmatrix} X(t) - X(t-1) \\ X(t-1) - X(t-2) \end{pmatrix} \quad (4)$$

with the set of four states  $\{(-1,-1), (-1,+1), (+1,-1), (+1,+1)\}$  with associated conditional probabilities and probability state vector  $\{\pi_1(t), \pi_2(t), \pi_3(t), \pi_4(t)\} \equiv \bar{\pi}(t)$ . The transition matrix for game  $B$  is therefore

$$\mathbf{B} = \begin{pmatrix} 1-q_1 & 0 & 1-q_3 & 0 \\ q_1 & 0 & q_3 & 0 \\ 0 & 1-q_2 & 0 & 1-q_4 \\ 0 & q_2 & 0 & q_4 \end{pmatrix} \quad (5)$$

and  $\bar{\pi}(t+1) = \mathbf{B}\bar{\pi}(t)$ . The stationary state distribution  $\bar{\pi}_{st}$  obeys

$$\mathbf{B}\bar{\pi}_{st} = \bar{\pi}_{st} \quad (6)$$

with

$$\bar{\pi}_{st} = \frac{1}{N} \begin{pmatrix} (1-q_3)(1-q_4) \\ (1-q_4)q_1 \\ (1-q_4)q_1 \\ q_1q_2 \end{pmatrix} \quad (7)$$

( $N$  is a normalization factor; we assume a similar set of equations exists for  $\mathbf{A}$ , a simple coin-toss.)

As explained in ref. [72], though game  $B$  as a whole is losing, the values of  $\{q_1, q_2, q_3, q_4\}$  in  $\mathbf{B}$  are such that  $B_2$  and  $B_3$  are independently losing,  $B_1$  and  $B_4$  winning. Then even if game  $A$  is losing ( $p < 1-p$ ), it perturbs the losing cycles of  $B$ , as in the original Parrondo effect, but with the sensitive moments transferred from capital to the state recorded in memory. This can no longer be analogized to the flashing ratchet poten-

tial, but rather occurs by construction, such that for certain values of  $p$  and  $\{q_1, q_2, q_3, q_4\}$  the winning games in  $B$  dominate. This can occur when

$$\begin{aligned} (1 - q_4)(1 - q_3) &> q_1 q_2 \\ (2 - q_4 - p)(2 - q_3 - p) &< (q_1 + p)(q_2 + p) \end{aligned} \quad (8)$$

For example, in ref. [72],  $p = \frac{1}{2} - \varepsilon$  and  $\{q_1, q_2, q_3, q_4\} = \{\frac{9}{10} - \varepsilon, \frac{1}{4} - \varepsilon, \frac{1}{4} - \varepsilon, \frac{7}{10} - \varepsilon\}$ . Then the conditions of (8) are met when  $0 < \varepsilon < \frac{1}{168}$ .

Ref. [73] extends the history-dependent PE further by showing that it may arise when Game  $A$  is redefined to have the same history-dependent structure as (5). A more complex set of equations define the conditions under which two losing games of this kind, each with four coins, generate winning results under random alternation. From this perspective the simple coin toss form for game  $A$  in ref. [72] may be reformulated with a specific set of parameters that fall within the more general parameter space analyzed in ref. [73], viz.:

$$A = \begin{pmatrix} 1 - (\frac{1}{2} - \varepsilon) & 0 & 1 - (\frac{1}{2} - \varepsilon) & 0 \\ \frac{1}{2} - \varepsilon & 0 & \frac{1}{2} - \varepsilon & 0 \\ 0 & 1 - (\frac{1}{2} - \varepsilon) & 0 & 1 - (\frac{1}{2} - \varepsilon) \\ 0 & \frac{1}{2} - \varepsilon & 0 & \frac{1}{2} - \varepsilon \end{pmatrix} = \begin{pmatrix} \frac{1}{2} + \varepsilon & 0 & \frac{1}{2} + \varepsilon & 0 \\ \frac{1}{2} - \varepsilon & 0 & \frac{1}{2} - \varepsilon & 0 \\ 0 & \frac{1}{2} + \varepsilon & 0 & \frac{1}{2} + \varepsilon \\ 0 & \frac{1}{2} - \varepsilon & 0 & \frac{1}{2} - \varepsilon \end{pmatrix} \quad (9)$$

For  $\varepsilon = 0.005$ , we obtain the following game matrices:

$$\mathbf{A} = \begin{pmatrix} 0.505 & 0 & 0.505 & 0 \\ 0.495 & 0 & 0.495 & 0 \\ 0 & 0.505 & 0 & 0.505 \\ 0 & 0.495 & 0 & 0.495 \end{pmatrix}; \mathbf{B} = \begin{pmatrix} 0.105 & 0 & 0.755 & 0 \\ 0.895 & 0 & 0.245 & 0 \\ 0 & 0.755 & 0 & 0.305 \\ 0 & 0.245 & 0 & 0.695 \end{pmatrix} \quad (10)$$

i.e.,  $\bar{\pi}^{(A)} = \{0.495, 0.495, 0.495, 0.495\}$  and  $\bar{\pi}^{(B)} = \{0.895, 0.245, 0.245, 0.695\}$ .

Solving the eigenvalue equation (6) for  $\mathbf{B}$  and the equivalent for  $\mathbf{A}$ , we obtain the respective steady state probabilities for the two independent games:

$$\bar{\pi}_{st}^{(A)} = \begin{pmatrix} 0.255 \\ 0.250 \\ 0.250 \\ 0.245 \end{pmatrix}; \bar{\pi}_{st}^{(B)} = \begin{pmatrix} 0.231 \\ 0.274 \\ 0.274 \\ 0.220 \end{pmatrix} \quad (11)$$

and the respective independent probabilities for winning:

$$\begin{aligned} P_{win}(A) &= \bar{\pi}_{st}^{(A)} \cdot \bar{\pi}^{(A)} = 0.495 \\ P_{win}(B) &= \bar{\pi}_{st}^{(B)} \cdot \bar{\pi}^{(B)} = 0.494 \end{aligned} \quad (12)$$

Naively, one might presume that with a mixing ratio of 1:1, a random alternation of the games would yield a winning probability equal to the mean of their winning probabilities,

but this is not so, i.e.  $P_{win}(\frac{1}{2}A, \frac{1}{2}B) \neq \frac{1}{2}[P_{win}(A) + P_{win}(B)] = 0.4945$ . Instead, the winning probability is determined by the probabilities and steady-state vector of the mean of the transition matrices. As detailed more generally in appendix Chapter A9 and ref. [35], the winning probability of a combination of Markovian transition (game) matrices is not generally equal to the mean of their independent winning probabilities. Thus:

$$\bar{\pi}^{(\frac{1}{2}A, \frac{1}{2}B)} = \{\frac{1}{2}(p_1 + q_1), \frac{1}{2}(p_2 + q_2), \frac{1}{2}(p_3 + q_3), \frac{1}{2}(p_4 + q_4)\} = \{0.695, 0.370, 0.370, 0.595\} \quad (13)$$

so that

$$\frac{1}{2}(\mathbf{A} + \mathbf{B}) = \begin{pmatrix} 0.305 & 0 & 0.630 & 0 \\ 0.695 & 0 & 0.370 & 0 \\ 0 & 0.630 & 0 & 0.405 \\ 0 & 0.370 & 0 & 0.595 \end{pmatrix} \quad (14)$$

and

$$P_{win}(\frac{1}{2}A, \frac{1}{2}B) = \bar{\pi}_{st}^{(\frac{1}{2}A, \frac{1}{2}B)} \cdot \bar{\pi}^{(\frac{1}{2}A, \frac{1}{2}B)} = 0.501 \quad (15)$$

### 3.2.3 Multiple-player capital-dependent Parrondo effect and its reversal under optimization

Many variants of the PE have been studied, including capital-dependent multi-player PG (MPPG) [74, 75]: At (every) time-step  $t$ , a constant-size subset of all participants is randomly re-selected actually to play. All participants keep individual track of their own capital but do not alternate games independently based on it. Instead, this data is used to select which game all participants must use at  $t$ . The chosen game is the one which, given the individual values of the capital at  $t-1$  and the known matrices of the two games and their linear convex combination, has the most positive expected *aggregate* gain in capital, summed over all participants. This rule may be thought of as a static optimization procedure—static in the sense that the “optimal” choice appears to be known in advance. It appears exactly quantifiable because of access to each player’s individual history. If the game is chosen at random, the change in wealth averaged over all participants is significantly positive. But when the “optimization” rule is employed, the gain becomes a loss significantly greater than that of either game alone. The intended “optimization” scheme actually reverses the positive (collective) PE. The reversal arises in this way: the “optimization” rule causes the system to spend much more time playing one of the games, and individually, any one game is losing.

### 3.2.4 Single-player capital-dependent Parrondo effect and its reversal under optimization

Here, we present a more natural illustration of the illusion of control in Parrondo games: While the MG is intrinsically collective, PGs are not. Neither the capital- nor the history-dependent variations require a collective setting for the PE to appear. Thus, the illusion of control effect is most clearly demonstrated in a single-player implementation with two games under the most natural kind of optimization rule: at time  $t$ , the player plays whichever game has accumulated the most points (wealth) over a sliding window of  $\tau$  prior

time-steps from  $t-1$  to  $t-\tau$ . Under this rule, a “current reversal” (reversal of a positive PE) appears. By construction, the individual games  $A$  and  $B$  played individually are both losing; random alternation between them is winning (the PE effect), but unexpectedly, choosing the previously best-performing game yields losses slightly less than either  $A$  or  $B$  individually: the PE is almost entirely eliminated. Furthermore, if instead the previously *worst* performing game is chosen, the player does better than either game and even much better than the PE from random game choice.

Under the choose-best optimization rule, two matrices  $\mathbf{A}$  and  $\mathbf{B}$  do not form a linear convex sum. Instead, the combined game is represented by an  $8 \times 8$  transition matrix  $\mathbf{M}$  with conditional winning probabilities  $m_j$ :

$$\mathbf{M} = \begin{pmatrix} 1-m_1 & 0 & 0 & 0 & 1-m_5 & 0 & 0 & 0 \\ m_1 & 0 & 0 & 0 & m_5 & 0 & 0 & 0 \\ 0 & 1-m_2 & 0 & 0 & 0 & 1-m_6 & 0 & 0 \\ 0 & m_2 & 0 & 0 & 0 & m_6 & 0 & 0 \\ 0 & 0 & 1-m_3 & 0 & 0 & 0 & 1-m_7 & 0 \\ 0 & 0 & m_3 & 0 & 0 & 0 & m_7 & 0 \\ 0 & 0 & 0 & 1-m_4 & 0 & 0 & 0 & 1-m_8 \\ 0 & 0 & 0 & m_4 & 0 & 0 & 0 & m_8 \end{pmatrix}$$

$$m_j = \frac{1}{2} \left\{ \pi_{\alpha(j)}^{(A)} \left[ 1 + \pi_{\beta(j)}^{(A)} - \pi_{\beta(j)}^{(B)} \right] + \pi_{\alpha(j)}^{(B)} \left[ 1 - \pi_{\beta(j)}^{(A)} + \pi_{\beta(j)}^{(A)} \right] \right\}$$

$$j = 1, 2, \dots, 8 \quad (16)$$

The indices on the individual conditional probabilities for game  $A$  and  $B$ ,  $\pi_i^{(A)}, \pi_i^{(B)}; i = 1, 2, \dots, 4$  are converted to indices  $\alpha(j)$  and  $\beta(j)$  with  $j = 1, 2, \dots, 8$  by the following:

$$\alpha(j) = \text{Mod}[j-1, 4] + 1, \quad \beta[j] = \frac{1}{2}(j - \text{Mod}[j-1, 2] + 1) \quad (17)$$

For the “choose worst” rule, eqn. (16) is replaced by:

$$m_j = \frac{1}{2} \left\{ \pi_{\alpha(j)}^{(A)} \left[ 1 - \pi_{\beta(j)}^{(A)} + \pi_{\beta(j)}^{(B)} \right] + \pi_{\alpha(j)}^{(B)} \left[ 1 + \pi_{\beta(j)}^{(A)} - \pi_{\beta(j)}^{(A)} \right] \right\}$$

$$j = 1, 2, \dots, 8 \quad (18)$$

Alternated at random in equal proportion under the “choose best rule”,  $P_{win}^{best(A,B)} = 0.496$ , while if “choose worst” is used,  $P_{win}^{worst(A,B)} = 0.507$  (Compare to eqns. (12) and (15)). The mechanism for this illusion-of-control effect characterized by the reversing of the PE under optimization is not the same as for the MG, as there is no collective effect and thus no-crowding out of strategies or games. Instead, the PE results from a distortion of the steady-state equilibrium distributions  $\bar{\pi}_{st}^{(A)}$  and  $\bar{\pi}_{st}^{(B)}$  into a vector  $\bar{\pi}_{st}^{(\frac{1}{2}A, \frac{1}{2}B)}$  which is more co-linear to the conditional winning probability vector  $\bar{\pi}^{(\frac{1}{2}A, \frac{1}{2}B)}$  than in the case of each individual game (this is just a geometric restatement of the fact that the combined game is

winning). One can say that each game alternatively acts at random so as to better align these two vectors on average under the action of the other game. Choosing the previously best performing game amounts to removing this combined effect, while choosing the previously worst performing game tends to intensify it.

### **3.3 Conclusions**

In this chapter we have identified two classes of mechanisms operating in Minority games and in Parrondo games respectively in which optimizing agents obtain suboptimal outcomes compared with non-optimizing agents. These examples suggest a general definition: the “illusion of control” effect occurs when low-entropy strategies (i.e. which use more information) under-perform random strategies (with maximal entropy). The illusion of control effect is related to bounded rationality as well as limited information [76] since, as we have shown, unbounded rational agents learn to converge to the symmetric mixed fully random strategies (at least in the MG). In this setting, it is only in the presence of bounded rationality that agents can stick with an optimization scheme on a subset of strategies. Our robust message is that, under bounded rationality, the simple (large-entropy) strategies may be preferred over more complex elaborated (low-entropy) strategies. This is a message that should appeal to managers and practitioners, who are well-aware in their everyday practice that simple solutions are preferable to complex ones, in the presence of the ubiquitous uncertainty.

More examples should be easy to find. For instance, control algorithms, which employ optimal parameter estimation based on past observations, have been shown to generate broad power law distributions of fluctuations and of their corresponding corrections in the control process, suggesting that, in certain situations, uncertainty and risk may be amplified by optimal control (ANN; see chapter 10) [77]. In the same spirit, more quality control in code development often decreases the overall quality which itself spurs more quality control leading to a vicious circle [78]. In finance, there are many studies suggesting that most fund managers perform worse than random [79] and strong evidence that over-trading leads to anomalously large financial volatility [80] (again, see Chapter 10). Let us also mention the interesting experiments in which optimizing humans are found to perform worse than rats [81]. We conjecture that the illusion-of-control effect should be widespread in many strategic and optimization games and perhaps in many real life situations. The contribution here is to put this question at a quantitative level so that it can be studied rigorously to help formulate better strategies and tools for management.

However, it is obviously not the case that optimization in the face of limited information inevitably yields worse results. We now compare the findings in the MG to two other closely related games, the MAJG and \$G in both their standard and TH forms.

## Chapter 4. Control in the Majority Game and Dollar Game

In chapter 4, we survey and summarize descriptively the main results of seeking an illusion of control in the two other classes of market-modeling games, the MAJG and the \$G. We discuss the important differences, both theoretical and applied, between the standard MAJG and \$G and their “time-horizon” variants. We discuss the different results with respect to the illusion of control in the MG/THMG versus the MAJG/THMAJG//\$G/TH\$G: The illusion is not present in these games; agents successfully optimize and outperform their strategies.

### 4.1 Optimization in the MAJG and \$G: No illusion

The surprising success of “counteradaptive” agents in the MG and THMG is an important marker of the “illusory” nature of the standard and intuitive optimization rule (“choose best”). But it also raises the following question: May one correctly think of such an inverted rule as equivalent to these agents’ playing a Majority Game (MAJG) instead? It would seem on the face of it that as they are not optimizing to be in the minority, they must be optimizing to be in the majority but failing, and rather inadvertently succeed in finding the minority remarkably often. By this reasoning it seems to follow that in a game where all agents are striving to be in the majority (MAJG), select agents that optimize instead to be in the minority will likewise succeed disproportionately—implying that the MAJG should also demonstrate an “illusion of control”. After all, one might further anticipate, the MAJG with its heterogeneous agents demonstrates frustration and multiple equilibria; it undergoes a phase transition at a critical value of  $m$ , indeed it may be mapped onto a Hopfield net [82]. A similar argument could be made, perhaps, with respect to the \$G since players of this game are rewarded according to a time-lagged majority rule. (The formal distinctions among the three types of games are specified in the following section.)

The goal of the present chapter is to clarify these questions. We demonstrate that agents who invert their optimization rule in the MG are not actually playing a MAJG and that no illusion of control is found in either the MAJG or the \$G. We discuss our results in terms of persistent versus anti-persistent characteristics of the time series generated by the various models. In [83] and in chapters 5. and 6., we relate these comparative results to different characteristics of markets—or different phases they may enter—as real-world agents alter their game-playing behavior.

We first briefly review the formal structure of each of the MG, MAJG and \$G. We then present the results of extensive numerical simulations. Finally we discuss the differences that emerge in terms of persistent versus anti-persistent time-series. Formal mathematical methods may be found appendix Chapter A9

### 4.2 Minority versus Majority and \$ Games

#### 4.2.1 Recap of the “illusion of control” in MG and Time-Horizon MG (THMG)

The characteristic of the MG that distinguishes it from the MAJG and the \$G is the in-

stantaneous payoff of agent  $i$  in, i.e.,  $-a_i(t)A(t)$ —the minus sign encoding the minority rule (and similarly for each strategy for which it is added to the  $\tau-1$  previous payoffs. As in eqn. (1), the payoff may also be more simply  $-Sgn[a_i(t)A(t)]$ . This does not change the fundamental dynamics of the game nor has it any effect on the question we are here studying). As the name of the game indicates, if a strategy is in the minority so that  $a_i(t)A(t) < 0$ , it is rewarded. In other words, agents in the MG and THMG try to be anti-imitative.

The phenomenon discussed in [84] and in Chapter 3. is that when optimizing, and averaged over all actual agents and their component strategies in a given realization, and then averaged over many such initial quenched disorder states, agents in the TH variant of the MG underperform the mean of their own measured strategy performance and do so in all phases for reasonable lengths of  $\tau$  at all  $m$ . In the MG proper, the same statement holds true for “reasonable” run lengths post initialization but pre-equilibrium. It holds true post-equilibrium as well except for  $m$  at or very near  $\alpha_c$ , but in this region the number of steps to equilibrium is extremely large.

#### 4.2.2 Definition and overview of the Majority Game (MAJG)

Mathematically, the MAJG game differs from the MG (and a THMAJG from the THMG) only by a change in sign for the individual agent payoff function: i.e.,  $g_i^{min}(t) = -a_i(t)A(t)$  or  $g_i^{min}(t) = -Sgn[a_i(t)A(t)]$  whereas  $g_i(t) = +a_i(t)A(t)$  or  $g_i^{maj}(t) = +Sgn[a_i(t)A(t)]$ . In consequence of the plus sign, agents are rewarded when they select the alternative selected by the majority of agents at the same time. Thus, agents strive to be imitative rather than anti-imitative. From the perspective of markets, agents in the MG are “pessimistic” in assuming that resources are limited so that there can be only a minority of winners; they are “contrarian” in attempting to do what they believe most others are not doing. Agents in the MAJG are “optimistic” in assuming that resources are boundless, price (and value) potentially rising simply by virtue of collective agreement, so that the majority wins; they are “conformist” in attempting to do what they believe most others are also doing. Agents in both types of games “believe” that their actions may be optimized by examining the past paper-performance of their strategies.

As only a minority of agents win in the MG, mean agent gain  $\bar{G}^{min}(t) = \frac{1}{N} \sum_{i=1}^N g_i^{min}(t) < 0$ .

Cumulative wealth tends to decrease over time. In the MAJG, a majority of agents win so that mean agent gain  $\bar{G}^{maj}(t) = \frac{1}{N} \sum_{i=1}^N g_i(t) > 0$ . Cumulative wealth tends to increase over time.

In the MG, the time series  $A(t)$  (the sum of all agents’ actions) is typically anti-persistent, paralleling the anti-imitative behavior of individual agents. In the MAJG the time series  $A(t)$  is typically persistent, paralleling the imitative behavior of individual agents. (Persistence versus anti-persistence in these games will be formalized and discussed at length

in Chapter 5. In brief, to recapitulate, persistence is a measure of the tendency within a series for a pattern to repeat.)

### 4.2.3 Definition and overview of the Dollar Game (\$G)

The dollar game was introduced in [39] in order to capture more accurately the behavior of a great many typical traders in markets, while keeping a framework as close as possible to the initial MG set-up. Above we commented that agents in the MG are “contrarian” in attempting to do what most others do not. But this is not what true contrarian traders attempt: First, they attempt to be in the majority when the market is rising. Second, they likewise attempt to be in the minority when it is falling or when there is a turning point. And this is exactly what non-contrarian traders are also attempting. Indeed every trader attempts to do this. Contrarians differ from conformists in their *reasoning* as to what the market trend will be in the immediate future. They make predictions that typically differ from the majorities’ prediction—but they may or may not be correct. Like all others, they will still hope that, if correct, it will lead them to be in the majority in one instance and the minority in the other, as is appropriate according to the corresponding market phase. A similar correction to the description of “conformist” traders can be made.

Thus, an agent with greater “real world” behaviors is precisely one that rationally (but within the bounds of his reasoning capacity and data) alternates between choosing what he believes will be the minority state and choosing what he believes will be the majority state. Ideally, he wants to start choosing to try to be in the majority state at the first moment the market begins a rise following a decline—i.e., at a convex inflection point. Likewise, he ideally wants to start choosing to try to be in the minority state just as the market begins a decline following a rise—i.e., at a concave inflection point. A prominent trader and technical theoretician explains that traders such as himself try to “...identify a trend reversal at a relatively early stage and ride on that trend until the weight of the evidence shows...that the trend has reversed”[85]. This behavior may be most simply captured by the following alteration in the rule for individual agent gain:  $g_i^s(t) = +a_i(t-1)A(t)$  or  $g_i^s(t) = +Sgn[a_i(t-1)A(t)]$ . That is, the action at the previous time step  $t-1$ , interpreted as a judgment about whether  $A(t)$  will be  $>0$  or  $<0$ , determines whether an agent gains or loses. The mean agent gain retains the same form:  $\bar{G}^s(t) = \frac{1}{N} \sum_{i=1}^N g_i^s(t)$  and we anticipate that  $\bar{G}^s(t) > 0$  because in spite of the time-lagged  $a_i(t-1)$ , the payoff function is preceded by a  $+$  sign, so intuitively should generate largely imitative behavior. This intuition is confirmed by the numerical simulations presented in ref. [39]

### 4.3. Main results on the “illusion of control” in the THMG v. MAJG and \$G

In chapter 3., our main result with respect to the THMG was stated from the perspective of utility theory: “Throughout the space of parameters  $(N, m, S, \tau \ll \tau_{eq})$ , the mean payoff of agents’ strategies not only surpasses the mean payoff of supposedly-optimizing agents, but the respective cumulative distribution functions (CDF) of payoffs show a

first-order stochastic dominance of strategies over agents. Thus, were the option available to them, agents would behave in a risk-averse fashion (concave utility function) by switching randomly between strategies rather than optimizing. Agents are supposed to enhance their performance by choosing adaptively between their available strategies. In fact, the opposite is true (with certain caveats for the MG proper). Two conditions must hold for the statement to be false: (1)  $m \geq m_c$ ; (2) the system must be allowed to reach equilibrium. Condition (2) requires an exceedingly large number of preliminary steps before agent selection begins, and orders of magnitude more steps if  $m = m_c$ ). In the real-world, observing the paper underperformance of his strategies, a trader would abandon the process of adaptively selecting among them.”

However, in both the MAJG and the \$G, we find that the reverse is true: The optimization method greatly enhances agent performance, with strategies’ virtual mean performance consisting of relatively small gains and agents’ mean performance consisting of significantly greater gains. In the language used above: “Throughout the space of parameters  $(N, m, S, \tau \ll \tau_{eq})$ , the mean payoff of agents’ strategies (as calculated by each agent averaged over all strategies and agents in a realization) not only underperforms the mean payoff of optimizing agents (averaged over all given agents), but the respective cumulative distribution functions (CDF) of payoffs show a first-order stochastic dominance of agents over strategies. Thus, were the option available to agents to behave in a risk-averse fashion (concave utility function) by switching randomly between strategies rather than optimizing, they would rationally avoid such risk in favor of the optimization procedure. Agents are supposed to enhance their performance by choosing adaptively between their available strategies and they in fact do so.”

### 4.3.1 Quantitative statement and tests

#### 4.3.1.1 Analytic Calculation versus Numeric Simulation

In the THMG, the “illusion of control” effect is observed for all  $N, m, S$  and  $\tau \ll \tau_{eq}$ . We use the same Markov chain formalism for the THMG [41, 84] and extend it to both a THMAJG and a TH\$G to obtain theoretical prediction for the gains  $\Delta W_{Agent}$ , averaged over all agents and  $\Delta W_{Strategy}$  averaged over all strategies respectively, of agents and of all strategies in a given realization for each of the MG, MAJG and \$G. Compare to eqn. (2) and (3):

$$\langle \Delta W_{agent}^{game} \rangle = \pm \frac{1}{N} |\bar{A}_D| \cdot \bar{\mu} \quad (19)$$

$$\langle \Delta W_{strategy}^{game} \rangle = \frac{1}{2N} (\hat{\mathbf{s}}_{\mu} \cdot \bar{\mathbf{k}}) \cdot \bar{\mu} \quad (20)$$

In eqn.s (19) and (20), the superscript “game” identifies the game type with  $game \in \{M, MAJ, \$\}$ . In eqn. (19), the minus sign is needed for the MG; otherwise not. Recall from chapter 3. that  $\mu$  is a  $(m + \tau)$ -bit “path history” [10] (sequence of 1-bit states);  $\bar{\mu}$  is the normalized steady-state probability vector for the history-dependent  $(m + \tau) \times (m + \tau)$  transition matrix  $\hat{\mathbf{T}}$ , where a given element  $T_{\mu, \mu_{-1}}$  represents the transi-

tion probability that  $\mu_{t-1}$  will be followed by  $\mu_t$ ;  $\bar{A}_D$  is a  $2^{(m+\tau)}$ -element vector listing the particular sum of decided values of  $A(t)$  associated with each path-history;  $\hat{s}_\mu$  is the table of points accumulated by each strategy for each path-history;  $\bar{\kappa}$  is a  $2^{m+1}$ -element (=  $\dim[\text{RSS}]$ ) vector listing the total number of times each strategy is represented in the collection of  $N$  agents. As shown in appendix chapter A9,  $\hat{\mathbf{T}}$  may be derived from  $\bar{A}_D$ ,  $\hat{s}_\mu$  and  $\bar{N}_U$ , the number of undecided agents associated with each path history.

Agreement is excellent between numerical simulations and the analytical predictions (19) and (20) for all of the THMG, THMAJG and TH\$G. For instance, for  $\{m, S, N, \tau\} = \{2, 2, 31, 1\}$  and one identical quenched disorder state  $\hat{\Omega}$ , **Table 1** shows the payoff per time step averaged over time and over all agents and all strategies for both analytic and numerical methods. In this numerical example, the average payoff of individual agents is smaller than for strategies by  $-0.15$  units per time step in the THMG, but larger by  $+0.35$  units in the THMAJG and by  $+0.33$  units in the TH\$G. Thus, in this example, optimization appears to agents as genuine in the THMAJG and TH\$G but would seem illusory in the THMG.


**Table 1:** Numeric and Analytic Results for a single typical quenched initial disorder state in the THMG, THMAJG and TH\$G

Numeric	$\langle \Delta W_{agent}^{game} \rangle$	$\langle \Delta W_{strategy}^{game} \rangle$	Analytic	$\langle \Delta W_{agent}^{game} \rangle$	$\langle \Delta W_{strategy}^{game} \rangle$
MG	-0.21	-0.06	MG	-0.21	-0.06
MAJG	+0.43	+0.08	MAJG	+0.43	+0.08
\$G	+0.39	+0.06	\$G	+0.40	+0.06


The above results illustrate primarily the close alignment of analytic and numerical methods in generating results. Of greater interest is the comparison of agent versus strategy gains among the MG, MAJG and \$G at various values of  $m$  below, at and above  $m_c$ , and at various values of  $\tau$  both for  $\tau \ll \tau_{eq}$ , for  $\tau < \tau_{eq}$  and for  $\tau \geq \tau_{eq}$ —all averaged over a large ensemble of randomly selected quenched disorder states. The computational resources required to evaluate the analytic expressions grows for  $\langle \Delta W_{agent}^{game} \rangle$  as  $\propto 2^{m+\tau}$ . We therefore report only the numerical results.

#### 4.3.1.2 Illusory versus Genuine Control for $\tau \ll \tau_{eq}$


Almost all results that hold for multiple values of  $\tau \ll \tau_{eq}$  are illustrated for  $\tau = 1$ . In **Figure 3**, **Figure 4** and **Figure 5** we therefore first present graphic representations of the ensemble average of 50 runs comparable to **Table 1** but over many values of  $m$ .



**Figure 3:** Agent versus Strategy mean per-step gain in the THMG at various  $m$  with  $\tau = 1$ . The phase transition expected at  $m = 4$  is not detectable; strategies outperform agents at all  $m$  as indicated by the black squares .i.e., agent performance is always negative relative to strategy performance. The optimization procedure employed by agents yields worse performance than their component strategies on the basis of which agents select which strategies to deploy at each time-step.



**Figure 4:** Agent versus Strategy mean per-step gain in the THMAJG at various  $m$  with  $\tau = 1$ . Agent performance is always positive and greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies.




**Figure 5:** Agent versus strategy mean per-step gain in the TH\$G at various  $m$  with  $\tau = 1$ . Agent performance is always greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies, but the gain approaches zero asymptotically with growing  $m$ .


We see that the illusion of control in the THMG persists at all values of  $m$ . Incidentally we note that the phase transition to be expected at  $m = 4$  is strongly suppressed in the sense that the present metric is not sensitive to it. For both the THMAJG and the TH\$G, the control exerted by agents is non-illusory: Agents outperform their constituent strategies at all  $m$ . Because of the non time-lagged implementation of a majority rule in the THMAJG ( $g_i^{maj}(t) = +a_i(t)A(t)$ ), strategies show consistent positive gain, even if less than agents. Strategies' gain tends toward a positive limit with agents' gain tending toward a greater value at all  $m$ . However, in the TH\$G, strategies on their own, in the aggregate, tend toward zero gain with increasing  $m$ , as would be expected from a realistic model of a market. Agents are superior to strategies at all  $m$ , but converge to the zero limit of strategy gain with increasing  $m$ . In other words, of the three variations, the TH\$G with very short  $\tau$  shows the most satisfying convergence toward neither net positive nor net negative gain for both strategies and agents as strategy complexity increases and begins to approximate random selection. It is especially interesting that this is so, given that the \$G rule remains a majority one, albeit time-lagged by one step to take into account the time lag between decision and return realization [39].

#### 4.3.1.3 Illusory versus Genuine Control for $\tau < \tau_{eq}$


In **Figure 6**, **Figure 7** and **Figure 8** we present graphic representations of the ensemble average of 50 runs of the MG, MAJG and \$G comparable to **Table 1** but over many values of  $m$  and with  $\tau < \tau_{eq}$ , i.e. with a time window of “reasonable” size, but smaller than the equilibrium value (except for  $m = 2$ , where  $\tau_{eq} = 800$ ).



**Figure 6:** Agent versus Strategy mean gain per-step in the THMG at various  $m$  with  $\tau = 1000$ . The phase transition expected at  $m = 4$  is clearly visible; strategies outperform agents at all  $m$  as indicated by the black square: Agent performance is always negative relative to strategy performance. Even with a very long lookback of historical data, the optimization procedure employed by agents yields worse performance than their component strategies on the basis of which agents select which strategies to deploy at each time-step.



**Figure 7:** Agent versus strategy mean gain per-step in the THMAJG at various  $m$  with  $\tau = 1000$ . Agent performance is always positive and greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies.



**Figure 8:** Agent versus strategy mean gain per-step in the TH\$G at various  $m$  with  $\tau = 1000$ . Agent performance is always greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies.

We see once again that the illusion of control in the THMG persists at all values of  $m$  in spite of the “reasonable” length (1000 time steps) of  $\tau$ . Note, however, that the MG phase transition at  $m = 4$  is now visible—even though at  $m = 4$  the system is still far from equilibrium. (Recall that away from  $m_c$ , we have  $\tau_{eq} \approx 100 \times 2^m$ ; while for  $m$  near  $m_c$ , we have  $\tau_{eq} \gg 100 \times 2^m$ .) In the MG proper, where  $\tau$  grows without bound and agent and strategy performance begins to be measured only after  $\tau_{eq}$  steps, agent performance will exceed strategy performance—optimization succeeds—but only for  $m \geq m_c$ . But even for a relatively small number of agents (e.g., 31, as here), at  $m = 10$  say,  $\tau_{eq} \approx 100 \times 2^{11} > 200,000$  steps is unrealistically large (for a comparison with standard technical investment strategies used for financial investments). For both the THMAJG and the TH\$G, the control exerted by agents is again non-illusory: Agents outperform their constituent strategies at all  $m$ . Strategies in the THMAJG again show consistent positive gain, if less than agents. Strategies’ gain likewise tends toward a positive limit with agents’ gain tending toward a greater value at all  $m$ , just as for  $\tau = 1$ . Likewise in the TH\$G once more: Strategies on their own, in the aggregate, tend toward zero gain with increasing  $m$ , as would be expected from a realistic model of a market. Agents are superior to strategies at all  $m$ , but converge to the zero limit of strategy gain with increasing  $m$ . We may draw a similar

conclusion for  $\tau = 1000$  as for  $\tau = 1$ : The TH\$G with reasonable  $\tau$  shows the most satisfying convergence toward neither net positive nor net negative gain for both strategies and agents as strategy complexity increases and begins to approximate random selection, in line with what would be expected from the efficient market hypothesis [86, 87].

#### **4.4. Interpretations: crowding-out, anti-optimizing agents and persistence**

##### **4.4.1 Illusion of control and the crowding-out mechanism**

As discussed in chapter 3., illusion-of-control effects in the THMG result from the fact that a strategy that has performed well in the past becomes crowded out in the future due to the minority mechanism: Performing well in the recent past, there is a larger probability for a strategy to be chosen by an increasing number of agents, which inevitably leads to its failing. The crowding-out effect likewise explains why anti-optimizing agents overperform [88]: Choosing their worst strategy ensures that it will be the least used by other agents in the next time step, which makes it more probable that they will be in the minority.


By contrast, in all of the MAJG, \$G, THMAJG and TH\$G, with their variants of a majority mechanism for agent gain, a strategy that has performed well in the past is likely to do so again in the future. The domain of successful optimization encompasses all  $m$ , but diminishing as  $m$  increases and strategies become widely dispersed in strategy space, approximating ever more closely a collection of random decision makers. The optimization procedure is most effective in the \$G where the positive bias present even for strategies alone in the MAJG appears neutralized by the time-delay factor: On their own, strategies show effectively neither gain nor loss—unsurprising if the market is neither positively nor negatively biased, as it is in the MG (to lose wealth on average over time) and MAJG (to gain wealth on average over time). After all, one would expect that in an unbiased market, the net gain/loss averaged over all of a set of randomly chosen an fixed strategies should be zero. Gains in the \$G are therefore due solely to the optimization procedure. Given this, we predict that anti-optimizing agents should show no advantage over their optimizing counterparts in the MAJG and \$G and will rather underperform. The next section presents results of simulations testing this prediction.

##### **4.4.2. Illusion of control and “anti-optimizing” agents**

We again select 3 of 31 agents to function “counteradaptively” (“C agents”) and the remaining to function in the standard fashion (“S agents”). C-agents “anti-optimize”—at each time-step they deploy that strategy with the *fewest* virtual points accumulated over  $\tau$ , rather than the strategy with the most points as do S-agents. We display results obtained for a wide range of  $m$  both less than, and greater than  $m_c$ , for the THMG, THMAJG and TH\$G with  $\tau = 100$ .  $\tau$  is long enough so that the phase transition in the MG is not suppressed at  $m_c$  ( $= 4$  for  $N = 31$ )

**Figure 9** shows C-agent minus S-agent mean per-step change in wealth for  $2 < m < 14$ , each averaged over 100 runs of 100 days post- $\tau = 400$ . In the THMG, in the crowded regime, the illusion of control effect is so strong that C-agents significantly outperform S-agents. Because we know that, for all  $m$  at this  $\tau$ , agents underperform strategies, we see

that the opposite is true for C-agents: In the act of “anti-optimizing”, they actually optimize. However, as the phase transition approaches, this becomes less true. Indeed at  $m_c$  and after  $m_c$ —that is, in the non-crowded regime—S-agents outperform C-agents, converging to zero difference with increasing  $m$ . We know however that these S-agents for large  $m$  are nonetheless underperforming their strategies. Thus, while the illusion of control effect remains present, it is not strong enough for C-agents to outperform S-agents in this regime.



**Figure 9:** Difference between C-agent and S-agent mean per-step change in wealth for 3 of 31 C-agents, averaged over 100 days and 200 runs with  $\tau=400$  in the THMG, THMAJG and TH\$G.

By contrast with the results for the THMG, C-agents in the THMAJG and TH\$G consistently underperform S-agents as predicted from the success of the optimization scheme at all  $m$  (again converging to zero difference at large  $m$ ). The size of this underperformance for anti-optimizing agents is consistent with the large degree of standard optimization success as shown in **Figure 7** and **Figure 8**.

## 4.5 Conclusions

The “illusion of control” is an unfortunate confounding effect that appears in many situations of “bounded rationality” where optimization occurs with limited (inadequate) information. However ubiquitous the illusion may be, it is not universal. In ref. [84] and chapter 3., we note that the illusion of control effect in the THMG is fundamentally due to three ingredients: (i) the minority mechanism (an agent or a strategy gains when in the minority and loses otherwise); (ii) the selection of strategies by many agents because they were previously in the minority, hence less likely to be so in the present; and (iii) the crowding of strategies (i.e., few strategies for many agents). In the following analysis of persistence, we will see that there is a close relationship among these three characteristics, a high degree of anti-persistence in the resulting time-series and the illusion of control. On the other hand, genuine control is more likely to be present when the underlying mechanism employed by agents is not of the minority type (as in the MAJG and \$G) and the resulting time-series is therefore more likely to be persistent. In another paper [83] and in chapters 5. and 6. we extend this analysis to the types of Hamiltonian cycles on graphs found associated with persistent and anti-persistent series, and employ these

methods for generating predictors of empirically generated time-series, both in models and in the real world.

## Chapter 5. Persistence and anti-persistence in the MG, MAJG and \$G

In Chapter 5. we first present the differences among the main time-series  $L^{game} \equiv 2Sgn\{A^{game}(0), A^{game}(1), \dots, A^{game}(t_{max}-1), A^{game}(t_{max})\} - 1$  generated by the three agent-based games in terms of persistence/anti-persistence. “Persistence”  $\mathfrak{P}$ ,  $0 \leq \mathfrak{P} \leq 1$ , is a quantitative measure of the tendency of patterns in a time-series to be followed by repetitions of that same pattern. Similarly, “anti-persistence” is a measure of the tendency of patterns to be followed rather by the pattern with the last binary digit reversed. As presented in full detail in ref. [89], we show that for the two games employing the majority rule (MAJG, \$G),  $L^{maj}, L^s$  is persistent ( $\mathfrak{P} > 0.5$ ), approaching the random limit asymptotically with increasing memory length  $m$  (i.e.,  $\lim_{m \rightarrow \infty} \mathfrak{P} = 0.5$ ); in the MG  $L^{min}$  crosses from anti-persistent for  $m < m_c$  ( $\mathfrak{P} < 0.5$ ) to persistent ( $\mathfrak{P} > 0.5$ ) at the well-known phase transition, attains a maximum and then declines asymptotically to ( $\mathfrak{P} = 0.5$ ) with further increasing  $m$ .

Second, we then present a new use of a cycle decomposition method that expresses the inherently probabilistic nature of a Markov chain as an exact superposition of deterministic sequences, extending ideas discussed in [90]. We demonstrate how the  $L^{game}$  generated by the THMG, THMAJG and TH\$G (all of which are Markovian processes) may be exactly represented as a weighted superposition of deterministic cycles on graphs;  $L^{game}$  generated by the MG, MAJG and \$G may be approximated by an empirically-derived form of the decomposition.

In other words, all three types of game-generated binary time-series as well as real world binary series may be reformulated as perturbations of a characteristic underlying dynamic that to the zero<sup>th</sup> order is wholly deterministic. I.e., any binary time series may be decomposed into a superposition of wholly deterministic Hamiltonian cycles on graphs. The cycle of greatest weight may be considered the dominant underlying determinism; cycles of lesser weight may be thought of as the higher-order perturbations. Probabilistic transitions from one history to another are in this view recast as probabilistic transitions among cycles. This cycle decomposition approach parallels the theories of dynamics systems and of deterministic chaos [91] in particular on the one hand and of quantum chaos [92] on the other hand, both based on the decomposition on unstable periodic orbits.

The representation of a series in deterministic cycles is related in complex fashion to the degree of “persistence” or of “anti-persistence” it displays. We discuss how the cycle decomposition of a time-series allows it to be understood as a perturbed deterministic process.

Third, we discuss how a decomposition of the respective series into such cycles on graphs reveals in highly intuitive fashion characteristic differences among the three types of games. These differences in cycle structure are consistent with, but further differentiate, the distinction between persistent and anti-persistent series.

We finally apply the cycle decomposition method to the prediction of game-generated time-series, based on a sliding window of past information.

### **5.1 Persistence versus Anti-persistence in the THMG, THMAJG and TH\$G**

As discussed in [84] and chapter 3., in the MG and THMG, the degree to which agents underperform their own strategies varies with the phase as parameterized by  $\alpha$ . As noted in [41], in the crowded phase ( $\alpha < \alpha_c, m < m_c$ , i.e. few available strategies relative to  $N$ ), the “crowd” of agents choosing an action at any given time-step acts like a single “super-agent”; the remaining agents as a (non-synchronized) “anti-crowd” whose actions will conform to the minority choice. Thus, for  $\alpha < \alpha_c, m < m_c$ , when a strategy is used, it is probably used by more than one agent, often by many agents. When many becomes enough, it becomes a losing strategy with large probability—precisely because so many agents “think” it’s the best choice and use it. This implies that at the next time step, agents will not use it. The time-series of determined choices  $\bar{A}_D$  therefore does not show trends (or persistence), but rather anti-persistence—here we use the term imprecisely and impressionistically, meaning “alternating”. Formally, anti-persistence is more complex and subtle than mere “alternating” (which, we will see, is formal anti-persistence at a scale of 1). Anti-persistence is scale-dependent and not equivalent to “random”.

Consider a binary time-series with an  $m$ -bit  $\mu(t)$  defined in the same way as we have in the MG, MAJG and \$G:  $\mu(t)$  is a sliding window of 1-bit states each of length  $m$ , sliding with a time step of one unit:  $s(t-m+1), \dots, s(t)$ . A *perfectly* anti-persistent binary series at scale  $m = 2$ , for example, is characterized as follows: Select any one instance of the four possible  $\mu(t) \in \{00, 01, 10, 11\}$ . Identify the following bit  $s(t+1) \in \{0, 1\}$ . Now identify the next instance of the selected  $\mu(t)$ . If the series is perfectly anti-persistent, the following bit will *always* be 1 if the previous following bit was 0, and 0 if the previous following bit was 1. (If  $m = 1$ , it follows that a perfectly anti-persistent series is 010101... or 101010..., i.e., alternating.) A perfectly anti-persistent series can be generated by two lookup tables indicating what bit follows which  $\mu(t)$ . Whatever bit is indicated by the first table, the opposite bit is indicated by the second. Whenever an entry in a table is used for a given  $\mu(t)$ , the other table is used when  $\mu(t)$  occurs again [93]. (These tables happen to be identical to strategy pairs at the maximum Hamming distance in the MG). No matter which of the  $2^{m+1} = 16$  possible strategies is used for the first table, and regardless of which of the  $2^{m+1} = 4$  possible  $\mu(t)$  are used to initiate it, the time series generated by these tables will rapidly settle into perfect anti-persistence.

The “persistence”  $\mathfrak{P}$  of a given series at scale  $m_s$  is thus simply the proportion of persistent such following bits, counting every instance of each of the  $2^m = 2^{m_s-1}$  possible histories. Its “anti-persistence”  $\tilde{\mathfrak{P}} = 1 - \mathfrak{P}$ . Another way of stating the same thing is that given

$m_s$ , the persistence  $\mathfrak{P}$  of a series is the proportion of times that histories of length  $m_s + 1$  end in bits 00 or 11; anti-persistence  $\tilde{\mathfrak{P}}$  is the proportion they end in 01 or 10.


The process generating a given empirical series may be unknown. This unknown process may itself be a memory-related process such as in the games we are discussing; it need not be (it could be, for example, completely random). The process may likewise be Markovian and memory-related as are the time-series generated by the TH games; it may be memory-related but non-Markovian as the non-TH version of these games. If the process is memory-related, whether Markovian or not, we need to distinguish between the unknown length  $m$  (or  $m + \tau$ ) underlying the process and a length we denote as  $m_s$  indicating the scale of our analysis. Intuitively, it would seem that choosing  $m_s = m$  or  $m_s = m + \tau$ , would yield the most robust analysis of persistence versus anti-persistence. But if the memory length of the process is unknown, this cannot be done. In the case of a TH game, all paths of length  $m + \tau$  transition to other paths of equal length with known probabilities as these games are Markovian. The scale  $m + \tau$  would seem even more natural since all quantities can be determined exactly using analytic methods, at least in principle. (See [94, 95] for illuminating studies on how to determine the optimal coarse-grained scale in simple cellular automata.)

However, for  $m$  or  $\tau$  large, the transition matrices become intractably large as well, scaling as  $(m + \tau)^2$ . We thus need to know whether the degree of persistence/antipersistence may be approximated at a lower effective  $m_s$ : I.e., given a binary time-series generated by an unknown process, may we usefully characterize its persistence by a small value of  $m_s$  to replace its “actual”  $m$  or  $m + \tau$ ? Before analyzing the degree of persistence and anti-persistence in the MG, MAJG and \$G, we first show that, in fact, relatively small values of  $m_s$  do successfully characterize persistence.

We implement an algorithm to characterize the persistence of a binary time series as described above. We find sharp differences in the degree of persistence between the time series generated by the MG on the one hand and the time series generated by the MAJG and \$G on the other. A less sharp distinction also emerges between the MAJG and the \$G. We find as well that characteristic distinctions arise at all reasonable  $m$ , attenuating as  $m$  grows large, and at all reasonable scales. This last point is important: While the degree of persistence is obviously best captured for the TH variants of these games as the natural scale  $m + \tau$  (since the TH games are Markovian), it is not obvious that a small- $m$  scale will effectively capture distinctions in the MG, MAJG and \$G proper as the natural scale is large and unbounded. It emerges that in general, if a significant degree of persistence or antipersistence is characteristic at a large- $m$  scale, it may be approximated by a low- $m$  analysis. We demonstrate this, and the differential characteristics of the respective time series in the following.


**Figure 10** illustrates graphically the mean degree of persistence or anti-persistence averaged over 25 different initializations identically shared by each of the THMG, THMAJG and TH\$G with  $N = 31$ ,  $S = 2$ ,  $\tau = 100$  and  $m \in \{2, 3, \dots, 10\}$ , scale  $m_s \in \{2, 3, \dots, 10\}$ . The generally darker shade in all squares of the grid representing the THMG implies values closer to or consistent with anti-persistence ( $< 0.5$ , with 0.5 representing the equal degree

of persistence and anti-persistence of a random sequence), those in the THMAJG and TH\$G with persistence ( $>0.5$ ). (We note for future investigation the telling fact that in the THMG,  $\mathfrak{P} \approx 0.5$  roughly uniformly, for *all* values of  $m_s > m$ . This is evidenced by the roughly uniform gray shading in the upper right region above the diagonal in the leftmost graphic in **Figure 10**. By contrast, in the MAJG,  $\mathfrak{P} \approx 1.0$ , again roughly uniformly, for all values of  $m_s < m$ , as shown in the lower-left region below the diagonal in the middle graphic. Both are consistent with the idea that  $m$  is the natural scale at which to measure persistence.)



**Figure 10:** Persistence (white)/Anti-Persistence (black) at various scales and memory lengths in the MG, MAJG and \$G for  $\{N, S\} = \{31, 2\}$ . The grey scale between 0 and 1 given to the right of the checkboards encodes the degree of persistence at the chosen scale (abscissa) for different  $m$  values (ordinate), calculated as described in the text, using game-generated binary histories of length 1000 over 100 different runs for each game type

The fraction of persistent sequences up cells of increasing  $m$  are roughly similar on a relative basis up the columns of different scales, but shifted toward the random, especially for the THMAJG and TH\$G. The fraction of persistent sequences up cells of increasing  $m$  in the THMG shows a shifting transition point. This feature is seen most sharply along the upward-and-to-the-right diagonal which represents the relation  $scale=m$ . **Figure 11** charts the degree of persistence along this diagonal for all three games.



**Figure 11:** Persistence and anti-persistence in MG, MAJG and \$G for  $m = m'$ , 100 runs of 1000 time steps for each game, at each  $m' = m$  Error bars show 1 SD (barely visible at this scale for the MAJG).

At the phase transition ( $m_c = 4$  for  $N = 31$ ), the time-series generated by the THMG when the scale equals  $m$  undergoes a transition from anti-persistence to persistence and then declines asymptotically to the random limit 0.5. When  $m_s \neq m$  this transition occurs at either smaller or larger values of  $m$ . Both the THMAJG and TH\$G generate persistent time-series exclusively. The degree of persistence declines monotonically to the random limit 0.5 with increasing  $m$ . Persistence in the THMAJG is always greater than in the TH\$G.

**Figure 11** also illustrates that the variability in persistence differs from game to game and for the MG by phase. Consistent with the fact that for  $m < m_c$  the time-series itself is highly variable from run to run and much less so for  $m \geq m_c$ , the variability in persistence from run to run likewise changes at the phase transition. Consistent with its high degree of persistence, there is little variability altogether for the MAJG. The variability in  $\mathfrak{P}$  for the \$G is roughly consistent throughout and much greater than for the MAJG.


We now compare the MG, MAJG and \$G in terms of a decomposition of the binary time-series they generate—i.e.,  $1/2\{\text{Sgn}[A(t)] + 1\}$ —into Hamiltonian cycles on graphs (at some given length  $m_s$ ).

## 5.2 Cycle structure of a binary series

To obtain a cycle decomposition of a binary series, we first encode a general binary series as its decimal equivalent  $+1$  given length  $m_s$  (a rolling window). For example, take some arbitrary binary series  $\{0, 1, 1, 1, 0, 0, 1, 0, \dots\}$ . Then at  $m_s = 2$ :


$$\begin{aligned} \{0,1,1,1,0,0,1,0,\dots\} &\rightarrow \{(0,1),(1,1),(1,1),(1,0),(0,0),(0,1),(1,0)\dots\} \\ &\rightarrow \{1,3,3,2,0,1,2,\dots\} + 1 = \{2,4,4,2,1,2,3,\dots\} \end{aligned} \tag{21}$$

All of the allowed transitions from one  $m_s = 2$ -bit state to the next form a complete binary de Bruijn graph of order 2 as shown in the middle graphic of **Figure 12**. (Examples of complete de Bruijn graphs of order  $m_s = 1$  and  $m = 3$  are shown to the left and right respectively.)



**Figure 12:** Complete binary de Bruijn graphs of orders 1, 2 and 3 from left to right. The vertices may be numbered as shown or by their decimal equivalents +1. In a complete graph, all possible states and transitions are represented.

If a binary sequence thus encoded touches no vertex more than once except upon returning to the first, the sequence is considered a cycle. **Figure 13** shows the cycle consisting of the last four digits in (21). The four digits are represented as three vertices (and the “edges” connecting them) because the last digit (vertex) repeats the first.



**Figure 13:** de Bruijn graph for  $m_s = 2$  showing only the cycle consisting of the last four states of (21)

In appendix Chapter A9, we demonstrate in detail that (and how) any sequence may be decomposed into a weighted superposition of such cycles, unique for each  $m_s$ . Each cycle represents a different deterministic binary-state process, implicit in the fact that no state in such a cycle can be reached by, nor may it transition to, more than one other state.

In the most general non-cyclic binary process, every  $m_s$ -bit state may be reached from two preceding states and may transition to two. This follows from the fact that all transitions may be considered as the motion of a sliding window along the binary string: One digit is dropped at the beginning, one is added at the end. Given a state at time  $t$ , the state at time  $t+1$  can end in only two possible states: the one ending in 0 or the one ending in 1. The converse is true as well reading the string in reverse order, hence any given state can be reached by two preceding ones. In ternary series, each state may transition to one of three possible states and may likewise be reached by three, and so on for higher bases.

Suppose a binary series represents transitions from one state to the next that are probabilistic. When a general process is decomposed into deterministic cycles, the transition probabilities between states are recast as transition probabilities to different cycles from those states common to more than one cycle. (In what follows, we take for granted the cycle decomposition and apply it; readers who prefer to understand the details first may turn to appendix Chapter A9).

### 5.3 Cycle distinctions among MG, MAJG and \$G

Figure 14 illustrates the typical weighting of cycles in the MG, MAJG and \$G for  $m = m_c = 4, m_s = 2$ , relative to the expected weighting in a completely random binary sequence. Bars in white represent cycle weights in excess of the random expectation value normalized to 1; bars in gray represent cycle weights less than 1.




Figure 14: Relative weighting of cycles in the MG, MAJG and \$G at  $m_s = 2$  for  $m = m_c = 4, N = 31, S = 2$  expressed as a fraction/multiple of the expected weighting (normalized to 1) of a cycle decomposition of a completely random binary sequence


There are many potential way of both characterizing and applying such a “taxonomy” of series. One simple metric is the pseudo-Euclidian distance between the cycle decomposition of a given series and that for a completely random series, providing a measure of non-randomness. Taken exactly, such a measure presumes that cycles in a decomposition are orthonormal and they are not. But by normalizing the relative weights  $w_j \in [0,1]$  (with all weights 0 for a random series), we get a somewhat better if still imprecise definition of “distance”. Presuming orthogonality, this distance  $d$  is simply:

$$d_{game} = \sqrt{\sum_{j=1}^k w_j^2} \quad (22)$$

As is evident from the absolute values of the distance from zero of each of the bars in **Figure 14**,  $d_{maj} = 0.270 > d_s = 0.133 > d_{min} = 0.028$ , i.e., the MAJG generates series that are, by this measure, the least random; the MG generates series that are closest to random (at  $m_c$ . A more exact use of this metric would require characterizing a series at many, theoretically all,  $m_s$ .) The same relationship among the three games are illustrated in **Figure 13** with respect to persistence and anti-persistence: While the MG is anti-persistent and the MAJG and \$Gs are persistent, the MG is closest to  $\mathfrak{P} = 0.5$ , the random limit, the MAJG farthest.

In **Figure 12** we present a side-by-side comparison of the raw weights for the non-TH MG, MAJG and \$G at  $2 \leq m \leq 8$ , decomposed at  $m_s = 2$ . The phase transition in the MG is evident in the change in the cycle structure at  $m = 4$ . It is important to note that these are ensemble averages over many different initial quenched disorder states  $\hat{\Omega}$ . The cycle distribution for any single  $\hat{\Omega}$  usually differs substantially from any other and typically departs widely from that for a random series. Thus, for the MG, the closeness of the distribution to the random distribution, especially for small  $m$ , represents not the typical distribution but the many widely varying distributions around a mean.

Even though the cycle distribution for the MG changes significantly as  $m$  varies, at all  $m$  the MG structure is different at a glance from the MAJG and \$G structures. Likewise, the \$G structure, with its disproportionately smallest representation of longer cycles is distinguishable from the MAJG structure at a glance.



**Figure 15:** Relative cycle weights given  $\{m, S, N\} = \{2, 2, 31\}$  for MG, MAJG and \$G from left to right and  $2 < m < 8$  from top to bottom. Within each graph, the numbers 1-6 on the x-axis represents the six cycles at  $m_s = 2$ , i.e.,  $(1, 2, 3, 4, 5, 6) \equiv (\{1, 1\}, \{4, 4\}, \{2, 3, 2\}, \{1, 2, 3, 1\}, \{2, 4, 3, 2\}, \{1, 2, 4, 3, 1\})$ . The cycles may likewise be converted into the outcome + or - following a sequence of directional changes, i.e., if the last number in the cycle label is odd, the height of the bar represents the probability relative to a random series that - follows the series of directional changes; if even that a + follows.

## 5.4 Cycle analysis of the illusion of control

Ref. [84] discusses another feature of MG's and THMG's: Even though agents' choice of strategy is designed to optimize their performance, the performance of agents (their mean per-step accumulation of wealth) is always on average poorer than the measured performance of their  $S$  strategies. Agents choose which strategy to use based on a comparison of how their strategies perform: The optimization rule employed is simply "deploy whichever strategy would have accumulated the largest number of points over the time horizon  $\tau$ ". Yet by this measure agents appear to do better by selecting a single strategy and sticking to it, or by randomly selecting strategies (or actions) at each time-step. Thus the control method in the MG and THMG yields what may be called an "illusion of control". The relative outperformance of strategies as thus defined can be seen as arising from the anti-persistence characteristic of MG-generated time-series. As discussed in ref. [89], MAJG, THMAJG, \$G and TH\$G display no illusion of control and the time-series they generate are persistent at all  $m$ .

Another way of understanding this phenomenon is by examining the performance of agents and strategies around deterministic cycles in the THMG. Depending on the weights, a single cycle may dominate the behavior of the THMG. In other instances just a few cycles dominate. It turns out that for a majority of  $\hat{\Omega}$ , most cycles show an illusion of control under the MG optimization rule for most strategy pairs. Therefore, the illusion is expected on average over many different initial configurations of  $\hat{\Omega}$ . As cycles are deterministic, we can examine the behavior of the optimization rule in particularly simple form. Table 2 illustrates the behavior of the MG optimization rule around all possible (allowed) cycles as averaged over every possible strategy and every agent (i.e., every possible pair of strategies) for  $\{m, S, N, \tau\} = \{2, 2, 31, 1\}$ . (This is the first occasion where [almost] all possible cycles based on paths of length 3 are listed. We take this opportunity to note that the two simplest cycles, (1,1) and (8,8), and only these two, are absent. This is a typical feature of the THMG and demonstrates immediately the tendency of its binary time-series toward anti-persistence, at least at the smallest possible  $m_s = 1$ .)

**Table 2:** Optimization of All Standard Strategy Pairs over All Cycles and time steps

Cycle	$\langle \Delta W_{agent, all\ pairs} \rangle$	$\langle \Delta W / \Delta t_{agent, all\ pairs} \rangle$	$\langle \Delta W_{strats, all} \rangle$	$\langle \Delta W / \Delta t_{strats, all} \rangle$
(3,6,3)	0.	0.	0.	0.
(2,3,5,2)	0.	0.	0.	0.
(4,7,6,4)	0.	0.	0.	0.
(1,2,3,5,1)	-0.5	-0.125	0.	0.
(2,4,7,5,2)	0.	0.	0.	0.
(4,8,7,6,4)	-0.5	-0.125	0.	0.
(1,2,4,7,5,1)	-0.5	-0.1	0.	0.
(2,4,8,7,5,2)	-0.5	-0.1	0.	0.
(1,2,4,8,7,5,1)	-1.0	-0.1667	0.	0.
(2,3,6,4,7,5,2)	0.	0.	0.	0.
(2,4,7,6,3,5,2)	0.	0.	0.	0.
(1,2,3,6,4,7,5,1)	-0	-0.0714	0.	0.
(1,2,4,7,6,3,5,1)	-0.5	-0.0714	0.	0.
(2,3,6,4,8,7,5,2)	-0.5	-0.0714	0.	0.
(2,4,8,7,6,3,5,2)	-0.5	-0.0741	0.	0.
(1,2,3,6,4,8,7,5,1)	-1.0	-0.125	0.	0.

(1,2,4,8,7,6,3,5,1)	-1.0	-0.125	0.	0.
<i>Mean, all cycles</i>		<b>-0.07682</b>	<b>0.</b>	<b>0.</b>

As noted before, from eqns. (27)-(31) (found in appendix chapter A9) it is straightforward to calculate the number of points gained or lost around cycles. We see that the per-step performance of strategies averaged over all possible strategies is neutral around every cycle ( $\Delta W = 0 = \Delta W / \Delta t$ ), while agent performance is  $\sim -0.08$  averaged over all cycles. (The unit step change in wealth must be adjusted for the differing number of steps in different cycles. This has been done in columns 3 and 5 of **Table 2** )

Note, too, that the change in wealth for strategies when averaged over all possible cycles and all possible strategies is zero For almost all given quenched disorder matrices  $\hat{\Omega}$  the distribution of strategies is asymmetrical, and the asymmetry of this distribution, in conjunction with the minority rule for winning, ensures an average loss (even if under most circumstances less of a loss than for agents [12]). In other words, a “crowd” of strategies in one region of strategy space insures that their choices will on average be the majority, hence losing, decisions. By averaging over all possible strategies, we eliminate the expected asymmetry in strategy distribution.

Similar patterns are found for all but a few exceptional values of  $\tau$  where mean agent performance is also neutral. (That this must be so follows from the fact that the mean performance of all agents in a given  $\hat{\Omega}$  is at best 0.) Thus, over many different  $\hat{\Omega}$ , the optimization rule of the THMG degrades mean performance relative to the measured performance of underlying strategies. In ref. [84], the inclusion of agents who select their worst-performing strategy is discussed. It turns out that these agents outperform not only standard agents, they outperform their underlying strategies and can even regularly attain net positive gain which for a standard agents in a MG structure is exceedingly rare. The performance of such agents around cycles have the same values as standard agents but with the opposite sign.

### **5.5 Cycle-based predictors for the MG, MAJG and \$G**

What **Figure** illustrates at a glance is that the overall departure from randomness of the MAJG and \$G is greater than for the MG and that the MAJG departs somewhat less from randomness than the \$G. On this basis we expect, and in fact find, that a prediction method based on a cycle decomposition should yield the best results for the MAJG and \$G relative to the MG.

The predictor itself consists of a table regenerated from all  $m_s=2$  cycles for every possible 250-day (-step) sliding window in sequence along a binary series. The cycles are rank-ordered by weight (frequency). The last state prior to the prediction day consists of 1 plus the last  $m$  days of binary data converted to decimal form. The predictor simply consists of the next state cycle following the present state. The prediction is the first number (1 or 0) in the binary representation of the state. If the present state is represented in more than one cycle, the cycle with the larger weight is used. (The rare persisting ties are settled by a fair coin toss.)

We start using days 1-250 predicting day 251 and then slide the 250 day window forward by one day and recalculate the decomposition, so that days 2-251 predict day 252. The

process is repeated through all days of data (minus the initial 250) and the percentage of correct predictions is calculated.

In Chapter 6, we create a toy predictor and apply it to real world financial time series. all three games, pretending that the sequences they return are actual market price-changes. We find that for the MG at  $m_s=2$ , the percent of correct predictions using a cycle-decomposition predictor is  $\sim 65\%$ ; for the MAJG the percent is  $\sim 72\%$  and the \$G  $\sim 66\%$ . Thus the predictability of the three game types by this test stand in the same relationship to one another as  $d_{game}$  and  $\mathfrak{P}_{game}$ .

## 5.6 Approximation of a Series as a Perturbed Dominant Cycle

As discussed in ref.'s [88, 89], the THMG, THMAJ and TH\$G may all be expressed analytically in the form of Markovian transition matrices. If these matrices are then decomposed into a weighted superposition of deterministic cycles, then the most heavily weighted cycle may be considered a zero<sup>th</sup> order approximation for the entire matrix. (In the limit of small  $m$  and  $\tau$ , it may often happen that more than one cycle has the same, largest weight; these may be linked or disjoint.) As detailed in ref.s [43, 88, 89] and below, the transition matrices for these games arises out of  $N_U(t)$ , the expression for the number of agents at timestep  $t$  whose contribution to the collective state  $A(t)$  needs to be determined by a coin toss (because their  $S$  strategies have accumulated the same number of points but would yield at least two differing predictions), and  $A_D(t)$ , the expression for the contribution to the value of  $A(t)$  that is wholly determined, contingent only upon the particular quenched disorder initializing the game and the particular “path history” at  $t$ , i.e.,  $\mu(t)$ . A path-history is simply the union of the  $m$ -bit binary history at  $t$  (which we denote as  $\mu_t$ ) and the preceding  $\tau$ -bit rolling window over which agent and strategy scores are maintained.

An agent’s contribution to  $A(t)$  will be fully determined (non-stochastic) when all of its strategies make the same prediction given  $\mu(t)$ . (This is more likely the smaller the value of  $S$ , i.e., the fewer the number of strategies per agent). This will happen only for some  $\mu(t)$  if an agent’s strategies differ, but of course for all  $\mu(t)$  if they happen to be identical.

Now,  $A(t) = A_D(t) + A_U(t)$ , where  $A_U(t)$  is the sum of all agent’s contributions determined by an unbiased coin-toss. For certain  $\mu(t)$  the absolute value of  $A_D(t)$  may be large enough (a sufficient proportion of agents contributing to it) so that  $A_U(t) < A_D(t)$ —even if all the remaining agents (the number of which =  $N_U(t)$ ) happen to vote the same way by chance and therefore generate the maximum possible  $A_U(t)$ . In these instances,  $Sgn[A(t)] = Sgn[A_D(t)]$ , i.e., the contribution to the series  $\mu(t) \rightarrow \mu(t+1)$  will be determined and that step in the series therefore deterministic. In

terms of a transition matrix for a binary series, the entries representing the two transition probabilities from a given state to each of two possible successor states will consist of 1 and 0, implying that the transitions are either present or absent with certainty (i.e., wholly determined). In fact, as we will detail shortly, the dominant cycle in a decomposition of a Markovian transition matrix arises directly from the expression for  $A_D(t)$ , excluding the term  $N_U(t)$ . Its weight is therefore a direct measure of the degree of determinism present in the series.

A transition matrix presumes that a series is in fact Markovian. Series generated by the THMG, THMAJG and TH\$G truly are. But the series generated by their non-TH variants, the MG, MAJG and \$G proper are effectively Markovian only at equilibrium, and with intractably large matrices—the window of past information grows without bound and becomes equivalent to a sliding window only when very remote information no longer has an effect. However, as discussed in ref. [89], a very high-dimensional Markovian series may be approximated by a series of much lower dimension, capturing at least some of the information of the full matrix. Likewise may non-Markovian process be approximated, in a fashion similar to the use of hidden Markov models in which a more complex process is approximated by a hidden switching between two different Markovian processes [96, 97]. (Indeed, the cycle decomposition method may be considered a simplification of the hidden Markov method in which the “switching” occurs among multiple deterministic matrices instead of between two probabilistic ones.)

The information contained in such an approximate transition matrix may be captured by analysis of a large window of preceding history. An efficient way of doing so is to create a cycle decomposition from this empirical data. If cycles emerge that are not merely somewhat but significantly weightier than all the others, one may hypothesize a significant degree of hidden determinism in the series as evidenced, for instance by the creation of predictors that successfully employ the weightiest element(s) of the decomposition.

Note, too, that in converting a continuously-valued (but time-discrete) series to a binary series, we are severely compressing the available data. Success or failure of a cycle-decomposition predictor in the teeth of real-world transaction costs provides a useful heuristic for the degree of data-preservation.

## Chapter 6. Cycle- and persistence-based real-world (toy) predictors

In Chapter 6, we apply the cycle decomposition method and analysis of persistence to real-life financial time series: Different series (or the same series in different periods) may be characterized by a signature cycle structure and/or degree of persistence. This fact leads to prediction methods based on cycles and on persistence, even for time-series which are not in fact Markovian.

### 6.1 Cycle-based predictor for the NASDAQ composite


Ultimately, the agent-based games we are studying are meant to illuminate the behavior of real-world financial markets and the time-series they generate. An important question that has been addressed on occasion is whether and how the light shed by studies of these games may be translated into methods applicable to, for example, the prediction of real-world time-series. For example, ref. [98] identifies “pockets of predictability” in the MG and shows how the reasoning as to the circumstances under which they occur may be transferred to a prediction method for the NASDAQ composite index (IXIC). Ref. [99] provides an example of prediction by “tuning” an ensemble of MGs to a real-world time-series. Ref. [100] uses the tuning method of [99] to predict the Shanghai stock market, but with an ensemble of “mix-games”, i.e., in each game in the ensemble, some agents play a MG and the rest a MAJG.

We likewise provide an example of how insights derived from these games may indeed be translated into real-world prediction methods. We do so by constructing a predictor based on a decomposition of the time-series into weighted deterministic cycles on graphs and by quantifying the degree of persistence or anti-persistence in our target series. For pedagogical purposes we here demonstrate only “toy” predictors of great simplicity. However, these methods may be readily generalized and improved, opening up new avenues for research and application, for instance, by encoding price-change histories in ternary or even more detailed form. The cycle decomposition method for such higher-order matrices is significantly more complex, however.

Following [98] to make possible comparison, we likewise examine the NASDAQ composite (IXIC) over its entire history. We then incorporate a persistence filter with improved results. We also examine the index divided in fourths over its history to analyze the different performance characteristics of the predictor and relate it to persistence during different time periods. We compare results of the cycle decomposition predictor with and without transaction costs both to a simple buy and hold strategy and to a Martingale strategy (MGL). The latter consists of going long on day  $t$  if the sign of the price change on day  $t-1$  is “+” and going short on day  $t$  if the sign of the price change on day  $t-1$  is “-”.

We first examine the structure of the series as a whole. We find that at  $m_s = 2$ , the overall degree of persistence is  $\sim 0.52$ , thus somewhat persistent like the series from the MAJG and \$G and contrary to the MG. Its relative-weighted cycle decomposition (again over its entire history) is shown in **Figure 12**. (The cycle decomposition for each fourth of the series is presented later.) The cycle with weight closest to zero is the last, Hamiltonian

cycle  $\{1,2,4,3,1\}$ . The one-step persistent cycles  $\{1,1\}$  and  $\{4,4\}$  are over-represented relative to a random series and the one-step anti-persistent cycle  $\{2,3,2\}$  is under-represented. We thus have a snapshot that suggests a somewhat predictable because somewhat persistent time series. (The perfectly anti-persistent one-step cycle, the last, is present with the lowest weight.) With  $d_{IXIC} = 0.118$ , the IXIC departs from randomness by this cycle measure to an extent midway between that of the MG and the \$G.



**Figure 12:** Cycle decomposition of the entire binary daily history of the NASDAQ composite index (IXIC) from February, 1971 through January, 2008. (No error bars as this is a single—the only actual-instantiation.)

From the over-representation in the IXIC of the two one-step persistent cycles one may guess informally at a simple strategy that has in fact been widely employed (by so-called “momentum traders”) to trade the IXIC, namely the one-step MGL: If the previous day’s change in price is positive, buy (go or stay long); if negative, sell (go or stay short). Over the entire nearly 10,000 day history of the IXIC this strategy yields 58% correct directional guesses (which needs to be compared to a net upward drift over the history of 56% up days). This small degree of excess predictability, however, yields an annualized return of 38% per year (exclusive of transaction costs) versus a “buy and hold” return of 9% per year.

However, at 50 basis points per “round-trip” transaction (i.e.,  $50/10000 = .005$  cost per unit value traded), the MGL strategy would have yielded an annualized return of ~7%, more than 22% less than the buy and hold return of ~9%. A significant amount of information has nonetheless been detected by the MGL strategy, but not enough to be of practical use—the predictor is highly inefficient, obtaining large gains at the cost of many changes in strategy direction. In the end, insufficient information has been recovered from the binary-compressed price-change series to compensate for the cost of applying that information.

An explicit hidden (switching) Markov model of tick data used to generate a predictor for foreign exchange rates [101] yielded similar results: While significant theoretical predictability was obtained for the USD/CHF exchange rate, it was insufficient to overcome transaction costs, a common problem that plagues financial market predictors.

We do better by building a more precise (if still “toy”) predictor based on the same cycle-decomposition for the IXIC index (NASDAQ composite; **Figure 12**), taking into account


the entire cycle structure (in a simple way to be described) and again obtain a correct prediction percentage of 58% versus 56% up. The excess predictability in this case yields an annualized return of 36% versus the “buy and hold” return of 9%—slightly worse than the MGL predictor. However, returns remain superior to “buy and hold” for “round-trip” transaction costs of up to 50 basis points. Thus, the cycle predictor even in “toy” form extracts information from the compressed binary time-series more efficiently than does the MGL predictor—enough arguably to warrant potential real world application.

Like most real-world series, the IXIC goes through periods of varying  $\mathfrak{P}$  (as measured over some time-scale). For example, we expect—and find—that during periods of low volatility as measured by proxy using the VIX (Chicago Board of Trade Volatility Index for the NASDAQ 100), the IXIC shows  $\mathfrak{P} > 0.5$ , while during periods of high volatility,  $\mathfrak{P} < 0.5$  (i.e., anti-persistence as here defined measures the probability of frequent reversals of direction). Furthermore, if very recent past periods (50 trading days) of anti-persistence ( $\mathfrak{P} < 0.45$ ) are excluded (no market exposure on such days, either long or short), the returns (annualized over days of market exposure) rises to 38% (equivalent to the MG predictor) and remain superior to buy and hold for up to 60 basis points per round trip. In addition, since one is exposed to the market for only ~8% of trading days using this measure, monies are freed up for other uses the rest of the time.

## **6.2 Cycle predictors for other time-series**


Similar results are obtained for the US Dollar/Japanese Yen foreign exchange rate characterized by  $\mathfrak{P} \approx 0.53$ . However, if we apply these same methods to a real-world series with  $\mathfrak{P} \approx 0.50$ , for example, the Philadelphia Exchange Gold and Silver Index (XAU), we find that they all fail.

It is worth comparing the cycle decomposition and its associated predictor to a simpler statistical analysis of dependencies and a comparable predictor we develop from these dependencies. In [102], Zhang studies the history of directional price changes in the NYSE Composite Index (NYA) of 400 stocks from 1966 to 1996. He displays the frequency that an up (+) daily price change occurs following each of the eight possible three-day sequences of directions of price changes. In **Figure 13** we duplicate for the IXIC his analysis of the NYA. (The figures are remarkably similar.)



**Figure 13:** Frequency with which a positive daily change in price occurs following each of the eight possible three-day sequences of the direction of price-change for the IXIC. White circles indicate the relative weighting of the sequence. “+” = up; “-” = down. The small number of days of no change are excluded

In **Figure 14** we transform the results of **Figure 13** into expectations relative to a random sequence with the same overall upward bias as the NYA itself, to make it comparable to the cycle decomposition of **Figure 12**. **Figure 14** makes evident the rationale for the Martingale predictor: Regardless of the preceding two states, if the last state is “+” the next one is more likely to be as well; similarly with “-”.



**Figure 14:** Relative probabilities of a positive daily price change following each of the 8 possible daily sequences of direction of price changes over three preceding days


To improve the comparison—and highlight differences—the cycles of **Figure 12** may be recast according as follows. The cycle  $\{4, 4\}$  is equivalent to  $11 \rightarrow 11 = 111 \equiv +++$ , i.e., “++” followed by “+” with frequency (probability) of  $0.50 + 0.03 = 0.53$ . The 0.03

comes from the second bar of **Figure 12**. The cycle  $\{1,1\}$  is equivalent to  $00 \rightarrow 00 = 000 \equiv ---$ , i.e., “--” followed by “-” with frequency (probability) of  $0.50 + 0.08 = 0.58$  (again from **Figure 12**) which is the same as followed by “+” with frequency (probability) of  $1 - (0.50 + 0.08) = 0.42$ . All the cycles of **Figure 12** may be recast according to **Table 1**. Note that not every possible sequence of states for any length other than 2 is included and that the selection of states based on the cycle decomposition includes states of four different lengths.

**Table 3:** Probability of + following state sequences equivalent to cycles

Cycle	Prior States	Calculation	P(+)
$\{1,1\}$	--	$1 - (0.50 + 0.08) =$	0.42
$\{4,4\}$	++	$0.50 + 0.03 =$	0.53
$\{2,3,2\}$	-+-	$0.50 - 0.06 =$	0.44
$\{1,2,3,1\}$	--+-	$1 - (0.50 - 0.04) =$	0.46
$\{2,4,3,2\}$	-++-	$0.50 - 0.03 =$	0.47
$\{1,2,4,3,1\}$	--++-	$0.50 + 0.51 =$	0.51

The results of **Table 3** are displayed in **Figure 15**.



**Figure 15:** Cycle decomposition of the IXIC recast in terms of frequencies that a given sequence of daily directions of price changes is followed by a positive change in price.

The departures from 0.50 in **Figure 15** show mean dependencies over the entire history of the IXIC that are noticeably smaller than the departures and dependencies in **Figure 13**. At first glance therefore one might expect that a predictor based on the cycle decomposition would be less powerful than a comparable one based simply on the eight possible three-day histories. (The six sequences in the transformed cycle decomposition have a


mean length of 3.3.) But the cycle decomposition arguably reveals more complex dependencies than the simpler statistical analysis.

In fact, using the eight three-day histories we construct a “dependency predictor” as close as possible in structure to the cycle-decomposition predictor and obtain the following results: Not taking into account transaction costs, the dependency predictor yields annualized returns of 26% — better than buy and hold (consistent with the arguments against the efficient market hypothesis made in [102]), but not so good as either the Martingale or cycle-decomposition predictor. (The superior performance of the Martingale predictor vis-à-vis the dependency predictor suggests that an approach based strictly on these dependencies may attain maximum performance at least for this application by utilizing histories of no longer than a single day.) Furthermore, while the cycle decomposition predictor retains performance superior to buy and hold (and to Martingale) for up to 50 basis points of transaction costs, dependency predictor gains—as well as effectively the entire initial investment—are wholly wiped out by such costs with an annualized loss of -19% per year. The dependency predictor retains performance superior to buy and hold only for transaction costs of no more than 17 basis points.

It is valuable to examine the performance of the predictor during a set of different arbitrarily selected time-periods. Following the procedure in [102] as applied to the NYA, we divide the IXIC price-change series in equal fourths and apply the cycle predictor to each. Results are shown in **Table 4** and in **Figure 16**.

**Table 4:** Summary of cycle predictor results for IXIC history divided into equal sequential fourths

Fourth	$\mathfrak{P}$	$\langle \mathfrak{P} \rangle$	Ann. Ret. Cyc.	Ann. Ret. BH	Ann. Ret. Cyc. – BH	$d$
1st	0.56	0.55	57%	2%	55%	0.16
2nd	0.52	0.54	40%	9%	31%	0.13
3rd	0.52	0.52	49%	23%	26%	0.11
4th	0.49	0.49	-7%	-6%	-1%	0.09



**Figure 16:** Annualized return of cycle predictor, by fourth of IXIC, in excess of buy and hold return showing correlation with  $d$  and  $\langle \mathfrak{P} \rangle$ .  $d$  is the pseudo-Euclidian distance between the cycle decomposition of the series and the expected decomposition of a random series with the same overall linear bias.  $\langle \mathfrak{P} \rangle$  is the mean persistence of the series, i.e., averaged over  $1 < m_s < 10$ .

For each fourth we determine the annualized raw return (exclusive of transaction costs) from the cycle predictor and the annualized buy and hold return. We compute the difference and denote the associated  $\mathfrak{P}$ ,  $\langle \mathfrak{P} \rangle$  (averaged over  $1 < m_s < 10$ ) and  $d$ . We see that the raw return from the cycle predictor in excess of the buy and hold return declines monotonically over time with monotonic declines in  $\mathfrak{P}$ ,  $\langle \mathfrak{P} \rangle$  and  $d$ . These results are consistent with the improved performance of the cycle predictor found previously with a persistence filter.

## Chapter 7. Supplementary findings in the MG

In Chapter 7, we illustrate certain additional findings and subtleties in the MG corollary to the main findings already presented, focusing on genuine and illusory control. First, we present additional findings for the MG/THMG in which a subset of agents choose their worst strategy. (Appendix Chapter A9 provides the mathematical details for modifying TH game expressions to account for any proportion of such “counteradaptive” or “anti-optimizing” agents in the TH –MG, -MAJG and -\$G.) Second, we present more extensive numerical simulations for the illusion of control in the MG/THMG for a wide range of  $\tau$ .

### 7.1 Choosing the worst strategy in the THMG/MG

In previous chapters we have discussed the fact that in the MG/THMG, if a small subset of agents invert the fundamental optimizing rule by always deploying their previously worst-performing strategy, these “anti-optimizing” (or “counteradaptive”, i.e. “C”) agents outperform standard (“optimizing”, “adaptive”, i.e. “S”) agents.

The MG literature is replete with specific phenomena that arise from the fact that by definition, the majority of agents on average must lose in the long run: Inter alia, these references discuss MG macrostate (and therefore mean agent) performance as a function of the critical MG parameters [4, 5, 59]; of evolutionary variations of the MG rules [103-105]; of agent-heterogeneity [106, 107] and of the limits on information available to agents [107-109]. Different forms of agent contrarian behavior have likewise been explored [107, 109]. But the primary fact of agent underperformance vis-à-vis their own strategies itself has remained only indirectly examined; nor has the experiment been done that follows logically once this observation is made—what is the effect of strictly reversing not agents’ choice per se (“contrarian” behavior in a naïve sense), but the basis on which they make choices.

Framing the experiment this way probes a fundamental presupposition of the MG—that choosing the best method of responding based on history is rational (if boundedly, which might therefore mean in fact mean “irrationally”. In other words, as in real life, one takes for granted that the most rational course of action is to continue doing that which has worked best in the past; irrational to deliberately do the reverse. The “illusion of control”, when present, illustrates how the opposite may prove true.) That the MG optimization procedure generates, on average (and under most circumstances, if not all), *worse* results for agents than for their strategies might lead a cynic to conclude that the fundamental MG rule for adaptation is instead *maladaptive* for agents, even if it is adaptive collectively (i.e., leads to cooperativity among agents).

The unexpected fact that a subset of agents selecting their worst strategy consistently outperform the rest by a very wide margin—and even generate net positive gain in a setting (minority wins) where mean loss is inevitable—is consistent with a related phenomenon explored in ref. [106, 110]: Therein, some or all agents choose among their strategies probabilistically as determined by a standard partition function of a temperature-like parameter. In effect, a degree of biased randomness is introduced into the optimization process. At any time-step a subset of agents will by chance choose their worst strategy. At each time-step the subset size varies about a mean as determined by the parti-

tion function, and in composition varies as a flat distribution among all agents allowed to select probabilistically. This has an overall salutary effect in lowering  $\sigma$ , as in [111]. By contrast we here examine a systematic, complete and non-probabilistic inversion of the optimization procedure for a fixed subset of agents. We focus primarily on the comparative performance characteristics of individual agents, secondarily on systemic effects.

There are also discussions in the literature of so-called “contrarian” agents [109], but these are agents who always vote contrary to their best strategy, not in accord with their worst. (This means that when all strategies dictate the same action, the agent chooses the opposite action). The distinction is not trivial: As a group, truly counteradaptive agents not only outperform standard agents as well as their own strategies, as noted, they often achieve net positive gain. By contrast, on average, “contrarians” achieve a maximum average winning probability of 0.485 [109]—still  $< 0.5$ —whereas truly counteradaptive agents attain a winning probability  $> 0.5$ .


### 7.1.1 Sample Results from the MG and THMG

For a single instantiation both numerically and analytically with  $N = 31$ ,  $S = 2$  and  $m = 2$  at  $\tau = 1$ , **Table 5** demonstrates how the improved (often net positive) gain among the C agents raises the overall performance of the agent pool. So long as the proportion of C agents is small enough the game as a whole behaves like standard minority game as evidenced by the still-negative gain of agents averaged over all agents.

**Table 5:** Numerical/Analytic Results of THMG with and without 3 C Agents 28 S Agents (= left value/right value) for  $N = 31$ ,  $S = 2$ ,  $m = 2$ ,  $\tau = 1$

	$\langle \Delta W_{Agent} \rangle$	$\langle \Delta W_{Strategy} \rangle$
<i>With 3 C-agents of 31</i>	-0.14/-0.14	-0.05/-0.05
<i>Without any C-agents</i>	-0.26/-0.26	-0.05/-0.05


Figure 17 shows an ensemble result over 100 runs. For both Table 5 and Figure 17, three agents are selected first to contribute as S agents and then, in an identical  $\hat{\Omega}_{31}$  as S agents. For the ensemble results of Figure 17, an identical sequence of pre-selected random numbers is used for each pair (i.e., of C-containing and all-S) of runs, at each instance where the net vote requires a coin-toss.




**Figure 17:** Mean agent (blue) versus strategy (black) performance over 100 runs for  $N = 31$ ,  $S = 2$ ,  $m = 2$  at  $\tau = 1000$ , both with (solid lines) and without (broken lines) 3 C-agents. Re-simulations use identical parameters and where required (and possible) identical random numbers.

The single example in **Table 5** is selected to be typical: That is, the three agents that are run once as S agents and then again as C agents are at the modal Hamming distance for the ensemble of **Figure 17**.

Both the typical and ensemble average for these agents—when containing C agents—not only outperform both their own strategies and the other S agents on average, they all generate net positive gain, approximately inverting the relation between Hamming distance between strategies and gain. We cannot tell from these findings whether the changed agents achieved actual positive gain. **Figure 18** demonstrates the individual agent and strategy performances for the 3 C agents in the single instantiation of a  $\hat{\Omega}_{31}$  of **Table 5**. **Figure 19** demonstrates the individual agent and strategy performances for the 3 C agents averaged over the 100 instantiations of  $\hat{\Omega}_{31}$  in **Figure 17**.



**Figure 18:** 3 of 31 agents use counteradaptive (“C”, choose worst) strategy selection, shown in red. Their previous underperformance when using standard selection rule (“S”, choose best) is shown in blue. The chart shows results for three agents from a single typical quenched disorder tensor  $\hat{\Omega}_{31}$ . Note the approximate inversion of the relation between  $d_H$  and  $\Delta W_{agent}$ . This individual example from among an ensemble of 100 differing  $\hat{\Omega}_{31}$  as typical: Its three C agents are at the mode of the three different Hamming distances in the ensemble.



**Figure 19:** Ensemble averages for 3 of 31 agents using counteradaptive (“C”) strategy selection are shown in red. Their previous underperformance when using standard selection rule (“S”, choose best) is shown in blue. The chart shows results for three agents from 100 randomly-selected quenched disorder tensors  $\hat{\Omega}_{31}$ . Note the approximate inversion of the relation between  $d_H$  (values in white) and  $\Delta W_{agent}$ .

The *hypothetical* outperformance of the “unused” relative to the “used” strategies in the MG was first observed in [44]. “Unused” strategies may be translated as “worst” for  $N = 2$  and as “not best” for  $N > 2$ . The results we discuss here are the same for  $N > 2$ : The worst strategy under these conditions performs best, the second worst performs second best, etc.

Ref. [56] discusses a MG in which the performance of the second-best strategy for  $S > 2$  is considered in agents’ decision as to which strategy to employ. The *global* performance is thereby enhanced. These results are arguably consistent with but not identical to what we discuss here: The principle that appears to emerge from [56] is that additional information—i.e., ranked strategy performance—enhances agent performance. But in fact, this conclusion is an artifact—an illusion—that arises from the simpler phenomenon demonstrated here.

### 7.1.2 Microscopic examination of agent performance with and without C agents in the THMG

As noted, that C agents may attain net positive gain on average is not immediately evident from the information presented so far: C agent performance that is merely better than S agent performance but still negative will also improve agent performance averaged over all agents. (From the actual values in **Figure 17**, however, one may deduce that the contribution of the three C agents must be on average positive in order to shift the performance curve upward as much as it does.) More explicitly, **Figure 20** presents the mean of many simulations in which positive net gain for C agents emerges clearly: The mean performance of these three agents only is shown as a solid red line, the mean for the same three agents’ performance when standard (S) is shown as a broken red line. In fact, in the symmetric phase ( $\alpha_c > \alpha = \frac{P}{N} = \frac{2^m}{N}$ ) counteradaptive agents end to be absolute winners. (The phase transition at  $\alpha_c$  is preserved when the number of counteradaptive agents is small as is the characteristic relation of the global efficiency,  $\frac{\sigma^2}{N}$ , to  $P$ .)

However, the effect of potentially differing histories with and without counteradaptive agents must be addressed. It seems likely that thus altering three of 31 agents must lead to significantly divergent histories: The sum of all agents’ actions will differ in the simulation pairs. With different histories, the timing and number of tie-breaks will also differ. Perhaps the outperformance of the counteradaptive agents arises in some way from a different history; and different histories would render meaningless the attempt to impose a fixed sequence of random numbers to break ties.

The divergences in histories can be large. But it turns out that there are surprisingly many instances when histories do not diverge at all within simulation “pairs” (i.e. between a given  $\hat{\Omega}_N$  run without C agents and the same  $\hat{\Omega}_N$  run with C agents). This is because (as we will address formally via the analytic formulation of the THMG) the counteradaptive vote of any given agent will differ from its standard vote for only certain “path histories”  $\mu(t)$  of length  $m + \tau$ , and not for all. Which  $\mu(t)$  are associated with a potentially changed vote depends upon the specific pair of strategies constituting the agent (and, of course, the summed actions of all other). Thus, for any given  $\mu(t)$  there will almost al-

ways be fewer changed individual “determined” votes (votes involving no tie break) than the total number of counteradaptive agents, and often none at all. (Hence a requirement that the proportion of C agents be “small”.) As detailed in the mathematical appendix which builds on [43], because this set of values is independent of agent action, the number of determined votes associated with each  $\mu(t)$  is the same in both simulations of a pair. The particular proportion of 1 and 0 votes contributing to the majority and minority determination may differ: But whether it can differ for a given simulation pair depends upon whether enough undetermined votes remain—votes that require a coin-toss—to alter the collective decision should all vote the same way; and whether it actually does differ if it can, depends upon the actual vote. (Note that if  $g_i(t) = -a_i(t)A(t)$  rather than  $g_i(t) = -\text{Sgn}[a_i(t)A(t)]$  as in our discussion, this principle still holds since each new (last) digit of the history  $\mu_i$  (and path-history  $\mu(t)$ ) is converted from  $A(t)$  to binary form, i.e.,  $D(t) = 1/2 \{-\text{Sgn}[A(t)] + 1\}$ )

If the proportion of counteradaptive agents is small enough, the total number of potentially changed individual determined votes between pairs in a simulation is therefore yet smaller. The sum of all changes alters the net voting imbalance but often not so much as to alter the minority position. Often the overall proportion of determined votes is so large, and the altered vote imbalance so small, that it will be impossible for the minority side to differ. When a difference is possible, it is most often unlikely: Even less likely if the randomly chosen value of the undetermined vote is made identical within pairs. Note, too, that if altering one or more agents changes  $\hat{\Omega}_N$ , the history is likely to be very different because the history is generated by  $\hat{\Omega}_N$ . This is what occurs when an agent or agents are given “optimal” strategies [57]; likewise if an agent always votes opposite to what it otherwise would, as this is equivalent to replacing each of a strategy-tuple with its most Hamming-distant “opposite.” Changing only agents’ method of strategy selection alone—even for all agents—has no effect on  $\hat{\Omega}_N$ ; changes in the history grow more likely therefore with the proportion of counteradaptive agents.

In both simulations the total number of undetermined voters following each possible history remains the same. In simulations using  $\{N, S, m\} = \{31, 2, 2\}$ , the proportion of undetermined votes is about 0.2. Over different  $\hat{\Omega}_N$ , and within simulation pairs, the average difference in the imbalance caused by the “determined” voters is dependent upon M, S and N and the distribution of the total of all choices requiring a coin-flip is binomial around zero. The difference varies from small to large and inversely with the range of variation in the total vote (i.e.,  $\propto \frac{1}{\sigma}$  for  $\alpha < \alpha_c$ ; see **Figure 22** and discussion): A large  $\sigma$  implies large imbalances, which in turn can arise within a given run consistently only by a large imbalance in the determined vote; and these are less likely to be reversed by the votes of a small number of undetermined voters—and may be impossible to reverse. (N.b., this analysis provides an impression of how the TH games in general may be treated as deterministic systems perturbed symmetrically about sequential determined states. Where the size of the perturbation is non-zero but small enough,  $A(t)$  will differ

from  $A_D(t)$  but not from  $2D(t) - 1$ . In these instances, perfectly deterministic sequences of states will arise, i.e., pockets of predictability as discussed in [98]. In this reference, it is the MG rather than the THMG that is at issue, but in many instances the equilibrium (unbounded  $\tau$ ) sequence of agent strategy choice, hence the equilibrium sequence of states  $D(t)$ , is little different than the choices and sequence of states for bounded  $\tau$ .

Therefore the possibility of divergent histories is reduced, sometimes to zero, when the undetermined agents are forced to make the same random choice in both simulations. Hence the distribution of divergences over many different initial states is not Gaussian: Divergences tend to persist; the absence of divergences likewise, especially, of course for those path-histories (possibly all) which are wholly determined. (See [98] for a discussion of other implications of determined portions of the overall history.) The presence of many instances when the histories within simulation pairs are identical both requires and makes meaningful the imposition of identical random number sequences in comparing them. We may then ask whether the outperformance of the counteradaptive agents is found in such instances as well.


For example, using a fixed sequence of random numbers in 100 typical simulations of 100 time-steps for  $N = 31$ , with and without 4 C agents, the minority group reversed in  $\sim 30\%$  of all time-steps averaged across the 100 simulation pairs. But in fully 20 of the 100 hundred runs all 100 time-step histories were identical within pairs (implying, as mentioned, that divergences tend to persist in those runs wherein they arise). The example in **Table 5** happens to show no divergence in the two runs in a pair, hence the strategy scores are identical.

Averaging over just the simulation pairs with identical histories, the agents selected to be counteradaptive did worse in that simulation where they were kept standard than in the simulation where they actually became counter adaptive; and even slightly worse than the agents kept standard in both runs. (And these did about the same in both). But when the selected agents became counteradaptive in fact, they significantly outperformed the standard agents (in both runs as well as themselves in the standard run).

Over all 100 runs the average outperformance of counteradaptive agents versus the standard ones was  $\sim 0.04$  wealth units per time-step. In just the 20 identical runs, their outperformance averaged  $\sim 0.05$  wealth units per time-step. Thus, the improvement showed by the agents becoming counteradaptive cannot be adequately explained by divergences in the collective history. For runs with identical history pairs, strategy performance does not (cannot) differ at all at any time-step, since the action of the strategies at all time-steps is wholly determined by the history. So, further, the improvement shown by the counteradaptive agents with respect to their own performance when kept standard holds as well for their performance with respect to their strategies. Henceforth we will drop any requirement that only runs with maximally-identical histories be compared.

The mean performance of all strategies with 3 agents choosing counteradaptively (black, solid line) in **Figure 17** is somewhat better than the mean performance of all strategies with no agents choosing counteradaptively (black, triangular points). The mean performance of all agents with 3 agents choosing counteradaptively (blue, solid line) is much better than the mean performance of all agents with no agents choosing counteradaptively

(black, square points). For both simulations mean agent performance is very much worse than mean strategy performance. These relations suggest what **Figure 20** breaks out in detail: That agents choosing counteradaptively perform significantly better than standard agents.



**Figure 20:** Comparative effect on agent v. strategy wealth of standard v. counteradaptive optimizations in the (TH)MG for  $\{N, m, S\} = \{31, 2, 2\}$  and  $\tau = 1$ .


**Figure 20** is composed as follows.

- ▲ All agents are standard, but this trajectory breaks out the mean wealth accumulation by strategies for those agents that are standard in both.
- ▲ All agents are standard, but this trajectory breaks out the mean wealth accumulation by strategies for those agents that in the second simulation will change to the counteradaptive optimization. In the standard run, the mean wealth accumulation by the strategies of these agents happens to be somewhat better the remaining agents' strategies. This is a matter of chance. It depends upon which agents are selected at random to use the other selection method in the second run.
- All agents are standard, but this trajectory breaks out the mean wealth accumulation by agents for those agents that are standard in both.
- All agents are standard, but this trajectory breaks out the mean wealth accumulation by agents for those agents that in the second simulation will change to the counteradaptive optimization. Note that for both groups in the standard run, wealth accumulation by agents significantly underperforms wealth accumulation by agents' strategies. This effect happens to be

especially striking for the 3 agents that in the second run will change their optimization method.

- 28 agents are standard, 3 counteradaptive. This trajectory breaks out the mean wealth accumulation by strategies for the 28 agents that remain standard in the second simulation.
- 28 agents are standard, 3 counteradaptive. This trajectory breaks out the mean wealth accumulation by strategies for the 3 agents that change to counteradaptive in the second simulation. Their mean is somewhat better than when they were standard, but remains negative.
- 28 agents are standard, 3 counteradaptive. This trajectory breaks out the mean wealth accumulation by agents for the 28 agents that remain standard in the second simulation.
- 28 agents are standard, 3 counteradaptive. This trajectory breaks out the mean wealth accumulation by agents for the 3 agents that change to counteradaptive in the second simulation. Their mean is dramatically better than when they were standard, and is in fact strongly positive in this case where they were the most strongly negative when standard. Most important is the simpler fact that these agents outperform rather than underperform their respective strategies.


**Figure 21** illustrates a similar effect averaged over 20 different randomly-selected  $\hat{\Omega}_N$  for  $\{N, m, S\} = \{101, 2, 2\}$  with 10 C agents. Here the counteradaptive agents outperform the standard ones but do not necessarily achieve a positive gain. Actual results for a particular agent/strategy configuration are highly sensitive both to  $\hat{\Omega}_N$  as well as to which and how many are C agents. If the fraction of C agents is very small ( $\leq 0.05$ ), on average the system-as-whole behaves very similarly to the same system with no C agents. But the performance of the C agents varies widely from run to run. If the fraction of C agents is large ( $\geq 0.25$ ), they constitute a more representative subset of all agents and their behavior varies less from run to run. But these agents then cause the behavior of the system as a whole to differ significantly in comparison to the same system with no counteradaptive agents.



**Figure 21:** Simulation of (TH)MG with 101 agents for 500 time-steps averaged over 20 randomly-selected agent/strategy configurations. 10 of the 101 agents are C agents, i.e., they choose their strategy counteradaptively.

In a mixed group of agents all with fixed memory the performance of the counteradaptive agents and standard agents both relative to each other and relative to zero mean gain depends in a consistent fashion both upon the total number and relative proportion of counteradaptive and standard agents. **Figure 22** illustrates the results of a THMG simulation with  $\tau = 1$  for each of  $N = 21, 31, \dots, 101$ , and for percentages of counteradaptive agents increasing from  $\sim 3\%$  to  $\sim 97\%$ . 10 runs of 100 time-steps were performed for each set of values: Every triple consisting of a blue, red and green triangle represents the mean of 10 final values at the end of a run. Every plot represents a different  $N$ . (The solid lines represent least-square polynomial fits of order 3). Blue represents the mean  $\Delta wealth/\Delta t$  for the standard agents at a given  $N$  at decreasing proportions of standard agents; red represents the mean  $\Delta wealth/\Delta t$  for the counteradaptive agents at a given  $N$  at increasing proportions of counterproductive agents; green represents the difference between the two.

Each of the  $\sim 1800$  runs were initialized independently with  $2N$  strategies randomly chosen from the same reduced strategy space in  $N$  pairs from the space of “typical” pairs restricted as discussed in the mathematical appendix, and with random initial histories.




**Figure 22:** Counteradaptive (CA) agent outperformance at different  $N$ , %C

Note the consistent shape of all three curves with varying % C agents, across  $N$ . For % C small, the C agents outperform the S agents and conversely (to a greater degree) at % C large. Equality of performance is reached consistently at % C  $\approx$  30. Note, too, that as the % C increases, it grows more likely that the game will settle into a Nash equilibrium and


the gain of the S agents not only exceeds that of the C ones, it is consistently positive and large.

More intriguing is the fact that the same is true in reverse, if much less markedly so at % C small ( $\approx 10\% - 15\%$ ): Not only do the counteradaptive agents outperform standard agents, within a narrow but consistent range of  $\{N, \%C\}$  they predictably achieve small absolute positive gain. The positive gain surfaces from the same data as in **Figure 22** are shown in **Figure 23**, the ungridded sections representing positive gain for counteradaptive agents, the gridded surface for standard agents. For range of  $\{N, \%CA\}$  with positive gain for the counteradaptive agents the game does not settle into Nash equilibria states. Hence the outperformance and positive gain of the counteradaptive agents is not associated with the loss of frustration in the system as for % CA large, but occurs in spite of it. It is important to emphasize that this positive gain is related neither to additional information or intelligence afforded the CA agents, nor to their adopting a different optimization procedure of evident advantage.



**Figure 23:** C (un-gridded) and S (gridded) agent performance surfaces

The underperformance of S agents relative to their strategies is found more predictably under different conditions than is the outperformance of C agents relative to S ones. And both phenomena appear even more consistently in the standard MG than in the THMG with  $\tau$  small (as in the examples just discussed). **Figure 24** illustrates the same phenomenon (and with largely positive absolute wealth gain) for the standard (non-TH) MG. But the standard MG with modest  $N$  is intractable in principle (and in practice the THMG with a very long time-horizon), to the kind of analytic reductions that will be used shortly, based on [41, 43], with which we can better understand how this paradoxical seeming effect arises.




**Figure 24:** Outperformance of 15 C agents by comparison to 86 S agents as a function of  $m$ , for  $N = 101$ ,  $S = 2$  over 10 runs showing high degree of absolute positive performance for  $m$  small. These results are for the standard MG in which strategy points accumulate without bound (over  $\tau$  growing with  $t$ ). The general form of the relation of  $\sigma$  to  $m$  (and of the dispersion of  $\sigma$ ) is present in spite of the fact that  $\sim 15\%$  of the agents are following the C rule. Outperformance disappears as  $\sigma$  reaches its minimum ( $\alpha \rightarrow \alpha_c$ ) and remains absent as the game enters the random or asymmetric phase ( $\alpha > \alpha_c$ ). The spread in the counteradaptive values at each  $m$  (red) is greater than the spread in the standard values (blue) because of the smaller number of agents for the former.

## 7.2 Illusion of control in the THMG at various $\tau$

We have emphasized that the illusion of control in the THMG is present for all “reasonable values” of  $\tau$ . **Figure 24** demonstrates that in the standard MG where the system is allowed to run to equilibrium and  $\tau$  is unbounded (and very large), the illusion is absent for values of  $\alpha \geq \alpha_c$  (where agents freeze). As the illusion vanishes, so too does the outperformance of C agents, as expected (since this outperformance is a powerful consequence of the illusion). It is instructive to examine the relation between the illusion and  $\tau$  in greater detail. This is shown in **Figure 25**. Here we increase  $\tau$  in approximately regular (exponentially growing) intervals from  $\tau = 1$  to  $\tau = 31620$  for a system with  $\{N, m, S\} = \{31, 2, 2\}$ . We also compare results for real histories to random histories. (We have not previously focused on this difference in this document. This graph demonstrates that in examining the illusion of control, the differences are quantitative but not qualitative: Random histories alter the position in the historical record of a given  $A(t)$  but not the response of the system to that  $A(t)$  with respect to agent and strategy wealth. As will be evident from the mathematical appendix, another difference is the proportional distribution of the various allowed path-histories  $\mu(t)$ . With random histories the vector  $\bar{\mu}$  is

composed of identical entries (steady states are equiprobable), equal to the mean of all  $\bar{\mu}$  averaged over an infinite number of randomly-selected  $\hat{\Omega}_N$  for given  $m, S$ .)



**Figure 25:** Illusion of control in the THMG at values of  $\tau$  from 1 to “reasonable” (31620), for both real histories (top 10 graphs) and random histories (bottom 10 graphs) with  $N = 31, m = 2, S = 2$ .

- $\sigma^2/N$
- $\Delta W_{strat} - \Delta W_{agent}$
- ◆  $\Delta W_{strat}$
- ▲  $\Delta W_{agent}$

We may note the following general features.


First, as  $\tau$  grows, the phase-transition at  $\alpha_c$  becomes increasingly evident ( $\sigma^2/N$ , the standard measure of inverse system cooperativity or wealth, is provided in black squares for convenience. It varies as the inverse square of  $\Delta W_{agent}$  shown as grey triangles).

Second, as  $\tau$  grows, the excess mean gain of strategies over (supposedly optimizing) agents at  $m_c$  — i.e.,  $\Delta W_{strat} - \Delta W_{agent}$  — shrinks. Because the maximum number of equilibrating steps, though very large, is less than the number of steps required to attain equilibrium (which at  $m_c$  is exceedingly large), it remains the case that, though small and at a minimum,  $\Delta W_{strat} - \Delta W_{agent} > 0$ , i.e., the illusion of control persists throughout.

Third, away from  $m_c$ , the illusion persists throughout and is large.

### 7.3 Performance of optimizing (S) versus random-choice agents in the THMG

**Figure 26** shows the performance (mean change in wealth per step) of a single optimizing agent embedded in a THMG-like game versus all other agents making a symmetric random choice. From left to right  $N = 11, 21, 31$ ;  $S = 2, 3$ ;  $m = 2, 3, 4, 5$ ; and  $\tau = 1$ . Random agents always outperform optimizing agents. Similar results are found for other values of  $N, m, S$  and  $\tau$ . Within statistical fluctuations typical for the number of runs (100) and for the random selection of strategies, results for anti-optimizing agents are identical.



**Figure 26:** Performance (mean change in wealth per step) of a single optimizing agent embedded in a THMG-like game versus all other agents making a symmetric random choice.

## Chapter 8. Discussion

In Chapter 8, we review and summarize our major findings from a bird's eye perspective. We discuss limitations in the present research and suggest further avenues of research.

### **8.1 Status of the illusion of control; anti-persistence**

“Illusion of control” is a problem that appears in many domains. We have here demonstrated certain formal conditions under which it appears: In multi-player MG and in an attempt to optimize the Parrondo effect in single-player games of the Parrondo type (PG). Both classes of games have experienced broad and growing acceptance as models for a very wide range of human activities, especially the MG and its variants. The effect is so powerful in the MG under most circumstances (particularly in the THMG variant) that the deliberate inversion of its optimization rule actually leads to superior results. (Consider how unlikely it is that a trader would deliberately choose to deploy that strategy which in his hypothetical simulations consistently demonstrated the worst results.)

However, the illusion of control is not present under all circumstances in the MG (even if very widely and under most realistic constraints) and it is entirely absent in the MAJ and \$G, whether TH-variant or not.

The mechanism for the illusion in PG's is unrelated to its mechanism in MG's. It is not possible to denote a single underlying similarity in cause and this doubtless reflects the fact that failure of optimization in real-world settings may be due to many different causes.

However, in that large class of phenomena that are characterized by a one-dimensional time series arising from the interaction of many different agents, i.e., complex adaptive systems (CAS), our results suggest that the presence of anti-persistence in particular may be a warning flag for would-be optimizers. Thus when applying the method of cycle decomposition to the creation of real-world predictors—which arises naturally in the study of the TH games—we find strong indicators that anti-persistence is at least one cause for periods of weaker success (or outright failure).

### **8.2 Cycle decomposition**

Complementary to the assessment of persistence in a binary time series, we have demonstrated a simple version of the cycle decomposition method as a rapid and visually appealing way to characterize time series and, in fact, to achieve a certain kind of optimization, e.g., prediction. While we have here restricted ourselves to elementary applications, we anticipate (and preliminary results not included support the impression) that this method may yield a rich harvest in future studies.

In particular, methods that employ higher-base decompositions (e.g., ternary instead of binary series) should be especially appealing in the financial forecasting arena since many financial series are much more usefully projected not onto a two-state spectrum of “up” and “down” changes, but onto a three-state one, “up”, “down” and “no change”. Clearly in the financial domain especially, the unusually large number of very small changes begs for a significant role for a “no change” element. (Financial time-series typically show both fatter tails than in a Gaussian distribution *and* sharper central peaks. Set-

ting a reasonable range of middle value to zero helps to emphasize the “fat tails” indicative of the strongest dependencies.) Nonetheless, the fact that we are able to make significant headway using “toy” predictors that lack this feature is an indication of the potential of the method.

We emphasize that there is probably nothing dramatically new in this approach that could not be achieved also, perhaps if only with greater toil, by more conventional statistical methods. An old argument comes to mind: For many years the statistics community routinely dismissed artificial neural networks (ANN) as nothing new—they simply codified in a different way various kinds of extant non-parametric statistical methods. Nonetheless, for many problems, the implementation of an ANN is the preferable solution method (noise-canceling headphones, sonar identification of foreign underwater targets, to give two current practical implementations).

Similarly, the statistical dependencies in financial markets discussed by Zhang [102] appear elsewhere in market prediction methods closely related to (and often applied as a part of) ANN methods, namely “cased-based reasoning” (CBR)[112]. In CBR, prior experiences are cataloged as “cases” along with their typical outcomes and the current situation is matched to the appropriate case to make a prediction. In so-called “fuzzy” CBR the likelihood of a certain next-step is estimated based on the statistics of the prior cases to which it is matched. This method of conceptualizing statistical dependencies is not inherently different than statistical methods proper but lends itself to machine learning (and to implementation as an ANN of the “vector learning” variety). Perhaps the introductory cycle-decomposition techniques we have presented here may be thought of in the same spirit as very modest analogs to the use of Feynman diagrams to express rapidly the dynamics presented in Schwinger’s equations for QCD

### **8.3 Limitations and future lines of research**

The literature on MG in particular has grown very rich as the field has matured. There are many “branch points” in the development of this work which would lead to obvious and important corollary studies. Here are some we have deferred for forthcoming efforts.

#### **8.3.1 Effects of $S > 2$ in MG, MAJG and \$G**

Much of the literature has made clear that the most significant phenomena associated with the MG, MAJG and \$G are present when each agent has only 2 strategies. We have spent time investigating what happens in these games when a subset of agents choose their worst strategy which, for  $S = 2$  means simply, “the other strategy”. At least one study has examined what differences arise when agents take into account on the predictions generated by their second best strategy (with  $S > 2$ , of course)[56]. It would be helpful, we think, to study more carefully what would happen were the subset of “C” agents to act solely in accord with their second-best, third-best, ... strategy.

#### **8.3.2 Grand Canonical MG, MAJG and \$G**

More importantly, we are interested in extending the analyses we have presented here to variations of the Grand Canonical (GC) MG/THMG (and MAJG, \$G). Many features of real-world markets that are absent from the non-GC version are present in the GC games. In GC games, agents may participate or not depending upon whether a certain perform-

ance threshold has been met. Allowing agents to *not* participate is analogous to ternary cycle decomposition. Just as extremely small changes are best classified as “0” rather than + or –, it makes sense that agents would be allowed to do nothing as part of the actions associated with their strategies.

Would the illusion of control, when present, be greater, less or not significantly different under these circumstances?

Related to this consideration is the role of “volume” in real world markets and various studies of so-called “momentum” trading. Volume may be studied in the CG version of agent-based games and related to the illusion of control.

### **8.3.3 Agents with mixed memory and/or lookback**

There is a large literature on the standard MG but with agents of mixed memory  $m$ , a variation that conforms to expectation that in real-world markets traders have a range of computational capacities (intelligence; see, for example ref.s [106, 107]). It would be valuable to examine the THMG, THMAJG and TH\$G under these conditions with an eye toward understanding where and why the illusion of control appears and disappears. Using the TH variants, we may also explore the role of having agents with different lookbacks, differentiating this condition from their having different strategy memories.

### **8.3.4 Other Games**

We have not explored the illusion of control in other games meant to mimic financial markets even more rigorously. Challet has recently introduced a variation of the three basic agent-based games we discuss that includes, importantly, 1) all three reward mechanisms and 2) holding periods (an obvious difference between MG, MAJG and \$G and real-world markets). This variation has only just begun to attract research attention. It would be instructive to determine if and when in this single game the illusion of control appears/disappears.

## Chapter A9. Appendix: mathematical methods

Chapter A9 is a mathematical methods appendix. We present formal analytic expressions for all three agent-based games. We present methods for illustrating the illusion of control effect in generalized single-player Parrondo games. We also present methods for determining the persistence/anti-persistence of a binary series, for decomposing a given binary series into a superposition of cycles on graphs and for obtaining analytically the cycle-decomposition of a known Markov transition matrix.

### A9.1: Analytic Methods for the TH –MG, -MAJG and -\$G

#### A9.1.1 The Time Horizon Minority Game: Choosing the Best Strategy

In the simplest version of the Minority Game (MG) with  $N$  agents, every agent has  $S = 2$  strategies and  $m = 2$ . In the Time Horizon Minority Game (THMG), the point (or score) table associated with strategies is not maintained from the beginning of the game and is not ever growing. It is a rolling window of finite length  $\tau$  (in the simplest case  $\tau = 1$ ). The standard MG reaches an equilibrium state after a finite number of steps  $t_{st}$ . At this point, the dynamics and the behavior of individual agents for a given initial quenched disorder in the MG are indistinguishable from an otherwise identical THMG with  $\tau \geq t_{st}$ .

The fundamental result of the MG is generally cast in terms of system volatility:  $\sigma^2/N$ , where  $\sigma$  is the variance over time in the agent vote (i.e., the sum of all agents' "actions"). All variations of agent and strategy reward functions depend on the negative sign of the majority vote. Therefore both agent and strategy "wealth" (points, whether "real" or hypothetical) are inverse or negative functions of the volatility: The lower the value of  $\sigma^2/N$ , the greater the mean "wealth" of the "system", i.e., of agents. However, this mean value is only rarely compared to the comparable value for the raw strategies of which agents are composed. Yet agents are supposed to enhance their performance by choosing adaptively between their available strategies. In fact, the opposite is true in the THMG: The optimization method is strictly a method for worsening performance, not necessarily for every agent, but averaged over all agents and all strategies in a given  $\hat{\Omega}$ , averaged over many  $\hat{\Omega}$ .

To emphasize the relation of the THMG to market-games and the illusion of optimization, we transform the fundamental result of the THMG from statements on the properties of  $\sigma^2/N$  to change in wealth, i.e.,  $\langle \Delta W / \Delta t \rangle$  for agents and  $\langle \Delta W / \Delta t \rangle$  for strategies. We again use the simplest possible formulation—if an agent's actual (or a strategy's hypothetical) vote places it in the minority, it scores +1 points, otherwise -1. Formally: At every discrete time-step  $t$ , each agent independently re-selects one of its  $S$  strategies. It "votes" as the selected strategy dictates by taking one of two "actions," designated by a binary value:

$$a_i(t) \in \{1, 0\}, \quad \forall i, t \quad (23)$$

The state of the system as a whole at time  $t$  is a mapping of the sum of all the agents' actions to the integer set  $\{2N_1 - N\}$ , where  $N_1$  is the number of 1 votes and  $N_0 = N - N_1$ . This mapping is defined as :

$$A(t) = 2 \sum_{i=1}^N a_i(t) - N = \sum_{i=1}^N [2a_i(t) - 1] = N_1 - N_0 \quad (24)$$

If  $A(t) > \frac{N}{2}$ , then the minority of agents will have chosen 0 at time  $t$  ( $N_0 < N_1$ ); if  $A(t) < \frac{N}{2}$ , then the minority of agents will have chosen 1 at time  $t$  ( $N_1 < N_0$ ). The minority choice is the “winning” decision for  $t$ . This is then mapped back to  $\{0,1\}$ :

$$D_{\text{sys}}(t) = -\text{Sgn}[A(t)] \quad \therefore D_{\text{sys}}(t) \in \{-1,+1\} \rightarrow \{0,1\} \quad (25)$$

For the MG, binary strings of length  $m$  form histories  $\mu(t)$ , with  $m = \dim[\mu(t)]$ . For the THMG, binary strings of length  $m + \tau$  form paths (or “path histories”), with  $m + \tau = \dim(\mu_t)$ , where we define  $\mu(t)$  as a history in the standard MG and  $\mu_t$  as a path in the THMG. Note that for memory  $m$ , there are  $2^{2^m}$  possible strategies from which agents select  $S$  strategies at random (with substitution). However detailed in ref. [4], the space of strategies can be minimally spanned by a subset of all possible strategies. This reduced strategy space [RSS] has dimension  $2^{m+1}$ . As in ref. [10] we may represent this quenched disorder in the allocation of strategies among agents (from the RSS) by a  $\text{dim} = \prod_{s=1}^S 2^{m+1}$  tensor,  $\hat{\Omega}$  (or from the full strategy space by a  $\text{dim} = \prod_{s=1}^S 2^{2^m}$  tensor). The  $2^{m+1}$  (or  $2^{2^m}$ ) strategies are arranged in numerical order along the edges of  $\hat{\Omega}$ . Each entry represents the number of agents with the set of strategies indicated by the element’s position. Then as demonstrated in [9], any THMG has a Markov chain formulation. For  $\{m, S, N\} = \{2, 2, 31\}$  and using the RSS, the typical initial quenched disorder in the strategies attributed to each of the  $N$  agents is represented by an  $8 \times 8$  matrix  $\hat{\Omega}$  and its symmetrized equivalent  $\hat{\Psi} = \frac{1}{2}(\hat{\Omega} + \hat{\Omega}^\top)$ . Positions along all  $S$  edges of  $\hat{\Omega}$  represent an ordered listing of all available strategies. The numerical values  $\Omega_{ij}$  ( $\Omega_{ij\dots}$  for  $S > 2$ ) in  $\hat{\Omega}$  indicate the number of times a specific strategy-tuple has been selected in the initial endowment of the  $S$  strategies to the  $N$  agents. (E.g., for two strategies per agent,  $S = 2$ ,  $\Omega_{2,5} = 3$  means that there are 3 agents with strategy 2 and strategy 5.) Without loss of generality, we may express  $\hat{\Omega}$  in upper-triangular form since the order of the strategies comprising an agent has no meaning. The example (26) is a typical such tensor  $\hat{\Omega}$  for  $\{m, S, N\} = \{2, 2, 31\}$ .

$$\hat{\Omega} = \begin{pmatrix} 1 & 2 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 3 & 3 & 1 & 1 \\ 0 & 0 & 2 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 2 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (26)$$

Actions are drawn from the reduced strategy space (RSS) of dimension  $2^{m+1}$  [4]. Each action is associated with a strategy  $k$  and a history  $\mu(t)$ . Together they can be represented in table form as a  $\dim(\text{RSS}) \times \dim[\mu(t)]$  binary matrix with elements converted for convenience from  $\{0,1\} \rightarrow \{-1,+1\}$ , i.e.,  $a_k^{\mu} \in \{-1,+1\}$ . For  $m = 2$ , there are  $2^2 = 4$  possible histories and  $r = 2^3$  reduced strategies. In this case, the table coding for all possible reduced strategies and paths reads:

$$\hat{a} \equiv \begin{pmatrix} -1 & -1 & -1 & -1 \\ -1 & -1 & +1 & +1 \\ -1 & +1 & -1 & +1 \\ -1 & +1 & +1 & -1 \\ +1 & -1 & -1 & +1 \\ +1 & -1 & +1 & -1 \\ +1 & +1 & -1 & -1 \\ +1 & +1 & +1 & +1 \end{pmatrix} \quad (27)$$

The change in wealth (point gain or loss) associated with each of the  $r = 8$  strategies for the 8 paths at any time  $t$  (i.e., 8 allowed transitions  $\mu(t-1) \rightarrow \mu(t)$  among the 4 histories) is then:

$$\delta \bar{S}_{\mu(t), \mu(t-1)} = \left( \hat{a}^T \right)_{\mu(t)} \times \{2 \text{Mod}[\mu(t-1), 2] - 1\} \quad (28)$$

$\text{Mod}[x, y]$  is “x modulo y”;  $\mu(t)$  and  $\mu(t-1)$  label each of the 4 histories  $\{00, 01, 10, 11\}$  hence take on one of values  $\{1, 2, 3, 4\}$ . Equation (28) picks out from (27) the correct change in wealth over a single step since the strategies are ordered in symmetrical sequence.

The change in points associated with each strategy for each of the allowed transitions between paths  $\mu_i$  of the last  $\tau$  time steps used to score the strategies is:

$$\bar{s}_{\mu_i} = \sum_{i=0}^{\tau-1} \delta \bar{S}_{\mu(t-i), \mu(t-i-1)} \quad (29)$$

For example, for  $m = 2$  and  $\tau = 1$ , the strategy scores are kept for only a single time-step. There is no summation so (29) in matrix form reduces to the score:

$$\bar{s}_{\mu_t} = \delta \bar{\mathbf{S}}_{\mu(t), \mu(t-1)} \quad (30)$$

or, listing the results for all 8 path histories:

$$\hat{\mathbf{s}}_{\mu} = \delta \hat{\mathbf{S}} \quad (31)$$

$\delta \hat{\mathbf{S}}$  is an  $8 \times 8$  matrix that can be read as a lookup table. It denotes the change in points accumulated over  $\tau = 1$  time steps for each of the 8 strategies over each of the 8 path-histories.

Instead of computing  $A(t)$ , we compute  $A(\mu_t)$ . Then for each of the  $2^{m+\tau} = 8$  possible  $\mu_t$ ,  $A(\mu_t)$  is composed of a subset of wholly determined agent votes and a subset of undetermined agents whose votes must be determined by a coin toss:

$$A(\mu_t) = A_D(\mu_t) + A_U(\mu_t) \quad (32)$$

Some agents are undetermined at time  $t$  because their strategies have the same score and the tie has to be broken with a coin toss.  $A_U(\mu_t)$  is a random variable characterized the binomial distribution. Its actual value varies with the number of undetermined agents. This number can be explicated (using an extension to the method employed in [9]) as :

$$N_U(\mu_t) = \left\{ \left( 1 - \left[ (\hat{a}_1^T)_{(\text{Mod}[\mu_t-1,4]+1)} \otimes_{\delta} (\hat{a}_1^T)_{(\text{Mod}[\mu_t-1,4]+1)} \right] \right) \circ (\bar{s}_{\mu_t} \otimes_{\delta} \bar{s}_{\mu_t}) \circ \hat{\mathbf{\Omega}} \right\}_{(\text{Mod}[\mu_t-1,2^m]+1)} \quad (33)$$

“ $\otimes_{\delta}$ ” is a generalized outer product, with the product being the Kronecker delta.  $\bar{N}_U$  constitutes a vector of such values. The summed value of all undetermined decisions for a given  $\mu_t$  is distributed binomially. Similarly:

$$A_D(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ (1 - \text{Sgn}[\bar{s}_{\mu_t} \ominus \bar{s}_{\mu_t}]) \circ \hat{\Psi} \right] \cdot \hat{a}_1 \right\}_r \right)_{(\text{Mod}[\mu_t-1,2^m]+1)} \quad (34)$$

An example of how (33) and (34) can be deduced is given later in the context of the original definition of alternate types of agents. Details may also be found in ref. [11]. We define  $\bar{A}_D$  as a vector of the determined contributions to  $A(t)$  for each path  $\mu_t$ . In expressions (33) and (34)  $\mu_t$  numbers paths from 1 to 8 and is therefore here an index.  $\bar{s}_{\mu_t}$  is the “ $\mu_t^{\text{th}}$ ” vector of net point gains or losses for each strategy when at  $t$  the system has traversed the path  $\mu_t$  ( i.e., it is the “ $\mu_t^{\text{th}}$ ” element of the matrix  $\hat{\mathbf{s}}_{\mu} = \delta \hat{\mathbf{S}}$  in (31)). “ $\ominus$ ” is a generalized outer product of two vectors with subtraction as the product. The two vectors in this instance are the same, i.e.,  $\bar{s}_{\mu_t}$ . “ $\circ$ ” is Hadamard (element-by-element) multiplica-

tion and “•” the standard inner product. The index  $r$  refers to strategies in the RSS. Summation over  $r$  transforms the base-ten code for  $\mu_t$  into  $\{1, 2, 3, 4, 1, 2, 3, 4\}$ . Selection of the proper number is indicated by the subscript expression on the entire right-hand side of (33). This expression yields an index number, i.e., selection takes place  $1 + \text{Modulo } 4$  with respect to the value of  $(\mu_t - 1)$ .

To obtain the transition matrix for the system as a whole, we require the  $2^{m+\tau} \times 2^{m+\tau}$  adjacency matrix that filters out disallowed transitions. Its elements are  $\Gamma_{\mu_t, \mu_{t-1}}$ :

$$\hat{\mathbf{\Gamma}} = \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \end{pmatrix} \quad (35)$$

Equations (33), (34) and (35) yield the history-dependent  $(m + \tau) \times (m + \tau)$  matrix  $\hat{\mathbf{T}}$  with elements  $T_{\mu_t, \mu_{t-1}}$ , representing the 16 allowed probabilities of transitions between the two sets of 8 path-histories  $\mu_t$  and  $\mu_{t-1}$ :

$$T_{\mu_t, \mu_{t-1}} = \Gamma_{\mu_t, \mu_{t-1}} \times \sum_{x=0}^{N_U(\mu_t)} \left\{ \binom{N_U(\mu_t)}{x} \left(\frac{1}{2}\right)^{N_U(\mu_t)} \times \delta \left[ \text{Sgn}(A_D(\mu_t) + 2x - N_U(\mu_t)) + (2 \text{Mod}\{\mu_{t-1}, 2\} - 1) \right] \right\} \quad (36)$$

The C expression  $\binom{N_U(\mu_t)}{x} \left(\frac{1}{2}\right)^{N_U(\mu_t)}$  in (36) represents the binomial distribution of undetermined outcomes under a fair coin-toss with mean =  $A_D(\mu_t)$ . Given a specific  $\hat{\mathbf{\Omega}}$ ,

$$\langle A(\mu_t) \rangle = A_D(\mu_t) \quad \forall \mu_t \quad (37)$$

We now tabulate the number of times each strategy is represented in  $\hat{\mathbf{\Omega}}$ , regardless of coupling (i.e., of which strategies are associated in forming agent  $S$ -tuples):

$$\bar{\mathbf{k}} \equiv \sum_{k=1}^{2^{m+\tau}} (\hat{\mathbf{\Omega}} + \mathbf{\Omega}^T)_k = 2 \sum_{k=1}^{2^{m+\tau}} \hat{\mathbf{\Psi}}_k = \{n(\sigma_1), n(\sigma_2), \dots, n(\sigma_{2^{m+\tau}})\} \quad (38)$$

where  $\sigma_k$  is the  $k^{\text{th}}$  strategy in the RSS,  $\hat{\mathbf{\Omega}}_k, \hat{\mathbf{\Omega}}_k^T$  and  $\hat{\mathbf{\Psi}}_k$  are the  $k^{\text{th}}$  element (vector) in each tensor and  $n(\sigma_k)$  represents the number of times  $\sigma_k$  is present across all strategy tuples. Therefore


$$\langle \Delta W_{Agent} \rangle = -\frac{1}{N} Abs(\bar{A}_D) \cdot \bar{\mu} \quad (39)$$

and

$$\langle \Delta W_{Strategy} \rangle = \frac{1}{2N} (\hat{\mathbf{s}}_{\mu} \cdot \bar{\mathbf{K}}) \cdot \bar{\mu} \quad (40)$$


with  $\bar{\mu}$  the normalized steady-state probability vector for  $\hat{\mathbf{T}}$ . Expression (39) states that the mean per-step change in wealth for agents equals  $-1$  times the probability-weighted sum of the (absolute value of the) *determined* vote imbalance associated with a given history. Expression (40) states that the mean per-step change in wealth for individual strategies equals the probability-weighted sum of the representation of each strategy (in a given  $\hat{\Omega}$ ) times the sum over the per-step wealth change associated with every history. The  $-1$  in (39) reflects the minority rule. I.e., the awarding of points is the negative of the direction of the vote imbalance. No minus sign is required in (40) as it is already accounted for in (27).

**Figure 27** shows the cumulative mean change in wealth for strategies versus agents over time, given (26).



**Figure 27:** Mean Strategy versus Agent Cumulative Change in Wealth in the THMG.  $\{m, S, N\} = \{2, 2, 31\}$ ; 100 time steps

As first studied in [12,13], and discussed in the body of the manuscript, agent performance is inversely proportional to the Hamming distance between strategies within agents. With the variation expected of a single example, our sample  $\hat{\Omega}$  given by (26) reproduces this relation as shown in **Figure 28**. Thus agent performance is distributed within  $\hat{\Omega}$  in orderly if complex fashion. The mean over many  $\hat{\Omega}$  corresponds to a “flat”  $\hat{\Omega}$ .



**Figure 28:** Agent wealth as a function of Hamming distance between strategy pairs in agents for the example simulation.

### A9.1.2 The Time Horizon Minority Game: Choosing the Worst Strategy

The Markov equations (33), (34) and (35) have been modified only slightly from their original form developed in [41, 43]. To account for the impact on the system of including an arbitrary selection of “counteradaptive” (or C agents, that select their previously worst-performing strategy; standard S agents select their best), the equations need a more significant extension. It is important, too, that the extended equations distinguish the C sub-population of agents from the S sub-population. We will see that the necessary analytic modifications are surprisingly simple and they yield results in as close concord with numerical simulation as the original equations do for the standard agents alone in the THMG (as illustrated in Chapter 7. We precede writing out the modified equations with a brief explanation of the reasoning that leads to them. This requires a “microscopic” illustration of the relation among strategies within agents following different path histories. Results differ in part depending on whether an agent’s “selection rule” is C or S, and requires a prior determination—unaffected by the “selection rule”—of whether an agent’s action at the next step is going to be “determined” or “undetermined”.

Regardless of whether C agents are present or not, computation of the transition matrix  $\hat{\mathbf{T}}$  requires all critical values: If we can derive it, we will have been able to obtain  $\bar{N}_U$ ,  $\bar{A}_D$  and  $\bar{\kappa}$  as well as characterize the steady state of the system. In the THMG without C agents  $\bar{N}_U$  and  $\bar{A}_D$  both derive directly from  $\hat{\mathbf{\Omega}}$  because two or more agents consisting of the same tuple of strategies are indistinguishable except in certain cases of tie-breaks. When C agents are included, and when they may distributed without restriction among the agent population, agents consisting of the same tuple of strategies need no longer necessarily be thus indistinguishable, but will behave differently often. Thus  $\hat{\mathbf{\Omega}}$  cannot be used so simply and directly when C agents are present as when they are absent. Of course, neither can  $\hat{\Psi} = \frac{1}{2}(\hat{\mathbf{\Omega}} + \hat{\mathbf{\Omega}}^T)$ . For example, the sum on the right-hand side of (34) must be carried out *after* taking into account whether each agent is C or S—that is, seemingly agent-by-agent.

But as we will see, the degree of complexity introduced by C agents is less than it may seem. First,  $\bar{N}_U$  follows unchanged from  $\hat{\Omega}$  regardless of whether the agents of each non-zero element in  $\Omega_{i,j}$  are C or S or both (iff  $\Omega_{i,j} > 1$ ). Second, only  $\bar{A}_D$  must be calculated differently than as in (34). The component values  $A(\mu_t)$  for each path  $\mu_t$  remain both determinable for both C and S agents. All other computations are identical to those in section A91.1. We want to determine the contribution of agent  $k$  to each component (by path) of  $\bar{A}_D$ , depending on whether  $k$  acts as a C-type or as an S-type agent. We will then be able to separately add all the C-type contributions to  $\bar{A}_D$  for each path and all the S-type contributions. In the end, we will find that we may simply redefine  $\hat{\Omega}$  as being composed of two tensors, one C-type, the other S-type. In the conventional THMG the initial quenched disorder is expressed by the single S-type matrix alone. For the THMG with C-agents, the two matrices may be handled separately using the same mathematical formalism, and the results summed with but a single change of sign at the appropriate location. Since the exact spot for the change is not intuitively evident at a glance, it is worth showing how it arises. This provides an opportunity to review the reasoning behind the equations.

Consider an arbitrary agent  $\{\sigma_1, \sigma_2\}_k = \{1, 2\}$ ,  $1 \leq k \leq N$  belonging to  $\hat{\Omega}_N$ . In this example we limit element  $\Omega_{1,2} = 1$  so that *only*  $\{\sigma_1, \sigma_2\}_k = \{1, 2\}$ . I.e., agent  $k$  (and for simplicity no other agent) consists of strategy 1 coupled with strategy 2. We will consider how agent  $k$  behaves depending upon the path history of the system and whether it S or C. The details are laid out in **Table 6**.

Suppose it is now time  $t$ . From (27)-(31) we know that when the system has just traversed path  $\mu_t$ , strategies  $\sigma_1 = 1$  and  $\sigma_2 = 2$  will be credited with a fixed and known number of points. In (31) (i.e.,  $\hat{s}_\mu$ ), these points are given by columns 1 and 2 respectively, with  $\mu_t$  numbering the rows. We have copied these columns as the first set of number-pairs in **Table 6**. The rows are labeled from 1 through 8 representing the 8 possible paths,  $\mu_t$ . Reading from left to right, the column headings show steps in computing agent  $k$ 's action at time  $t$ , for each prior path history. Thus, this first column of number-pairs is headed  $s_{\mu_t}$ , and sub-headed  $(\sigma_1, \sigma_2)$  for the scores accumulated by the respective strategies of agent  $k$ .

The “virtual action” that defines a strategy—how every strategy would vote at every time  $t$  based on  $\mu(t)$ —is encoded as its numeric label. (In this case, the numeric labels for strategies run from 1 through 8 and so do the labels for paths. That both numbers are 8 is of course only coincidental, because we have chosen the simplest  $\tau = 1$  so that  $2^{m+\tau} = 8 = 2^{m+1}$ .) In **Table 6** we represent this action for  $\sigma_1 = 1$  and  $\sigma_2 = 2$ , by path history, in the second column of number-pairs. This column is headed  $a^{\mu(t)}$  and again sub-headed  $(\sigma_1, \sigma_2)$ . Note that if we write  $a_k(t)$ , we mean to indicate the “actual” action taken by an agent—agent  $k$ .  $a_k(t)$  will contribute to a sum of actions constituting a state

of the system of the whole. For convenience, strategy “votes” have been mapped from their “native” 1/0 form (in which the minority state is generally counted) to +1/−1.

The next step is to find out if  $a_k(t)$ , agent  $k$ 's action at time  $t$ , is determined or undetermined by  $\mu_t$ . For its action to be undetermined, both strategies must have accumulated the same number of points following  $\mu_t$ , yet the action dictated by these strategies must differ. (Under these conditions, there is no basis for choosing one strategy over the other and a fair coin toss is used. It will make no difference whether the agent is C or S—every possible action has a fair shot at being selected at random. Thus also, the calculation of  $N_U$  and the results of each coin-toss are unaffected by whether an agent is C or S.)

The converse is therefore also true: If the strategies have accumulated a different number of points; or if the actions dictated by both strategies are the same (or both), then the agent's action at time  $t$  is determined. (Note for later comment that these are three logically distinguishable conditions.)

We define the “determinacy” of an agent  $\{\sigma_1, \sigma_2\}_k$  at  $t$  by the expression  $D_{\sigma_1, \sigma_2}(t; \mu_t) \in \{1, 0\}$ . (More precisely,  $D_{\sigma_1, \sigma_2}$  is the determinacy of the agent's action  $a_k(t)$ .) Iff  $D_{\sigma_1, \sigma_2}(t; \mu_t) = 1$ ,  $a_k(t)$  is determined and :

$$D_{1,2}(t; \mu_t) = 1 - \left( \delta_{1,2}(s_{\mu_t}) \times \left\{ 1 - \delta_{1,2} \left[ a^{\mu(t-1)} \right] \right\} \right) \quad (41)$$

The third column of number-pairs in **Table 6** is composed of the two Kronecker delta functions that are part of the product on the left-hand side of (41). The column is headed  $\delta_{1,2}$  and the pairs are sub-headed  $(s_{\mu_t}, a^{\mu(t)})$ . The first number in each pair is 1 if  $s_{\mu_t}(\sigma_1) = s_{\mu_t}(\sigma_2)$  and 0 otherwise; the second number is 1 if  $a_{\sigma_1}^{\mu(t)} = a_{\sigma_2}^{\mu(t)}$  and 0 otherwise. Column 4 calculates the results of (41) with 1 meaning  $a_k(t)$  is determined, 0 that it is undetermined. All the information is now at hand either to compute the action  $a_k(t)$  for all  $D_{\sigma_1, \sigma_2}(t; \mu_t) = 1$ , or to replace it with equiprobable random values if  $D_{\sigma_1, \sigma_2}(t; \mu_t) = 0$ . Determinacy and action may be similarly computed in brute fashion for all possible  $(2^{m+\tau})^S$ -tuples.

**Table 6:** Possible determined or undetermined actions  $a_k(t)$  for a given agent  $k$ . The agent is composed of strategies  $\sigma_1 = 1$ ,  $\sigma_2 = 2$ , and  $a_k(t)$  is shown assuming both standard (S) and a counteradaptive (C) selection rule. Both  $a_k(t)$  and all intermediate steps depend upon the prior path history  $\mu_t$  for the system-as-a-whole.  $a_k(t)$  for C-agents differs from that for S-agents only when both the accumulated points and the history-dictated actions of the component strategies differ. For the particular strategy-tuple shown, this occurs for  $\mu_t = 7$  or 8. Otherwise  $a_k(t)$  is either determined and identical for both C-agents and S-agents, or it is undetermined for both with two-state  $p = 0.5$ .

$path$	$s_{\mu_t}$	$a^{\mu(t)}$	$\delta_{1,2}$	$1 - \delta_{1,2} s_{\mu_t} \times$ $[1 - \delta_{1,2} a^{\mu(t)}]$	$a_k(t)$
$\mu_t$	$(\sigma_1, \sigma_2)$	$(\sigma_1, \sigma_2)$	$(s_{\mu_p}, a^{\mu(t)})$	$D_{1,2}(t, \mu_t)$	<b>(S, C)</b>
1	(+1, +1)	(+1, +1)	(1, 1)	1	(+1, +1)
2	(-1, -1)	(+1, +1)	(1, 1)	1	(+1, +1)
3	(+1, +1)	(+1, -1)	(1, 0)	0	( $\pm 1$ , $\pm 1$ )
4	(-1, -1)	(+1, -1)	(1, 0)	0	( $\pm 1$ , $\pm 1$ )
5	(+1, -1)	(+1, +1)	(0, 1)	1	(+1, +1)
6	(-1, +1)	(+1, +1)	(0, 1)	1	(+1, +1)
7	(+1, -1)	(+1, -1)	(0, 0)	1	<b>(+1, -1)</b>
8	(-1, +1)	(+1, -1)	(0, 0)	1	<b>(-1, +1)</b>

The last column of **Table 6** provides the computed actions for agent  $k$  by path. Equiprobable undetermined actions are indicated as  $\pm 1$ , determined actions as either +1 or -1. As mentioned above, there are three logically distinguishable conditions that yield a determined outcome for  $a_k(t)$ :

1. The strategies comprising the agent have accumulated the same number of points, and their strategies dictate the same actions;
2. The strategies comprising the agent have accumulated a different number of points, but their strategies dictate the same action;
3. The strategies comprising the agent have accumulated a different number of points, and their strategies dictate different actions;

When neither of conditions 1., 2. nor 3. hold,  $a_k(t)$  is undetermined. Strategy selection and agent action results from a random process.

When either of conditions 1. and 2. hold, the accumulated strategy scores are irrelevant (and differences between strategies). No selection takes place and the action is determined by:

$$a_k(t) = a_{\sigma_1}^{\mu(t)} = a_{\sigma_2}^{\mu(t)} \quad (42)$$

Only when condition 3. holds does strategy selection play a role in determining  $a_k(t)$ . Let us convert to 1/0 form the accumulated point values in the column of pairs labeled

$s_{\mu_t}$  in **Table 6** and make a dot product of each pair with the corresponding pair in the next column labeled  $a^{\mu(t)}$ . Multiply this by the corresponding value (in the same row) of  $D_{\sigma_1, \sigma_2}$ . The result is the action for agent  $k = \{\sigma_1, \sigma_2\}_k$ , for each path, assuming that agent  $k$  is S and not C. Determined actions will be expressed as +1 or -1; undetermined action as 0:

$$a_k(t) = D_{\sigma_1, \sigma_2}(t; \mu_t) \times \{2s_{\mu_t}(\sigma_1) - 1, 2s_{\mu_t}(\sigma_2) - 1\} \bullet \{a_{\sigma_1}^{\eta(t)}, a_{\sigma_2}^{\eta(t)}\} \quad (43)$$

Summing over all actual agents in  $\hat{\Omega}$  yields  $A_{D_\mu}$ . These steps are condensed in (34) (for S agents only). In **Table 6**, 0's have been replaced with  $\pm 1$ . These are accounted for probabilistically in equation (36) which remains unchanged for S agents.

Because C agents behave differently under condition 3., two additional steps are required to account for them here and in modifying (34). We must identify when condition 3. arises and when it does we must treat it properly. For example, we could insert an additional column in **Table 6** that adds the pairs of the Kronecker deltas in the column labeled  $\delta_{1,2}$  and takes the sign of this sum. and subtracts it from 1. If the result is 0, condition 3. obtains. The actions for all agents when undetermined, or when determined under conditions 1. and 2., can be handled by multiplying the right-hand side of (43) by the value in the appropriate row (which will be 1). Subtracting the values from 1 assigns 1 to condition 1.. It may then be handled by multiplying the right hand side of (43) by this assigned value and by -1.

Then, importantly, the reversal of sign picks out the position of the strategy with the lower point score. The preceding agent-by-agent computation may be condensed and made general with a straightforward analytic extension to (34) that incorporates the reversal of sign for C agents.

First, we re-cast the initial quenched disorder on the set of strategies attributed to the  $N$  agents in a given game realization as a two-component tensor  $\hat{\Omega} = \{\hat{\Omega}^+, \hat{\Omega}^-\}$ .  $\hat{\Omega}^+$  represents standard (S) agents that adapt as before;  $\hat{\Omega}^-$  represents ‘‘counteradaptive’’ (C) agents that instead select their worst-performing strategies. In our example (26) then, suppose we select at random 3 agents to use the C rule, one each at  $\Omega_{1,2}, \Omega_{2,6}$  and  $\Omega_{7,8}$  and redefine  $\hat{\Omega}$  so that:

$$\hat{\Omega} \equiv \{\hat{\Omega}^+, \hat{\Omega}^-\} = \left( \begin{array}{c} \left( \begin{array}{cccccccc} 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 3 & 2 & 1 & 1 \\ 0 & 0 & 2 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 2 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right) \left( \begin{array}{cccccccc} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right) \end{array} \right) \quad (44)$$

(Thus, the original  $\hat{\Omega}$  of eqn. (26) would be  $\hat{\Omega}^+ + \hat{\Omega}^-$ )

We likewise modify the definition of  $\hat{\Psi}$  :

$$\hat{\Omega} = \{\hat{\Omega}^+, \hat{\Omega}^-\} \quad (45)$$

$$\hat{\Psi} = \{\hat{\Psi}^+, \hat{\Psi}^-\} \quad (46)$$

$$\hat{\Psi}^+ = \frac{1}{2}(\hat{\Omega}^+ + \hat{\Omega}^{+\top}); \quad \hat{\Psi}^- = \frac{1}{2}(\hat{\Omega}^- + \hat{\Omega}^{-\top}) \quad (47)$$

The elements of  $\hat{\Omega}^+$  are the S agents of  $\hat{\Omega}$  as previously defined and discussed. The elements of  $\hat{\Omega}^-$  represent agents that use the C strategy selection rule. We then re-write (34) as :

$$A_D(\mu_i) = \left( \sum_{r=1}^8 \left\{ \left[ (1 - \text{Sgn}[\bar{s}_{\mu_i} \ominus \bar{s}_{\mu_i}]) \circ \hat{\Psi}^+ + (1 + \text{Sgn}[\bar{s}_{\mu_i} \ominus \bar{s}_{\mu_i}]) \circ \hat{\Psi}^- \right] \cdot \hat{a}_1 \right\}_r \right)_{(\text{Mod}[\mu_i - 1, 2^m] + 1)} \quad (48)$$

If we use (48) with the same 3 arbitrarily selected C agents as in **Figure 17**, and with  $\hat{\Omega}^+ + \hat{\Omega}^-$  identical to that used to generate the results of **Figure 17** and **Table 5**, we obtain the satisfying analytic/numeric agreement for agent versus strategy mean per-step change in wealth shown in Table 7.

**Table 7:** Numerical versus Analytical Results of THMG with 3 C agents.

	$\langle \Delta W_{Agent} \rangle$	$\langle \Delta W_{Strategy} \rangle$
<i>Numerical</i>	-0.16	-0.05
<i>Analytical</i>	-0.16	-0.05

The 3 C agents of 31 now perform so well that they significantly raise the overall performance of the system as shown in **Figure 17** and detailed in **Figure 18**. They not only outperform both their own strategies and the other S agents on average, they generate net positive gain. The *hypothetical* outperformance of unused relative to used strategies in

the MG was first observed in [13]. But the explicit generation of positive results, by agents simply deploying their unused strategies (without privileging) has here been tested. (In the case of  $S = 2$ , “unused” are by definition the “worst-performing”.)

### A9.1.3 Analytic methods for the THMAJG and TH\$G

The Markov transition matrices for the THMAJ and TH\$G (for any proportion of C agents) may be obtained by straightforward extensions to the methods presented above.

We know that for the MG, MAJG and \$G respectively, the gain (loss) accrued by agents are given by the following equations

$$g_i^{min}(t) = -a_i(t)A(t) \text{ or } g_i^{min}(t) = -Sgn[a_i(t)A(t)] \quad (49)$$

$$g_i^{maj}(t) = +a_i(t)A(t) \text{ or } g_i^{maj}(t) = +Sgn[a_i(t)A(t)] \quad (50)$$

$$g_i^s(t) = +a_i(t-1)A(t) \text{ or } g_i^{min}(t) = -Sgn[a_i(t)A(t)] \quad (51)$$

Extending eqns (27) and (28), the change in wealth (point gain or loss) associated with each of the strategies for the allowed transitions among the 4 histories) at any time  $t$  for each of the three games is then (note sign and subscript differences corresponding to eqns (49)-(51)):

$$\delta\bar{S}_{\mu(t),\mu(t-1)}^{min} = +(\hat{a}^\top)_{\mu(t)} \times \{2Mod[\mu(t-1),2] - 1\} \quad (52)$$

$$\delta\bar{S}_{\mu(t),\mu(t-1)}^{maj} = -(\hat{a}^\top)_{\mu(t)} \times \{2Mod[\mu(t-1),2] - 1\} \quad (53)$$

$$\delta\bar{S}_{\mu(t),\mu(t-1)}^s = -(\hat{a}^\top)_{\mu(t-1)} \times \{2Mod[\mu(t-1),2] - 1\} \quad (54)$$

Once again,  $Mod[x,y]$  is “x modulo y”;  $\mu(t)$  and  $\mu(t-1)$  label each of the 4 histories  $\{00,01,10,11\}$  hence take on one of values  $\{1,2,3,4\}$ . Equations (52)-(54) pick out from (27) the correct change in wealth over a single step since the strategies are ordered in symmetrical sequence.

With  $game \in \{min, maj, \$\}$ , the change in points associated with each strategy for each of the allowed transitions along all the  $\tau$  histories (i.e., along the path  $\mu_t \equiv \mu(t-1) \rightarrow \mu(t)$ , accounting for the last  $\tau$  time steps used to score the strategies) is:

$$\bar{S}_{\mu_t}^{game} \equiv \sum_{i=0}^{\tau-1} \delta\bar{S}_{\mu(t-i),\mu(t-i-1)}^{game} \quad (55)$$

Eqn (55) accounts for the change in points along path  $\mu_t$  by summing them over all transitions on the path. For example, with  $m = 2$  and  $\tau = 1$ , the strategy scores are kept for only a single time-step and  $\tau - 1 = 0$  so the sum in (55) reduces to a single term. (55) in matrix form therefore reduces to the score (-by-strategy) vector:

$$\bar{S}_{\mu_t}^{game} \equiv \delta\bar{S}_{\mu(t),\mu(t-1)}^{game} \quad (56)$$

or, listing the results for all 8 path histories:

$$\hat{\mathbf{s}}_{\mu}^{game} = \delta \hat{\mathbf{S}}^{game} \quad (57)$$

$\delta \hat{\mathbf{S}}^{game}$  is an  $8 \times 8$  matrix that can be read as a lookup table. It denotes the change in points accumulated over  $\tau = 1$  time steps for each of the 8 strategies over each of the 8 path-histories.

Instead of computing  $A^{game}(t)$ , we compute  $A^{game}(\mu_t)$ . Then for each of the  $2^{m+\tau} = 8$  possible  $\mu_t$ ,  $A(\mu_t)$  is composed of a subset of wholly determined agent votes and a subset of undetermined agents whose votes must be determined by a coin toss:

$$A^{game}(\mu_t) = A_D^{game}(\mu_t) + A_U^{game}(\mu_t) \quad (58)$$

Some agents are undetermined at time  $t$  because their strategies have the same score and the tie has to be broken with a coin toss.  $A_U^{game}(\mu_t)$  is a random variable characterized by the binomial distribution. Its actual value varies with the number of undetermined agents  $N_U^{game}$ . For each of the three games, this number can be denoted as:

$$N_U^{min}(\mu_t) = \left\{ \left( 1 - \left[ (\hat{a}^T)_{(\text{Mod}[\mu_t-1,4]+1)} \otimes_{\delta} (\hat{a}^T)_{(\text{Mod}[\mu_t-1,4]+1)} \right] \right) \circ (\bar{s}_{\mu_t}^{min} \otimes_{\delta} \bar{s}_{\mu_t}^{min}) \circ \hat{\Omega} \right\}_{(\text{Mod}[\mu_t-1,2^m]+1)} \quad (59)$$

$$N_U^{maj}(\mu_t) = \left\{ \left( 1 - \left[ (\hat{a}^T)_{(\text{Mod}[\mu_t-1,4]+1)} \otimes_{\delta} (\hat{a}^T)_{(\text{Mod}[\mu_t-1,4]+1)} \right] \right) \circ (\bar{s}_{\mu_t}^{maj} \otimes_{\delta} \bar{s}_{\mu_t}^{maj}) \circ \hat{\Omega} \right\}_{(\text{Mod}[\mu_t-1,2^m]+1)} \quad (60)$$

$$N_U^s(\mu_t) = \left\{ \left( 1 - \left[ (\hat{a}^T)_{(\text{Mod}[\mu_{t-1}-1,4]+1)} \otimes_{\delta} (\hat{a}^T)_{(\text{Mod}[\mu_{t-1}-1,4]+1)} \right] \right) \circ (\bar{s}_{\mu_t}^s \otimes_{\delta} \bar{s}_{\mu_t}^s) \circ \hat{\Omega} \right\}_{(\text{Mod}[\mu_{t-1}-1,2^m]+1)} \quad (61)$$

Note that (59) and (60) are structurally identical while (61) differs from these in that the indices on  $(\hat{a}^T)$  and on the entire expression reference path-histories  $\mu_{t-1}$  rather than  $\mu_t$ , reflecting the one-step time-lag in the payoff for the \$G. The sign differences on  $(\hat{a}^T)$  vanish again because of the products. Similarly:

$$A_D^{min}(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ \left( 1 - \text{Sgn} \left[ \bar{s}_{\mu_t}^{min} \ominus \bar{s}_{\mu_t}^{min} \right] \right) \circ \hat{\Psi} \right] \cdot \hat{a} \right\}_r \right)_{(\text{Mod}[\mu_t-1,2^m]+1)} \quad (62)$$

$$A_D^{maj}(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ \left( 1 - \text{Sgn} \left[ \bar{s}_{\mu_t}^{maj} \ominus \bar{s}_{\mu_t}^{maj} \right] \right) \circ \hat{\Psi} \right] \cdot (-\hat{a}) \right\}_r \right)_{(\text{Mod}[\mu_t - 1, 2^m] + 1)} \quad (63)$$

$$A_D^s(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ \left( 1 - \text{Sgn} \left[ \bar{s}_{\mu_t}^s \ominus \bar{s}_{\mu_t}^s \right] \right) \circ \hat{\Psi} \right] \cdot (-\hat{a}) \right\}_r \right)_{(\text{Mod}[\mu_t - 1, 2^m] + 1)} \quad (64)$$

If C agents are present, these expressions must be modified accordingly, i.e.:

$$A_D^{min}(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ \left( 1 - \text{Sgn} \left[ \bar{s}_{\mu_t}^{min} \ominus \bar{s}_{\mu_t}^{min} \right] \right) \circ \hat{\Psi}^+ + \left( 1 + \text{Sgn} \left[ \bar{s}_{\mu_t}^{min} \ominus \bar{s}_{\mu_t}^{min} \right] \right) \circ \hat{\Psi}^- \right] \cdot \hat{a} \right\}_r \right)_{(\text{Mod}[\mu_t - 1, 2^m] + 1)} \quad (65)$$

$$A_D^{maj}(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ \left( 1 - \text{Sgn} \left[ \bar{s}_{\mu_t}^{maj} \ominus \bar{s}_{\mu_t}^{maj} \right] \right) \circ \hat{\Psi}^+ + \left( 1 + \text{Sgn} \left[ \bar{s}_{\mu_t}^{maj} \ominus \bar{s}_{\mu_t}^{maj} \right] \right) \circ \hat{\Psi}^- \right] \cdot (-\hat{a}) \right\}_r \right)_{(\text{Mod}[\mu_t - 1, 2^m] + 1)} \quad (66)$$

$$A_D^s(\mu_t) = \left( \sum_{r=1}^8 \left\{ \left[ \left( 1 - \text{Sgn} \left[ \bar{s}_{\mu_t}^s \ominus \bar{s}_{\mu_t}^s \right] \right) \circ \hat{\Psi}^+ + \left( 1 + \text{Sgn} \left[ \bar{s}_{\mu_t}^s \ominus \bar{s}_{\mu_t}^s \right] \right) \circ \hat{\Psi}^- \right] \cdot (-\hat{a}) \right\}_r \right)_{(\text{Mod}[\mu_t - 1, 2^m] + 1)} \quad (67)$$

Once again, the index  $r$  refers to strategies in the RSS and summation over  $r$  transforms the base-ten code for  $\mu_t$  into  $\{1, 2, 3, 4, 1, 2, 3, 4\}$ . Selection of the proper number is indicated by the subscript expression on the entire right-hand sides of eqns (62)-(64). This expression yields an index number, i.e., selection takes place  $1 + \text{Modulo } 4$  with respect to the value of  $(\mu_t - 1)$  for the THMG and with respect to the value of  $(\mu_{t-1} - 1)$  for the THMAJG and TH\$G.

Equations (58) through (67) yield the history-dependent  $(m + \tau) \times (m + \tau)$  matrix  $\hat{\mathbf{T}}$  with elements  $T_{\mu_t, \mu_{t-1}}$ , representing the 16 allowed probabilities of transitions between the two sets of 8 path-histories  $\mu_t$  and  $\mu_{t-1}$  (the game-type superscripts on  $A_D$  and  $N_U$  are understood in context):

$$T_{\mu_t, \mu_{t-1}}^{game} = \Gamma_{\mu_t, \mu_{t-1}}^{game} \times \sum_{x=0}^{N_U^{game}(\mu_t)} \left\{ \binom{N_U^{game}(\mu_t)}{x} \left( \frac{1}{2} \right)^{N_U^{game}(\mu_t)} \times \delta \left[ \text{Sgn} \left( A_D^{game}(\mu_t) + 2x - N_U^{game}(\mu_t) \right) + \left( 2 \text{Mod} \{ \mu_{t-1}, 2 \} - 1 \right) \right] \right\} \quad (68)$$

Given a specific  $\hat{\Omega}$ ,

$$\langle A^{game}(\mu_t) \rangle = A_D^{game}(\mu_t) \quad \forall \mu_t \quad (69)$$

$\bar{\kappa}$ , the list of the number of times each strategy is represented in  $\hat{\Omega}$ , is identical in all three games of course, so that:

$$\langle \Delta W_{Agent} \rangle = \pm \frac{1}{N} Abs(\bar{A}_D) \cdot \bar{\mu} \quad (70)$$

with the minus sign for the MG, otherwise not, i.e., the awarding of points is the negative of the direction of the vote imbalance for the MG, and in the direction of the imbalance in the MAJG and \$G. And

$$\langle \Delta W_{Strategy} \rangle = \frac{1}{2N} (\hat{\mathbf{s}}_\mu \cdot \bar{\kappa}) \cdot \bar{\mu} \quad (71)$$

#### A9.1.4 Analytic methods for the cycle decomposition of binary series

##### A9.1.4.1. Cycle Decomposition of TH games

In [90], the MG proper (not the THMG) is represented as a deterministic system with perturbations due to tie-breaks. The underlying deterministic system is likewise characterized as an Eulerian path on a de Bruijn graph even though the system's path is not Markovian. As described above, a close approximation of the underlying determinism in the THMG, THMAJG and TH\$G may be achieved simply by rounding every element of  $\hat{\mathbf{T}}^{game}$  (eqn. (68) to either 0 or 1. When  $N$  is relatively large,  $\hat{\mathbf{T}}$  is relatively unlikely to have elements equal to  $\frac{1}{2}$  and the rounded  $\hat{\mathbf{T}}$  ( $\equiv \hat{\mathbf{T}}_0$ ) well-characterizes the system, even for small  $m$  and  $\tau$ . But for many interesting values of  $N$  (i.e.,  $N = 31$ ), elements equal to  $\frac{1}{2}$  are common in  $\hat{\mathbf{T}}^{game}$  and there are different (binomially-distributed)  $\hat{\mathbf{T}}_0^{game}$ .

Note that elements rounded to 0 are eliminated from  $\hat{\mathbf{T}}^{game}$ . If no values other than 0,1 are allowed,  $\hat{\mathbf{T}}_0^{game}$  will therefore either represent one cycle from the De Bruijn graph, or if more than one, the cycles will be disjoint. If elements exactly equal to  $\frac{1}{2}$  are left unchanged, then for every such element there will be a cycle. The cycles will not necessarily be disjoint and the dynamic represented by  $\hat{\mathbf{T}}_0^{game}$  may then be viewed as two (or more) deterministic orbits in state-space with a 0.5 probability of switching from one to another upon exiting any the state with probability  $\frac{1}{2}$ . Similarly,  $\hat{\mathbf{T}}^{game}$  unmodified (with arbitrary elements  $T_{i,j}^{game} \in [0,1]$ ) can be thought of as a (probabilistically) weighted superposition of all possible deterministic cycles present in  $\hat{\mathbf{T}}^{game}$ . In general, the cycle with the greatest weight provides a good first-order approximation to  $\hat{\mathbf{T}}^{game}$ .

The upper limit on the number of terms (cycles) required to decompose a finite-sized transition matrix is also finite but grows super-exponentially with  $m + \tau$ . The accuracy of the approximation grows with the number of cycles included. But more important for our purposes is the fact that deterministic cycles are much easier to study individually than their composite, and characteristics emerge in the aggregate that shed light on features of the latter.

#### A9.1.4.2.Extraction of Cycles from a Binary Series

We require first a method of extracting cycles from a given binary data series. (Note that if we create a cycle decomposition of an arbitrary series, we are recasting it in a form that treats the series as Markovian, whether it is or is not, with a transition matrix of dimensionality equal to some memory scale,  $m_s$ .) Consider an arbitrary finite binary history (not necessarily derived from a TH  $\hat{\mathbf{T}}^{game}$ , but a purely arbitrary history selected at random):

$$\{1,0,0,1,0,0,0,1,1,1,0,1,1,1,0,1,0,0,1,1,0,0,0,1,1,0,1,1,0,0\dots\} \quad (72)$$

This is equivalent to the following path-history (binary converted to digital) when paths are defined as having length  $m_s = 3$  (i.e., were eqn. (72) in fact a TH game history, then “ $m + \tau$ ” would likewise = 3):

$$\begin{aligned} &\{(1,0,0),(0,0,1),(0,1,0),(1,0,0),(0,0,0),(0,0,1),(0,1,1),(1,1,1), \\ &(1,1,0),(1,0,1),(0,1,1),(1,1,1),(1,1,0),(1,0,1),(0,1,0),(1,0,0), \\ &(0,0,1),(0,1,1),(1,1,0),(1,0,0),(0,0,0),(0,0,1),(0,1,1),(1,1,0), \\ &(1,0,1),(0,1,1),(1,1,0),(1,0,0)\}\dots \rightarrow \\ &\{5,2,3,5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5\dots\} \end{aligned} \quad (73)$$

A cycle consists of any sequence of digital states of length  $\leq 9$  ( $= 2^{m_s} + 1$ ) that begins and ends with the same digit, and within which no digit is otherwise found more than once. The first cycle in this series happens to coincide with the first four decimal digits and is  $\{5,2,3,5\}$ . Extract and tabulate this cycle and replace it with 5, the repeated initial digit, i.e.,

$$\begin{aligned} &\{5,2,3,5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5\dots\} \rightarrow \\ &5, \{5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5\dots\} \end{aligned} \quad (74)$$

The next cycle is  $\{4,8,7,6,4\}$ . Extract this cycle and replace it with 4. This leaves  $5,1,2,4,8,7,6,3,5,2\dots$ . Continue extracting and tabulating until no cycles are left. Ignore any remaining digits. Then repeat the process seven more times (i.e., for a total of  $2^{m_s}$  times), each time dropping one more digit from the beginning of the series. For the beginning of the series in (73) **Table 8** shows the first four of the seven extractions.

**Table 8:** Extraction of Cycles from a Series

Remaining Sequence	Extracted Cycles
(5,2,3,5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (5,2,3,5)
(5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (4,8,7,6,4)
(5,1,2,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (5,1,2,4,8,7,6,3,5)
(5,2,4,7,5,1,2,4,7,6,4,7,5)	– (5,2,4,7,5)
(5,1,2,4,7,6,4,7,5)	– (4,7,6,4)
(5,1,2,4,7,5)	– (5,1,2,4,7,5)
(5)	
(2,3,5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (2,3,5,1,2)
(2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (4,8,7,6,4)
(2,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (2,4,8,7,6,3,5,2)
(2,4,7,5,1,2,4,7,6,4,7,5)	– (2,4,7,5,1,2)
(2,4,7,6,4,7,5)	– (4,7,6,4)
(2,4,7,5)	
(3,5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (4,8,7,6,4)
(3,5,1,2,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (3,5,1,2,4,8,7,6,3)
(3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (5,2,4,7,5)
(3,5,1,2,4,7,6,4,7,5)	– (4,7,6,4)
(3,5,1,2,4,7,5)	– (5,1,2,4,7,5)
(3,5)	
(5,1,2,4,8,7,6,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (4,8,7,6,4)
(5,1,2,4,8,7,6,3,5,2,4,7,5,1,2,4,7,6,4,7,5)	– (2,4,8,7,6,3,5,2)
(5,1,2,4,7,5,1,2,4,7,6,4,7,5)	– (1,2,4,7,5,1)
(5,1,2,4,7,6,4,7,5)	– (4,7,6,4)
(5,1,2,4,7,5)	– (5,1,2,4,7,5)
(5)	

Categorize and count the number of times each cycle appears and compute its proportion as a fraction all cycle types. Adjust this proportion by treating cyclic permutations of a cycle as the same type and combine them into one category. Results for the partial series (four of eight extractions) in (73) are shown in **Table 9**. The proportional representation of a cycle in the *complete* extraction constitutes its “weight”. **Table 9:** Computation of Cycle Weights

Cycle	Raw Count	Adj. Count	Weight (Proportion)
(4,7,6,4)	4	4	0.190
(5,2,3,5)	1	1	0.048
(2,3,5,1,2)	1	1	0.048
(4,8,7,6,4)	4	4	0.190
(5,2,4,7,5)	2	2	0.095
(1,2,4,7,5,1)	1	5	0.238
(2,4,7,5,1,2)	1	–	–
(5,1,2,4,7,5)	3	–	–
(2,4,8,7,6,3,5,2)	2	2	0.095
(3,5,1,2,4,8,7,6,3)	1	2	0.095
(5,1,2,4,8,7,6,3,5)	1	–	–

The cycles in column 1 of **Table 9** can be represented as adjacency matrices within the space of all possible cycles on the binary deBruijn graph of order 3 as shown in **Figure 29**. Call this column vector  $\bar{\mathbf{J}}$ .

$$\left\{ \begin{array}{c} \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{array} \right) \end{array} \right\}$$

$$\left\{ \begin{array}{c} \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right), \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{array} \right), \left( \begin{array}{cccccccc} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{array} \right) \end{array} \right\}$$

**Figure 29:**  $\bar{\mathbf{J}}$  = the 8 non-zero-weight cycles of (73) in adjacency matrix form.

Call column 4 of **Table 9**  $\bar{\omega}$ , the weights of every cycle. The dot product  $\bar{\mathbf{J}} \cdot \bar{\omega}$  yields (75), the transition matrix for the series in (73):

$$\bar{\mathbf{J}} \cdot \bar{\omega} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0.38 & 0 & 0 & 0 \\ 0.38 & 0 & 0 & 0 & 0.24 & 0 & 0 & 0 \\ 0 & 0.10 & 0 & 0 & 0 & 0.19 & 0 & 0 \\ 0 & 0.52 & 0 & 0 & 0 & 0.38 & 0 & 0 \\ 0 & 0 & 0.29 & 0 & 0 & 0 & 0.33 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.57 & 0 \\ 0 & 0 & 0 & 0.52 & 0 & 0 & 0 & 0.38 \\ 0 & 0 & 0 & 0.38 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (75)$$

In this case the resulting matrix has complementary entries that do not add to 1 because the series is so short. But as proven formally in [113, 114], in the limit of a sufficiently long such series, the transition values obtained by a complete extraction are accurate and arbitrarily precise, i.e., for every pair of non-zero entries  $\{J_{k_i,j}, J_{k_{i+1},j}\}$  representing an allowed transition in the  $R \times R$  transition matrix  $J_k$  (where  $\{J_1, J_2, \dots, J_k, \dots, J_R\} \equiv \bar{\mathbf{J}}$ ),  $\lim_{L \rightarrow \infty} (J_{k_i,j} + J_{k_{i+1},j}) = 1$ , with  $L$  the series' length. In any series of such finite length, the complementary entries  $\{J_{k_i,j}, J_{k_{i+1},j}\}$  may be normalized so their sum equals 1 to convert actual transition frequencies to approximate transition probabilities:

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0.38 & 0 & 0 & 0 \\ 0.38 & 0 & 0 & 0 & 0.24 & 0 & 0 & 0 \\ 0 & 0.10 & 0 & 0 & 0 & 0.19 & 0 & 0 \\ 0 & 0.52 & 0 & 0 & 0 & 0.38 & 0 & 0 \\ 0 & 0 & 0.29 & 0 & 0 & 0 & 0.33 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.57 & 0 \\ 0 & 0 & 0 & 0.52 & 0 & 0 & 0 & 0.38 \\ 0 & 0 & 0 & 0.38 & 0 & 0 & 0 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 0 & 0 & 0.61 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0.39 & 0 & 0 & 0 \\ 0 & 0.16 & 0 & 0 & 0 & 0.33 & 0 & 0 \\ 0 & 0.84 & 0 & 0 & 0 & 0.67 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0.37 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.63 & 0 \\ 0 & 0 & 0 & 0.58 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0.42 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (76)$$

Note, too, that the weightiest cycle treated as having weight = 1 provides a 0<sup>th</sup> order (and therefore wholly deterministic) approximation for the series. (The closer its actual weight is to 1 the more accurate the approximation). In the above instance:

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (77)$$

We may create higher order (no longer wholly deterministic) approximations by adding the second weightiest matrix, third weightiest, etc. (and again normalizing the entries to 1).

now demonstrate the use of this decomposition in the series generated by the THMG by analytically decomposing the transition matrix  $\hat{\mathbf{T}}^{min}$  derived from an initial quenched disorder tensor  $\hat{\Omega}$ , in particular the tensor of eqn. (26). Identical methods will decompose  $\hat{\mathbf{T}}^{maj}$  and  $\hat{\mathbf{T}}^s$ . We will later use the above empirical method to approximate series generated by the (non-TH) MG, MAJG and \$G as well as real-world series. Note that the decomposition of a given finite series produces an approximation of a transition matrix whereas an actual transition matrix generates a series of unlimited length.

#### A9.1.4.3 Analytic Form of a Cycle Decomposition

The decomposition of a known transition matrix into its weighted cycle structure was first proposed in ref. [113]. For any stochastic matrix  $\hat{\mathbf{T}}$  (of which the transition matrix for THMG is an instance of size  $2^{m+\tau} \times 2^{m+\tau}$ ), and the set  $C$  of all possible directed cycles  $J^c$  of the adjacency matrix  $\hat{\Gamma}$  of  $\hat{\mathbf{T}}$ , it can be shown that [115]:

$$T_{i,j} = \frac{\sum_c \omega_c J_{i,j}^c}{\sum_c \omega_c J_i^c}, \quad 1 \leq c \leq \dim(C), \quad 1 \leq i, j \leq 2^{m+\tau} \quad (78)$$

(recall that the  $\omega_c$  are the cycle weights) where

$$J_i^c = \sum_{j=1}^{2^{m+\tau}} J_{i,j}^c = \sum_{i=1}^{2^{m+\tau}} J_{i,j}^c. \quad (79)$$

$J_{i,j}^c = 1$  if  $(i,j)$  is a directed edge of  $J^c$ , 0 otherwise. In addition,

$$\pi_i = \sum_c \omega_c J_i^c \quad (80)$$

so that

$$\pi_i T_{i,j} = \sum_c \omega_c J_{i,j}^c; \quad \hat{\mathbf{T}} \cdot \bar{\pi} = \bar{\pi} = \sum_c J^c \cdot \bar{\omega} \quad (81)$$

where the  $\pi_i$  are the  $2^{m+\tau}$  steady-state probabilities derivable from  $\hat{\mathbf{T}}$ . (81) represents  $\dim(C)$  simultaneous matrix equations to be solved for the  $\omega_c$ .

For the  $\hat{\mathbf{T}}^{min}$  of equation (26) (and  $\hat{\Gamma}$  of equation (35)) which has only 4 cycles (and only  $S$  agents),

$$\bar{\mathbf{J}} \equiv \{J^1, J^2, J^3, J^4\} = \quad (82)$$

$$\left\{ \begin{array}{c} \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \end{array} \right\}$$

corresponding to  $\{(4,7,6,4), (4,8,7,6,4), (1,2,4,7,5,1), (1,2,4,8,7,5,1)\}$  and their cyclic permutations. (In general, a random binary series converted to paths of length 3 can, and if long enough, will have 19 unique cycles. The truncated series (73) we used as a sample contains only the 8 cycles shown in adjacency matrix form in **Figure 29**, of which we showed the extraction of 4 in detail. In the present example derived from a THMG tensor, there are only 4 extractable cycles in toto reflecting the high degree of determinism in the time series generated by  $\hat{\Omega}$ .) From (79):

$$\hat{\mathbf{J}} \equiv \begin{pmatrix} J_1^1 & \cdots & J_1^4 \\ \vdots & \ddots & \vdots \\ J_8^1 & \cdots & J_8^4 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{pmatrix} \quad (83)$$

Solving (81) (equivalently,  $\bar{\pi} = \hat{\mathbf{J}} \cdot \bar{\omega}$ ), for the  $\omega_c$  using the method of cofactor expansion, we obtain the approximate values shown in **Table 10**.

**Table 10: Cycle Weights**

	Cycle	Weight
$\omega_1$	(4,7,6,4)	0.072
$\omega_2$	(4,8,7,6,4)	0.428
$\omega_3$	(1,2,4,7,5,1)	0.072
$\omega_4$	(1,2,4,8,7,5,1)	0.428

To check on the correctness of the solution for equation (81), equivalently  $\bar{\mathbf{J}} \cdot \bar{\omega} \equiv \hat{\mathbf{T}}_{\text{cyc}}$ :

$$\hat{\mathbf{T}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.1445 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0.8555 & 0 & 0 & 0 & 0 \end{pmatrix}; \hat{\mathbf{T}}_{\text{cyc}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.1445 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0.8555 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (84)$$

Note the one pair of entries that differ between the two matrices in (84). The matrix equation (81) may be solved by a number of methods (e.g., division-free row reduction, one-step row reduction; we used cofactor expansion). All typically also have difficulty with probabilities = 0.5. We see that this has happened here. Formally, the convergence of  $\hat{\mathbf{T}}_{\text{cyc}}$  to  $\hat{\mathbf{T}}$  is in general weak and “almost sure” [116]. We will see by comparison to numerical simulation that the small error thus introduced has no effect on our use of the decomposition. The difference between the two matrices in (84) provides an illustration of one reason why this has no effect: The original matrix  $\hat{\mathbf{T}}$  on the left of (84) shows an equiprobable transition from path 3 [from(0,1,0)] to path 5 or to path 6 [to (1,0,0) or to(1,0,1)]. But looking at row 3 of either matrix, we see that state 3 is inaccessible altogether (no transitions from any other state to it), hence no transition *from* state 3 is possible. This is reflected in  $\hat{\mathbf{T}}_{\text{cyc}}$  (on the right) by the two 0 zero entries in place of 0.5. Thus, the two matrices, being at equilibrium, are equivalent.

#### A9.1.4.4 Comparison to Numerical Simulation

Each cycle in the decomposition is a purely deterministic, discrete periodic process. As noted before, the elements of  $\hat{\mathbf{T}}_{\text{cyc}}$  represent the probabilities of transitions between cycles rather than (equivalently in  $\hat{\mathbf{T}}$ ) transitions between path histories. All the dynamics of the full THMG can therefore be derived from the deterministic dynamics of the THMG around just one of every cycle, properly composed and weighted. For example, if  $\hat{\mathbf{T}}_{\text{cyc}}$  were composed of just two cycles of respective lengths 5 and 4, the total change in agent and strategy wealth over these will be computed over twenty (the lowest common multiple) steps—4 of the first cycle times 5 of the second, each times its appropriate

weight, and then summed. The net per-step gain (or loss) in wealth for both agents and for their underlying strategies is easily computable around any cycle using (27)-(31). For the  $\hat{\mathbf{T}}$  and  $\hat{\mathbf{T}}_{cyc}$  of (84), with 31 agents, overall results are shown in

**Table 11.**

**Table 11:** Elements of the Cycle Decomposition of the THMG with S and C agents

Cycle	Weight	$\Delta W$
<b>(4,7,6,4)</b>		
$\Delta W_{agent}$	0.08(4)	-0.10
$\Delta W_{strategy}$		-0.09
<b>(4,8,7,6,4)</b>		
$\Delta W_{agent}$	0.85(4)	-0.23
$\Delta W_{strategy}$		-0.06
<b>(1,2,4,7,5,1)</b>		
$\Delta W_{agent}$	0.06(3)	-0.11
$\Delta W_{strategy}$		-0.04
<b>(1,2,4,8,7,5,1)</b>		
$\Delta W_{agent}$	0.00(0)	-0.19
$\Delta W_{strategy}$		-0.03
<b>Weighted Composite</b>		
$\Delta W_{agent}$	<b>1.00(0)</b>	<b>-0.21</b>
$\Delta W_{strategy}$		<b>-0.06</b>

Comparison to numerical and the standard analytic results are in **able 12** showing the close agreement among methods. (“Analytic” results here refer to the mean per-step change in wealth averaged over all agents and all strategies respectively, i.e.,  $\Delta W_{agent}, \Delta W_{strategy}$  as detailed in section 3.2 as well as in [43, 84]. Numeric results refer to numeric simulations of the THMG. In both cases we refer here to a single initial quenched disorder matrix. As explained in [84], this example is typical with changes in wealth close to the average over many initial quenched disorders.)

**able 12:** Comparison of cycle decomposition to analytic and numeric results

	$\Delta W_{agent}$	$\Delta W_{strategy}$
<i>Numeric</i>	-0.20	-0.07
<i>Analytic</i>	-0.22	-0.06
<i>Cycle</i>	-0.21	-0.06

## A9.2 Analytic expression of the general Parrondo effect

Consider  $N > 1$   $s$ -state Markov games  $G_i$ ,  $i \in \{1, 2, \dots, N\}$ , and their  $N$   $s \times s$  transition matrices,  $\hat{\mathbf{M}}^{(i)}$ . For every  $\hat{\mathbf{M}}^{(i)}$ , denote its vector of  $s$  winning probabilities conditional on each of the  $s$  states as  $\vec{\mathbf{p}}^{(i)} = \{p_1^{(i)}, p_2^{(i)}, \dots, p_s^{(i)}\}$  and its steady-state equilibrium distribution vector as  $\vec{\Pi}^{(i)} = \{\pi_1^{(i)}, \pi_2^{(i)}, \dots, \pi_s^{(i)}\}$ . For each game, the steady-state probability of winning is therefore  $P_{win}^{(i)} = \vec{\mathbf{p}}^{(i)} \cdot \vec{\Pi}^{(i)}$ . Consider also a sequence of randomly alternating  $G_i$  with individual time-averaged proportion of play  $\gamma_i \in [0, 1]$ ,  $\sum_{i=1}^N \gamma_i = 1$ . The transition matrix for the combined sequence of games is the convex linear combination

$\hat{\mathbf{M}}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} \equiv \sum_{i=1}^N \gamma_i \hat{\mathbf{M}}^{(i)}$  with conditional winning probability vector  $\bar{\mathbf{p}}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} = \sum_{i=1}^n \gamma_i \bar{\mathbf{p}}^{(i)}$  and steady-state probability vector  $\bar{\Pi}^{(\gamma_1, \gamma_2, \dots, \gamma_N)}$  (which is a complex nonlinear mixture of the  $\bar{\Pi}^{(i)}$ 's). The steady-state probability of winning for the combined game is therefore

$$P_{win}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} = \bar{\mathbf{p}}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} \cdot \bar{\Pi}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} \quad (85)$$

A PE occurs whenever (and in general it is the case that)

$$\sum_{i=1}^N \gamma_i P_{win}^{(i)} \neq P_{win}^{(\gamma_1, \gamma_2, \dots, \gamma_N)}, \text{ i.e., } \sum_{i=1}^N \gamma_i \bar{\mathbf{p}}^{(i)} \cdot \bar{\Pi}^{(i)} \neq \bar{\mathbf{p}}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} \cdot \bar{\Pi}^{(\gamma_1, \gamma_2, \dots, \gamma_N)} \quad (86)$$

hence the PE, or “paradox”, when the left hand sides of (86) are less than zero and the right-hand sides greater.

## **Chapter A10. Appendix: anti-persistence using an equity-ranking predictor**

A10 is an appendix chapter that describes an alternate prediction method based on neural networks and building on the so-called “Value Line anomaly.” It forms the basis for a working paper. The method illustrates a certain degree of predictive control along with many instances of the “illusion of control”. Anti-persistence reappears as an accompaniment to the overall success of the predictor, yielding high levels of volatility.

### ***A10.1 Introduction***

Using an artificial neural network (ANN), a fixed universe of ~1500 equities from the Value Line index are rank-ordered by their predicted price changes over the next quarter. Inputs to the network consist only of the ten prior quarterly percentage changes in price and in earnings for each equity (by quarter, not accumulated), converted to a relative rank scaled around zero. Thirty simulated portfolios are constructed respectively of the 10, 20, ..., and 100 top ranking equities (long portfolios), the 10, 20, ..., 100 bottom ranking equities (short portfolios) and their hedged sets (long-short portfolios). In a 29-quarter simulation from the end of the third quarter of 1994 through the fourth quarter of 2001 that duplicates real-world trading of the same method employed during 2002, all portfolios are held fixed for one quarter. Results are compared to the S&P 500, the Value Line universe itself, trading the universe of equities using the proprietary “Value Line Ranking System” (to which this method is in some ways similar), and to a Martingale method of ranking the same equities. The cumulative returns generated by the network predictor significantly exceed those generated by the S&P 500, the overall universe, the Martingale and Value Line prediction methods and are not eroded by trading costs. The ANN shows significantly positive Jensen’s alpha. All three active trading methods result in very high levels of volatility. But the network method exhibits a distinct kind of volatility: Though overall it does the best job of segregating equities in advance into those that will rise and those that will fall relative to one another, there are many quarters when it does not merely fail, but rather “inverts”: It disproportionately predicts an inverse rank ordering and therefore generates unusually large losses in those quarters. The same phenomenon occurs, but to a greater degree, with the VL system itself and with a one-step Martingale predictor. An examination of the quarter to quarter performance of the actual and predicted rankings of the change in equity prices suggests while the network is capturing, after a delay, changes in the market sampled by the equities in the Value Line index (enough to generate substantial gains), it also fails in large measure to keep up with the fluctuating data, leading the predictor to be often “out of phase” with the market. A time series of its global performance thus shows antipersistence. However, its performance is significantly better than a simple one-step Martingale predictor, than the Value Line system itself and than a simple buy and hold strategy, even when transaction costs are accounted for.

### ***A10.2 Background***

A wealth of technical and fundamental information on a representative universe of publicly-traded equities is updated each week, in principle, for every company in the well-

known *Value Line Investment Survey* (VL, “the Survey”) of approximately 1700 primarily-U.S. companies. According to a proprietary and not necessarily static formula known to depend disproportionately upon recent percentage changes in the price of an equity, the recent percentage change in its earnings, and especially on an intermittently-generated and more loosely quantified “earnings surprise factor”[117], the Survey updates and assigns to each equity every week a “Timeliness Rank” from 1 to 5 (In fact, not every equity is updated every week in consequence of a certain “slippage” in the VL system). This rank is a measure of future “price performance.” Stocks assigned a rank of 1 are predicted to experience the largest positive *long-* and *intermediate-*term price change (six to twelve months), 5 the least (or greatest decline).

Because the VL survey appears to provide information on equities with at least some predictive power, it has been the object of a significant amount of academic study, beginning with Shelton in 1969 [118], but most notably Fisher Black’s 1973 paper, “Yes, Virginia, there is hope: Tests of the Value Line ranking system” [119] and a subsequent more detailed dissertation by a student of Black’s at M.I.T. [120]. Other widely-cited studies have been performed again in 1973 (with a focus on risk [121]), and in 1981 (testing aggressive investing using VL ranks [122]), 1985 (testing the inverse effect: How VL rank changes affect stock prices [123]), 1987 (relating VL rank to firm size [124]), 1990 (discussing the implications for the efficient market hypothesis—EMH), 1992 (relating the VL effect to post- announcement earnings changes [125]), 2000 (finding a positive effect even controlling for post-announcement earnings changes [126]) and 2008 (examining the predictive value of other data in the VL Survey apart from the ranking system proper)[127].

Until relatively recently, with the advent of more extensive computerized financial data services, the Value Line survey was one of the most widely-used for professional analysts’ forecasts. It has been shown to provide some of the most accurate forecasts of analysts’ predicted excess return, especially in comparison to other widely used sources (e.g., IBES, S&P) [128]. Fisher Black is reported to have offered the following advice in 1983: “One of the best ways for an investment firm to pilot a portfolio through the vicissitudes of the market would be to fire all the financial analysts, save one, and make that one read Value Line. [129]”

Even though VL defines its rankings to predict long-term price appreciation, most of these studies have concluded that its predictive power is real chiefly for the short-term only and only doubtfully so once transaction costs are included. Nonetheless, given the power of the efficient market hypothesis, the presumed predictive capacity of the VL system is impressive: “In a world with no end of people hawking investment advice, the *Value Line Investment Survey* has captured the imagination of the finance community like few others[126].”

Studies of the VL ranking system consistently demonstrate that it is at least theoretically effective (before trading costs); and occasionally demonstrate that it is practically effective when its predictive range is carefully analyzed and applied such that trading costs do not erode gains[120].

Because the Survey claims to heavily weight “earnings surprises”, and because these are known to affect prices, this factor has been offered as the explanation for how there could

be some predictive power in the ranking system [130, 131]. But this explanation contrasts with VL's own arguments on its behalf, since the system claims to incorporate more than simply earnings surprises in arriving at its rankings (See for example, [117]). Indeed, most rank assignments are made without any earnings surprises. It appears that much of the outside research testing the VL ranking system on the basis of earnings' surprises presumes that the EMH is effectively correct—all *available* information is instantaneously incorporated into the present price of a stock; earlier price and earnings data therefore has no predictive power; an earnings surprise represents new (by definition *unavailable*) information that requires some time to be reflected in the current price; during this time the surprise therefore has (rapidly declining), short-term predictive power.

Nonetheless the question has been hotly debated in the academic community as to whether the VL ranking system as a whole can provide more than a theoretical refutation of the EMH. The semi-strong form of the EMH precludes the ability to profit from VL information as all such information—including VL's forecast of future price appreciation—would already be in principle incorporated into the present price of a security and thus discounted against future gains [132, 133]

A review of the literature makes it clear that while a significant majority of researchers have in fact detected a VL “anomaly” or “enigma”, most find the size of the anomaly likely to be too small to be exploited given transaction costs.[119, 120, 122, 126, 134, 135], regardless of whether it is attributed to the earnings' surprise factor or not.

The VL data and rankings are used both by analysts and traders: Given the relatively modest price of a subscription to the survey, VL could scarcely be a working business, especially for as long as it has been—since the 1960's at least—if its subscription base were only analysts). This leads naturally to another important consideration which has also been the object of study: the possibility of feedback between the weekly release of the VL rankings, especially therefore rank changes, and short-term price changes among equities undergoing rank changes. The EMH could remain in principle true, yet brief departures from it could occur simply because of the market response to rank changes—whether or not these changes accurately reflect underlying fundamentals.

Indeed, there is evidence that analysts will herd significantly (and thus their clients will trade accordingly) based on VL recommendations, thus amplifying any direct effect from VL subscribers. Among analysts who publish newsletters this herding unsurprisingly occurs when signal correlation among them is high. It also occurs, also unsurprisingly, when their measured performance ability is low. Perhaps more surprising is the fact that significant herding based on VL occurs when the analysts' *reputation* is high [136]. The surprise fades when one considers that high reputation may be as much an effect of herding, as independent of underlying fundamentals (ability) as may be prices.

In any event, there is evidence in the literature that a measurable component of the change in certain stock prices may be due a herding effect mediated by VL.

It should be noted, however, that the influence of VL may be declining. First, very little research on VL has taken place after the year 2000. Second, by 2000, popular discussion frequently noted a decline in the belief “on the street” that VL recommendations were still of use. In 2000, for example, *Money* magazine published an article decrying its decline[137]. The article made no academic claims, but may well have both represented and


reinforced a popular belief that made its claims accurate. (In light of the above discussion, whether they were or not would depend of course in large part on how widely the belief remained that they were, or whether a consensus was arising that they weren't.)

In any case, the anomalous performance of the VL ranking system leads naturally to the question as to whether it may not be improved upon.

### A10.3 Review of the VL ranking system and its performance

The distribution of equities in each ranking “bin” is not flat. There are approximately 100 1’s, 300 2’s, 600 3’s, 300 4’s and 100 5’s = 1500. The number of equities characterized as “3” varies the most, with other equities dropping in and out of the survey altogether over time. If the middle of the 3’s is treated as a zero line, then the cumulative distribution of the ranks approximates a coarse-grained hyperbolic tangent: It is thus a natural way of quantizing the (discrete) rank-ordering by predicted price change for every stock, given the natural distribution of percentage price changes around zero, both positive and negative, ignoring the greater asymmetries at the extremes. (Since the smallest number of equities in the two 100-size bins are predicted to experience the largest price changes, respectively up and down, the next-sized bins the next largest price changes and the largest middle bin the least, the distribution of price changes—rather than bin sizes—conforms to a coarse-grained arc-tanh: If the largest positive changes are represented on the left side of a chart, the curve is actually a negative arc-tanh.)

During the period covered by this study (beginning of 1994 – end of 2001, i.e., capturing the run-up to, the peak and the drop-off following the 2000 “bubble”), VL rank 1 (T100, i.e., top 100), rank 5 (B100, i.e., bottom 100, Simulates long investments in Rank 5 stocks) or long-short (H100) hedged portfolios (long Rank 1, short Rank 5) would have performed as shown in **Figure 30**, assuming no transaction costs:



**Figure 30:** Returns from the VL Ranking System 1994-2001

During this period, reasonable costs would have eroded all competitive gains assuming weekly portfolio restructuring. Notice as well that while the slope of the compounded returns for the VL-defined T100 equities (rank 1) is  $> 1$  for four-plus of the six-plus years, the slope of the compounded returns for the VL defined B100 equities (rank 5) is actually up for five+ years. More pertinent are the segments of relative slope, indicated qualitatively by the H100 return curve. This consists of five contiguous segments of positive slope followed by two of sharply negative slope where the ranking system not only fails, it inverts. As noticed by others and at other scales, inversions such as these are typical in the VL system (“In 1983 the average annual returns of stocks ranked four or five at the beginning of 1983 were higher than the corresponding average returns for stocks ranked one, two or three at the beginning of the year” [134]). The erosion of its competitive advantage is not primarily caused, for example, by intermittent “statistical” failures of the B100 portfolio to decline in absolute value and regress toward the market mean—especially easy to do in the face of a generally rising market. Rather, the failures and the erosion are abrupt and rapid—much faster than the gains as may be seen at a glance—and is caused rather by *reversals* in the model’s performance. During these periods it seems, the model is not merely “not working”, it is working “in reverse”, as it were. This observation of the VL failure, while a mere impression from such a coarse example, provides an initial hypothesis for the consideration of “regime-change” in market models of a particular kind, where relative valuation of a group of equities is used as a general method for handling problems of normalization of data.

The VL investment survey includes many other kinds of rankings for its equities (e.g., “safety”) and for broad categories of equities as well (e.g., performance and safety rankings for market “sectors”). Thus, another peculiarity of the VL approach to the financial domain—perhaps underestimated because of its simplicity—is that by combining and recasting so many numerical quantities into ranks, VL indeed performs a crude “renormalization” which—given the amount of noise, uncertainty and error in financial data—may nonetheless be quite effective, even if it was never conceptualized in such formal terms.

That is, VL’s weekly “Timeliness Rank” (a function of relative change of price and relative change in earnings, inter alia) *might* be an effective method of weekly renormalization on a basis if it consisted of (or at least began with) a complete rank-ordering of the universe of equities, instead of a mere 5-bin coarse-grained version. The hand method of ranking originated by Arnold Bernhard has now been extended and computerized by his 86%-still-family-owned firm. The quasi non-linear, perhaps somewhat adaptive algorithms now employed remain proprietary. It is possible that the firm privately creates a complete and effective single-step ranking, but sells only the coarse-grained version to the public. If so, the remarkably poor performance of the VL *Fund* (a mutual fund that is said to employ the ranking system) argues against its efficacy. Or perhaps these proprietary methods involve the creation of the coarse-grained version only, or of something between it and a full rank-ordering. In any case, the coarse-grained version appears to contain enough information to be of academic interest and to provide a large incentive to subscribers, but to be at or just below the borderline of profitability once rigorously scrutinized, especially of late.

For example, we performed a preliminary Monte Carlo simulation on 1400 artificial stock prices undergoing randomized price changes drawn from the actual distribution of the Value Universe of price changes from 1994-2001, and then scattered into ranking partitions that mimic the 100-300-600-300-100 VL structure. This simulation shows that typically, 40-60% of the rank changes reported in the weekly VL survey can be attributed simply to price perturbations near the rank boundaries. Since on many weeks there may be only one or two changes, there are many weeks when no changes are caused by anything other than such perturbations. The Survey itself refers to this phenomenon somewhat misleadingly as the “dynamism” of the ranking system.

Of much greater interest—and pertinent to this study—is the fact that there is an immediate, very strong post-release effect on the price of a stock whose rank has changed once the change is announced. This effect has been noted and exploited by many (subscribers, analysts and their advisees)—and that can be exploited by VL company insiders completely legally in advance. It also has been examined and argued *not* to be caused primarily by a preceding earnings announcement: For instance, Thomas et al. [6] examined the impact of VL timeliness rank changes on stock prices while controlling for contemporaneous earnings releases, and found that the market response is consistent with increased liquidity in the shares.

Furthermore, our Monte Carlo simulation shows that stocks occasionally can change even two ranks stochastically because of the “tanh”-like distribution of rank bin sizes. Indeed, the distribution of the bin size ensures that these “meaningless” large events both happen and that they are disproportionately more likely to occur as moves both in and out of ranks 1 and 5—the ones that in turn have the largest post-release impact on the market they are meant to predict.

Thus, if there is indeed any genuine information contained in as coarse-grained a ranking system as that publicly available in the VL investment survey, it makes sense to attempt to create a more fine-grained version to extract it. The long history of the VL system with its tantalizing successes and failures; its longstanding and successful use of various ranking methods as naïve ways to handle scaling and normalization problems with financial data; the specific successes and kinds of failures of its “Timeliness Ranks” as a method for the short-term prediction of *relative* stock price changes; the observation of possible “phase” or “regime” changes causing catastrophic failure even in so crude a model as this; the tanh-like distribution of the bin sizes being suggestive of a coarse-grained re-normalization of a finer-grained ranking; the fact that even though the VL system was developed nearly half a century ago, its core “inputs”—recent percent change in price and recent percent change in earnings—continue to dominate: All of these ideas suggest that it should be possible to create *de novo* a complete ranking of equities from the VL universe.

All of the above suggest that a very simple neural network architecture could be used to generate an equity ranking system that might provide insight into the phenomenon of *abrupt performance inversions* characteristic of the VL system and perhaps also improve upon the VL system itself by replacing its coarse-grained ranking with a more fine-grained version.

An initial system was developed for trading purposes and employed successfully as part of a number of hedge funds and funds of funds in different configurations in 2002. However, because of the very high volatility at all time scales, in spite of continuing good returns, it was decided to perform a much longer and detailed set of studies of the method and the nature of the volatility of which this paper is part.

## **A10.4 Methods**

### **A10.4.1 Equities**

More than 1600 equities from the VL universe were selected with data eventually collected both by hand and from electronic sources from March 1, 1992 through December 1, 2001. All data was checked against two independent sources for consistency<sup>1</sup>, and a third if discrepant (The primary source of data was the Value Line Investment Survey itself (the print version). The secondary source was Bloomberg, inc. Tertiary data sources were chiefly WRDS and Telescan). Equities with irreconcilable data discrepancy rates > 0.5% were eliminated from use. This resulted in a significantly smaller pool of equities than VL itself routinely uses in its ranking system and a much cleaner data set. From this universe, a permanently fixed set of 1452 equities were identified for data extraction.

However, in any given quarter, fewer than 1452 stocks may actually be ranked. This is always because of the listing of new corporations and the delisting of existing ones. To have included only equities that were listed throughout the test period (plus 10 prior training quarters for the first out-of-sample prediction = 39 quarters  $\approx$  10 years would have resulted in a very reduced set of equities highly biased toward large capitalization corporations unrepresentative of the VL universe.

### **A10.4.2 Input data and outputs**

Inputs to the network consist of the ten preceding quarterly percent price and earnings changes (not accumulated) transformed as ranks. Outputs are the predicted next quarter's percentage price changes. All  $\sim$ 1452 stocks are then ranked in descending order by the ANN's predicted percentage price change for the next (out of sample) quarter. (The MGL predictor simply uses the prior quarter's actual price-change rank as the best estimate for the next quarter.)

From this output, for each successive out of sample quarter, twenty portfolios are constructed (and ten more from hedged combinations among these twenty). The twenty portfolios represent cumulative deciles from 10 to 100 from the top and bottom ends of the ranking. A T10 portfolio consists of the 10 equities predicted to perform best, the T20 the twenty equities predicted to perform best, and so on to T100. The deciles are cumulative in the sense that the T20 portfolio consist of the T10 portfolio plus the next 10 best and so on. The B10,...,B100 portfolios are constructed similarly but from the bottom of the ranking up. H10,...,H100 portfolios represented combinations of the respective T and B cumulative deciles with the T equities bought and the B equities sold short.

Depending on the week or month that the data is drawn from, raw price data will vary, of course, within a given quarter, whereas raw earnings data will either be unavailable, available, or will be available and then modified after the fact. Only original earnings reports were used, and only for those weeks and months in a cycle when they would actually have been available. Furthermore, based both on the well-known date of release

problem in historical data, and on a variety of other glitches that arise in real-time trading (that were uncovered during experience with this method in 2002) approximately 30% of earnings reports that look as though they would have been available on a given day actually are not. Therefore, the final column of earnings data is not used at all as input for those weekly or monthly date cycles when it couldn't be available at all, and in all earnings input columns, 30% of the earnings figures are removed at random before ranking to simulate other real-life problems.

#### **A10.4.3 Selection of trading period start**

Given a starting quarter, there would in principle be (on average) ~thirteen different weekly periods of data all starting in that quarter and sharing the same change in earnings value; or three monthly periods. (VL reports changes in its ranking system on a weekly basis.) The data structures for each of these cycles differ in their relation to earnings releases both with regard to the availability in relation to pricing data and from company to company. All of these considerations have been addressed, but because of the complexity of the task in back testing (by contrast to collecting data in real time going forward), the study reported here is limited to a single cycle of properly collected and error-checked data rather than an aggregation of between two and twelve weeks of data with an unknown amount of error and anachronism. The completed and fully error checked data set is simply the single best one able to be completed with the available resources. It contains no *known* errors. Back tests on other incomplete cycles show qualitatively similar results. The data period reported on here makes its first prediction for June 1, 1993 and its last for December 1, 2001 (roughly comparable to the report on hedge fund performance referenced below [138]).

#### **A10.4.4 Network Architecture**

The results reported on here are obtained using a simple back propagation network with a single hidden layer and recurrence. The results of multiple initializations are aggregated to obtain a final ranking. Exact net architecture and parameters are optimized independently on each new data set using a genetic algorithm but with extremely tight constraints. No variable deletion is allowed. Only a hidden single hidden layer is allowed. In general, minimal searching is permitted.

#### **A10.4.5 Training**

Training and testing sets are selected a-priori at random for optimization of training iterations. Under-fitting is greatly preferred to over fitting. Results are relatively insensitive to training lengths between six and twenty quarters. The results shown here are at about the median.

Two points should be emphasized here. First, as is an appropriate procedure in the use of ANNs, the network is always freshly trained on (ten quarters of) data that is out of the (one quarter) prediction sample. The ANN never has access to data from within the period it is predicting, hence the special care required with respect to using historical earnings data as explained in section 3.2.

Second, while the VL method as received by a subscriber appears static in the sense of implementing no evident adaptive or learning mechanisms, we know informally from a private meeting with the founder of the VL system that the regression-like formulae em-

ployed by VL are updated over time. (Thus any decline in its performance over time cannot be attributed to its algorithms become outdated solely because they are static).

## 5 Results

### A10.5.1 Hedged Returns


#### A10.5.1.1 Overall results

**Figure 31** provides a concise graphic snapshot of results, demonstrating the superior performance of the ANN predictor (not including transaction costs) relative to a one-step Martingale (MGL) predictor and (for all 100 equities) to the VL system itself, as well as to the S&P 500 index. Once a month, the ANN and MGL methods are used to predict and rank-order the top 10, top 20, ..., top 100 equities (i.e., “cumulative decile”: T10, T20, ..., T100, from among the universe of 1452 stocks based on the inputs as described in section A10.4.2), as well as the bottom deciles: B10, B20, ..., B100. A portfolio is composed of matched T deciles held long and B deciles sold short (resulting in fully hedged “market-neutral” portfolios H10, H20, ..., H100). Every portfolio is readjusted once per month. To adjust for possible monthly or seasonal effects, results are averaged over all three possible monthly starting points in a quarter. These results are compared to the actual VL selection of T100 and B100 stocks (groups “1” and “5” respectively) adjusted every quarter similarly (rotated and averaged), and to the S&P 500 Index over the entire out-of-sample range of 29 quarters. As shown here, trading costs are not included (to be discussed later). Note, however, that for small portfolios (i.e., H10, H20, H30), even were the turnover to be 100%, such costs are relatively moderate as they occur only at quarterly intervals.

**Figure 31** illustrates that at the end of the 29-quarter period, the ANN predictor succeeds at separating high-performing from low-performing stocks sufficiently well to generate substantial returns for all hedged decile ANN portfolios. A 6% annual risk-free rate of return has been assumed (high, therefore conservative). Furthermore, the internal progressively layered relations among the top 10, top 20, ..., top 100 are very well-preserved by the network predictor: Returns generally fall off by cumulative decile (CD) implying that the separation is highly significant (These relations are not preserved by the MGL predictor which generates roughly the same loss for all cumulative deciles.

On the right edge, we see that the fully hedged  $T100+B100 = H100$  equity portfolio using the network predictor yields annualized returns of 16.7% (10.7% in excess of the risk-free return). The MGL predictor for the same hedged portfolio yields excess annualized returns of 4.0% while the VL ranking system yields  $-0.8\%$

During this period both the broad market index (the S&P 500) and the (unweighted) VL universe performed roughly comparably, i.e., flat to slightly negative, thus we have eliminated any overall bias during the test period, but this result is composed of a period of rapid growth followed by a short period of high volatility followed by a period of rapid decline (before and after the 2000 market bubble), a challenging stretch of time for any model. The ANN not only does a superior job of ranking stocks than the VL method itself (while employing what is likely very similar inputs), in addition it parses the ranking more finely. (Both the ANN and MGL predictors provide explicit rankings for all 1452 stocks.)




**Figure 31:** Comparative annualized excess returns from the ANN and MGL predictors for hedged (top and bottom 10, 20, ..., 100) cumulative decile (CD) baskets of securities and for the VL system hedged Rank 1 and Rank 5 stocks.

To understand the relation of the ANN and MGL predictors to changes of regime, it will later prove instructive to decompose the hedge into its constituent long and short components.

#### *A10.5.1.2 Long Returns*

**Figure 32** demonstrates the long-only ANN, MGL and VL system portfolios. Both the MGL and ANN predictors successfully generate the progressive relationship among cumulative deciles, but the ANN predictor generates superior returns in 9 of 10 instances. Comparing only the ANN Top 100, the MGL Top 100 and the VL Top 100, we find that the ANN returns 19.6%, the MGL 11.6% and the VL system 9.9% compared to the VL universe of 1452 stocks which returned  $-0.16\%$  over this period. (The S&P 500, by comparison, returned  $-0.35\%$ , i.e., both VL and SP are comparably flat). Thus, in a head-to-head comparison with the VL system (using all T100 stocks), the network predictor performs best by a large margin, the MGL predictor next best, the VL system comparably to the MGL and all three significantly better than the VL universe as a whole. Comparing the hedged to long only results, we see that the VL system has succeeded in identifying rising stocks but not in identifying poor-performing or declining ones. This fact is made evident by examining the short side of the results.




**Figure 32:** Comparative annualized excess returns from the ANN and MGL predictors for long (top 10, 20, ..., 100) cumulative decile (CD) baskets of securities and for the VL system hedged Rank 1 and Rank 5 stocks.

### A10.5.1.3 Short Returns

In **Figure 33**, we illustrate the short-only ANN, MGL and VL system portfolios (results show the actual returns, not their negatives as is required for a short position). In other words, the lower the return the more desirable it is. We see that no system succeeds in generating absolutely negative results (it is typically far more difficult to predict stocks falling in price than rising ones). But the value of a hedged portfolio is not in amplifying gains by succeeding in shorting falling stocks. The goal is rather to create a portfolio that is “market-neutral” so as to neutralize price changes that may attributed to changes in the market as a whole. Captured gains therefore presumably arise from the intelligent selection of a portfolio of strategically chosen equities from among the available choices. (We will quantify the degree of success achieved by the ANN predictor in the next section.)

We see that the ANN successfully preserves the appropriate progressive relations among CDs: The bottom 10 are the worst performers (best for shorting), the bottom 100 the best (worst for shorting). The MGL predictor is not nearly so good as it was in identifying stocks for the long component—though it does a better job than the VL Rank 5 selection, its short baskets are all much better performing (therefore worse for shorting) than the ANN for all CD portfolios and the progressive relations among CDs are not preserved. Note, however, that the scale for the short selections is much compressed vis-à-vis the scale for long selection. The superior performance of the ANN versus the MGL predictor for hedged portfolios is therefore attributable both to its superior selection of top-performing sticks and bottom-performing stocks. The VL system fails altogether—its selection of bottom-performing stocks does better than the VL universe and even better than its selection of top-performing stocks.




**Figure 33:** Comparative annualized excess returns from the ANN and MGL predictors for short (bottom 10, 20, ..., 100) cumulative decile (CD) baskets of securities and for the VL system hedged Rank and Rank 5 stocks. Returns shown are as though long, not their negatives (shorted) as computed in the hedged portfolios. For the short component of the hedged portfolio, therefore, the lower the return the more desirable.

## A10.5.2 Risk-adjusted returns

### A10.5.2.1 Sharpe ratios

The simple hedged returns by cumulative decile provide an excellent test of the capacity of the ANN to extract information about future stock performance from prior price and earnings changes. Note, however, that this model does not attempt to predict or to minimize risk, i.e., volatility. We do not necessarily expect that such a model will produce superior *risk-adjusted* returns. Indeed, a well-known correlation between performance and risk is the *bête noir* of most aggressive approaches to achieving superior returns. A perhaps overly simple but widely-used measure of risk-adjusted performance is the “Sharpe ratio”, i.e., the ratio of the excess annualized return to the annualized standard deviation of the price. **Figure 34** provides a comparison of Sharpe ratios for the same categories as the excess returns in **Figure 31**. The general fall-off with cumulative decile is preserved once again for the ANN and more weakly for the MGL predictor. Except for the first cumulative decile, the resulting Sharpe ratios for the ANN predictor are lower than for the S&P 500 reflecting the relatively high volatility of the equities selected within the top 100. (Post-hoc examination reveals that these are very disproportionately in the technology sector which over this time period experienced unusually dramatic volatility in fact.) Thus, the superior returns generated by the ANN come at the cost of high volatility. (Note, however, that by this measure the ANN significantly outperforms the VL methodology for selecting high performance stocks as well as the MGL predictor and the VL universe as a whole.



**Figure 34:** Comparative Sharpe ratios (risk-adjusted returns) for the ANN and MGL predictors for hedged baskets of securities; for the VL system proper; and for the VL and S&P 500 universes of equities. (VL H SR [green] and All VL SR [blue] have the same value.)

Hedge funds in general claim to aim for and achieve Sharpe ratios of between 1 and 2. Again assuming a 6% risk-free rate of return, between January of 1990 and June of 2003, equity long/short funds reported annualized excess returns of ~12.07% [139] and an important study of eleven major market-neutral hedge funds from May, 1990 through April 2000 (hence during a period of generally rising markets only) reported Sharpe ratios of 1.1 [140]. By this measure, the ANN predictor falls short—unless one considers the very much more difficult and volatile period encompassed by the study. (It is very likely, furthermore that Sharpe ratios of 1-2 are less common than measured by various hedge fund tracking reports (for example, Hedge Fund Research) because data on failed funds is often unavailable, especially those that fail relatively quickly [139].

Another and arguably more accurate indicator of performance is provided by Jensen's alpha, and by the associated beta, a measure not of absolute volatility in terms of simple arithmetic or logarithmic price change, but of expected volatility given the volatility of the universe of stocks from which a portfolio is selected. From this perspective, the ANN succeeds remarkably well.

#### *A10.5.2.2 Jensen's alpha*

The most widely used tool for assessing risk-adjusted investment performance is Jensen's alpha ( $\alpha$ ).  $\alpha$  is designed to quantify how an investment performs not absolutely but relative to the volatility of the actual market universe from which it is drawn. It has been widely demonstrated that high volatility investments with a large possibility of large losses are likely to demonstrate larger gains (upon success) than low volatility ones: The

risk of losing a great deal is compensated for as a larger “risk premium”. Thus, no investment “intelligence” is required to attain large gains by simply investing in a very high-risk vehicle. For example, a majority of start-up companies fail altogether. But those that succeed return a very high premium to their investors who assume the high risk. Well-established, so-called “blue chip” corporations offer relatively modest returns to purchasers of their stocks. But in exchange, shareholders may anticipate a relatively low risk of large losses.

The value added by a portfolio investment strategy (or manager) is therefore associated with gains beyond that attributable to the simple volatility of the appropriate universe of choices from which the strategy is drawn.  $\alpha$  reflects the difference between the investment strategy’s actual performance and the performance expected based simply on inherent risk. An investment that produces the expected return for the level of risk has an  $\alpha$  of zero.  $\alpha > 0$  implies that the strategy produced a return greater than expected for the risk taken.  $\alpha < 0$  indicates that the strategy has produced a return smaller than expected relative to the risk.

It has been widely observed that ~75% of stock investment managers fail to improve on the performance of someone who had simply invested in a market-weighted basket of every stock. This phenomenon has been argued to be due to the “efficiency” of markets. This belief yielded market capitalization weighted index funds that seek to replicate broad market indices (i.e., baskets of securities representative of the entire pool of securities from which the selected ones are drawn), i.e., to reproduce the returns of those indices, hence aim for  $\alpha = 0$ .

$\alpha$  thus depends upon a measure of risk that is relative to a given market denoted Beta ( $\beta$ ). The  $\beta$  of an investment strategy is defined as:

$$\beta_{strat} = \frac{Cov[R_{strat}, R_{mkt}]}{Var[R_{mkt}]} \quad (87)$$

where  $R_{strat}$  is the return of the strategy and  $R_{mkt}$  is the return of the market from which the strategy is drawn. In other words,  $\beta_{strat}$  is the slope of the linear fit of a scatter-plot with  $R_{strat}$  the abscissa,  $R_{mkt}$  the ordinate.

Stock index funds established a tacit standard of performance for investment strategies (and their managers): The successful strategy is one that yields performance in excess of the passive strategy of investing in everything equally since the latter strategy is more likely to be the better one, statistically.  $\alpha$  is thus any additional return above the expected return of the  $\beta$ -adjusted return of the appropriate market.

The formal expression for investment  $\alpha$  is derived from the Capital Asset Pricing Model (CAPM), wherein the estimated return  $R_i$  on a security  $s_i$  is given by three terms:

$$R_i = \alpha_i + \beta_i R_{mkt} + \varepsilon_i \quad (88)$$


$\alpha_i$  (in CAPM) is considered a constant drift unique to the  $i^{\text{th}}$  asset in the market with asset weights  $mkt \equiv \{s_1, s_2, \dots, s_i, \dots, s_N\}$ ;  $\beta_i$  is the volatility of asset  $i$  as defined in (87);  $\varepsilon_i$  is

a random variation unique to asset  $i$  with mean zero (hence as the number of different assets increases in a portfolio of many securities, the combined  $\varepsilon_i$  for the portfolio vanishes). Thus:

$$\alpha_{strat} \equiv R_{strat} - \left[ R_{rf} + \beta_{strat} (R_{mkt} - R_{rf}) \right] \quad (89)$$

where returns have been adjusted by the risk-free rate of return,  $R_{rf}$ . By this standard,  $\alpha_{strat} > 0$  represents a successful strategy.

The  $\alpha$  results for the ANN predictor are especially illuminating, given that the study period was one of unusual volatility. These are presented in **Figure 35**. The results show  $0.4 > \alpha_{ANN} > 0.2$ , roughly half the value reported by eleven major market-neutral funds between the bull-market of May, 1990 and April, 2000 ( $\alpha \approx 0.6$ ), but still significantly positive (and in line with the Sharpe values presented above). More importantly the preservation of progressive relations among CDs is preserved more rigorously and for this measure of risk-adjusted performance, a power-law fall-off has a high correlation coefficient. (The market for which we calculate  $\beta_{ANN}$  to derive  $\alpha_{ANN}$  is the VL universe.) As we may anticipate from **Figure 34**,  $\alpha$  for the MGL predictor is poor, hovering at or below zero throughout, a reflection of the excess volatility associated with its returns.



**Figure 35:**  $\alpha$  for the ANN with  $\beta$  for the VL 1452. (MGL results not shown as they are evidently worse)

We now turn to the question of how regime changes affect predictor performance which is the major purpose of this chapter.

### **A10.6 Antipersistence in predictor behavior**

This study was initially motivated by two facts: First, that a subtle feature of the VL system is that even if it has a history of in general working (i.e., making successful predictions contrary to the EMH), when it fails its failures are especially striking. Second, that the ANN predictor described above, which created a more fine-grained version of the VL rankings, and which was traded successfully in the real-world (with net positive post-

transaction-cost gains) likewise suffered from high volatility—i.e., its periods of success were remarkably large but were interspersed with periods of large failure.

In section 4., studying a much longer period of time, we have shown that in spite of high volatility, the ANN predictor nonetheless is able to generate significantly positive  $\alpha$ . Thus, its gains, and implied predictive capacity, can not be attributed solely to the volatility ( ) of the underlying market in which it trades.


We wish therefore to understand what is the nature of the swings in the returns generated by the predictor. To this we end, we first examine in finer detail the structure of the rankings generated by the ANN across all ~1500 equities. We argue that the peculiar distribution of actual returns by predicted rank may be understood *within each time period* (quarter) as a weighted mixture of highly successful and highly unsuccessful predictions. Whether a time-period's out-of-sample prediction is on balance successful or not depends upon the relative weights of the successful and unsuccessful components of the ranking. The respective components may be represented mathematically as negative and positive ArcTanh functions.

We then examine the decomposition of the rankings across time and show that the balance between successful (net negative ArcTanh distribution of returns by ranking) and unsuccessful (net positive ArcTanh distribution of returns), varies to a degree that is greater than what should be expected by chance, in particular showing “anti-persistence”. Persistence is defined as a measure of the tendency of a binary series to exhibit repeating patterns; anti-persistence the tendency to specifically avoid repeating patterns. A random sequence exhibits neither persistence nor anti-persistence.

We then concentrate on just the extreme left (T10-T100 portfolios) and right (B10-B100 portfolios) ends of the distributions, combining the data across time into two sets: one set representing all the net-successful quarters, a second set representing all the net unsuccessful quarters. From the hypothetical returns generated by these two contrasting sets, we find that the nature of the failure of the ANN predictor, when it fails, is different from and less dramatic than the more complete inversion of results of the MGL predictor, especially in choosing equities for the B10-B100 portfolios.


#### **A10.6.1 Distribution of ANN predictions as a composite of success and failure**

For each of the 29 out-of-sample quarter, both the MGL and ANN predictors forecast the quarterly price change for every stock. The stocks are then sorted in descending order by this forecast to obtain the predicted rank-ordering. Consider quarter 26 as a first example. **Figure 36** shows the distribution of actual returns from for all equities ranked by the ANN predictor. Those ranked highest are to the left, those lowest to the right. The butterfly-shaped plot is typical for both the MGL predictor and the ANN predictor. The precise shape and density of the distribution determines whether the predictor is successful for long, short or hedged portfolios and for which CDs.




**Figure 36:** All 1452 returns (vertical axis) by rank (horizontal axis) in quarter 26 ranked by the ANN predictor

Now suppose that the predictor were 100% successful: i.e., the rank-ordering predicted by the ANN corresponded exactly to the descending ordering of actual price-changes. In that case **Figure 36** would look rather like **Figure 37**.



**Figure 37:** All equities for quarter 26 in descending (exactly correct) order of actual price change.

If the predictor were 100% unsuccessful, **Figure 36** would look instead like



**Figure 38:** All equities for quarter 26 in ascending (exactly wrong) order of actual price change

The actual butterfly distribution then may be considered a mixture of the fully correct and fully incorrect distributions. To find this empirically, we first fit both the correct and incorrect distributions to respective ArcTanh functions with parameters that thus reflect the actual distribution for the quarter. We find that:

$$\Delta p_{correct} = -0.254 + 0.346 \times \text{ArcTanh}(1 - 0.000689 r_{correct}) \quad (90)$$

where  $\Delta p_{correct}$  is the actual price change as predicted by the correct ranking and  $r_{correct}$  is the correctly predicted rank.

Likewise:


$$\Delta p_{incorrect} = 0.227 - 0.350 \times \text{ArcTanh}(1 - 0.00689 r_{incorrect}) \quad (91)$$

with  $r_{incorrect}$  the rank predicted maximally incorrectly. (In both (90) and (91) the subscript on  $\Delta p$  refers to the source of the price change, either the correct or incorrect ranking. In both cases the  $\Delta p$  in question is one of the actual ones.)

We may now fit the actual distribution of price changes by the actual ANN prediction/ranking to a weighted sum of these two functions (for which the ArcTanh terms are of course the same):

$$\Delta p = 0.510 \Delta p_{correct} + 0.490 \Delta p_{incorrect} \quad (92)$$

which yields the curve shown in Figure 39. (Only a part of the range is shown—the mid-range of **Figure 36**—because the wide scatter would otherwise obscure the shape of the fit. The same fit may be obtained directly from a single ArcTanh fit, but this would not yield the weights of the decomposed correct and incorrect components.)



**Figure 39:** ArcTanh fit to the actual price change by ANN predicted rank for the out-of-sample data in quarter 26.


The fact that the fit curve is generally declining from left to right implies that over the entire list of equities the predictor is on average successful. The flatter region at the right end implies a smaller degree of success in identifying stocks that will decline relative to the universe; the sharper upturn at the left end implies a larger degree of success in identifying stocks that will rise relative to the universe. The fact that the curve is almost entirely below the zero line illustrates that the universe of equities declined overall this quarter

Each quarter's rank ordering constructed by the ANN may be similarly treated as a mixture or "superposition" of a globally successful and a globally unsuccessful distribution. The net balance of each yields both the distinctive final result and generates the typical butterfly distribution.

Note that Figure 39 shows only a fit for quarter 26 (as an illustration) and only for the ANN. The question may be raised as to whether such a fit is meaningful. The most straightforward evidence that it is (averaged across all quarters)—and that it is also both meaningful in the MGL equivalent and in comparison shows evidence of the MGL predictor's poorer performance follows from the three facts that (1) the ANN generates significantly positive returns for the hedged 10, 20, ... 100 portfolios (left and right ends); that (2) the same is true for the MGL predictor, but less so; and (3) that the returns among the 10 CDs for both the ANN and MGL generally fall off by CD. Furthermore, we discuss in section 5.4 similar graphed results examining *only* the top and bottom extremes of the distribution but aggregated for all quarters, segregated by winning and losing quarters. A negative ArcTanh-like distribution is visibly evident for the aggregated winning quarters likewise a positive Arc-Tanh-like distribution for aggregated losing ones. Comparing these similar results between the ANN and MGL we find in the different structure of the distribution an explanation for how the ANN outperforms the MGL.

### A10.6.2 Anti-persistence in predictor performance

The weights for the successful and unsuccessful components in each quarter’s prediction always add to 1. We may therefore track the weight of the successful component to glean a snapshot of ANN performance over time as in **Figure 40** (and in ranking *all* ~1500 equities, not just the extremes.)



**Figure 40:** Proportion of successful component ( $W_{success}$ ) of ANN predictor by quarter.

Impressionistically, we see that the predictor appears to show relatively wide and frequent swings between being globally successful and not. We may quantify this impression by considering the sequence of “successful” ( $W_{success} > 0$ ) and “unsuccessful” ( $W_{success} < 0$ ) quarters as a binary series and examining this series for persistence.

As discussed in [83, 93, 141], persistence is a formally defined measure on  $[0,1]$  of the extent to which patterns in a binary series tend to repeat, anti-persistence of their tendency to avoid repetition. The persistence  $\mathfrak{P}$  of a maximally repetitive sequence equals 0, of a maximally anti-persistent sequence 1. A random sequence has persistence  $\mathfrak{P} = 0.5$ .

In brief, persistence is determined at a particular scale  $m_s$  by examining all possible binary  $m_s$ -bit subsequences and counting the proportion of times that when a particular subsequence is followed in the series by 0, it is followed by 0 again upon its next occurrence; likewise for a following 1. Details of the calculation and examples from other domains may be found in [83, 93, 141]. At a scale of 1 (which is the only scale at which so short a sequence can have a meaningful persistence measure),  $\mathfrak{P}$  measures the tendency of a series not to alternate ( $1 - \mathfrak{P}$ , its anti-persistence, the tendency to alternate, with the series 1,0,1,0,1,0,... or 0,1,0,1,0,1,... having  $\mathfrak{P} = 0.5$ ).

We find that the 29 quarters of predictor data in **Figure 40** have  $\mathfrak{P}_{MGL} = 0.143$  at  $m_s = 1$  which is highly anti-persistent. It might appear that with only 29 binary values, the measured  $\mathfrak{P}$  could not be statistically significant. But in fact, measuring  $\mathfrak{P}$  on 1,000,000 random binary sequences each of length 29 yields  $p \leq 0.00403$ . A similar set of analyses

performed on the results from the MGL predictor yields  $\mathfrak{P}_{MGL} = 0.357$  with  $p \leq 0.00127$ .

The implication of a significant degree of antipersistence in a predictor's results is that whether returns are in general positive or not, they are associated with a high degree of volatility. The MGL predictor's higher volatility as measured by its Sharpe ratio is consistent with its somewhat greater anti-persistence. We see in the comparison between anti-persistence in the ANN and in the MGL the fact that while high performance tends to be associated with high volatility, it is possible for subtler methods to yield a more satisfactory relationship between volatility and return as reflected in the measure of alpha.

A natural question is whether there is correlation between successful versus unsuccessful predictions by quarter and the direction of the market. Perhaps the ANN is successful when the market rises and unsuccessful when it falls—a common complication of naïve predictors. If the direction of the overall mean VL universe by quarter—equivalent to a buy and hold strategy—is converted into a binary series, this series is also anti-persistent with  $\mathfrak{P}_{All} = 0.283$ . For a series of this length, this degree of antipersistence is not highly statistically significant  $p \leq 0.035$ , suggesting that the variation in quarterly mean returns may well be effectively random. Furthermore there is no correlation between success by quarter for the ANN and mean gain by quarter:  $r^2 < 0.00001$ .

### **A10.6.3 Effects of overall predictor performance on the top and bottom ends of the predictor rankings**

The severity of the models failures when they fail (both the ANN predictor and the MGL) reflects the fact that the state of the model is at times of failure at best one time-step behind the phase of the market. Nonetheless, if its inputs have been properly chosen so as to reflect actual feedback characteristics of the market in question such that by tacitly learning patterns of market change (which a MGL predictor cannot do), the ANN model should be able to change its state rapidly enough to compensate for those time periods when its state is out of phase with the market and generate a net cumulative positive return that is significantly greater than any control, as has occurred.

We know that at least some real-world feedback has been demonstrated between the selection of top-ranked VL stocks and changes in its price and earnings reports. The ANN inputs have been chosen with this in mind and look backward over ten quarters' worth of prior earnings and price rankings, developing a tacit relation among these in determining its prediction for the subsequent ranking. Unlike the MGL predictor, it may detect a pattern not only of ranking but of change of rank structure over time. (We do not here report on a similar predictor that does not use earnings data. Results are degraded in the direction of the MGL predictor but remain superior to it.)


In order to quantify results efficiently, and test this hypothesis using so limited a number of quarters, the quarterly results can be segregated by positive or negative net return for the market as a whole (the VL universe) and then aggregated. If there is a strong tendency for quarters with net positive returns to show a negative ArcTanh-like distribution (and an positive ArcTanh for quarters with negative returns), aggregation should amplify this effect. Since quarters with results close to zero are included, considerable noise is introduced, strengthening the significance of the aggregation while ensuring that no se-

lection bias is involved in the segregation. We do not attempt to fit our results, since as will be demonstrated, we are combing data from the extremes of the distribution, creating a discontinuity.

Aggregated results by top, bottom and hedged 10 through 100 deciles can be compared to the MGL predictor for the same groupings likewise segregated into the same two major classes (rising quarters versus declining quarters).

In preparation for this analysis, we first provide some relevant global measures of performance.

**Figure 41** illustrates the cumulative returns over time for the hedged ANN portfolios (log scale). Note the tendency for the returns to be relatively extreme in both directions; for quarters with sharply positive (negative) returns to be followed by quarters with sharply negative (positive) returns; for the progressive relationship by decile to be relatively well-preserved through time; and for the net cumulative return in all fully balanced long-short portfolios to be significantly positive. Note too how the progressive relations among cumulative deciles are quite well preserved over time.



**Figure 41:** Semi-log Chart of Multiple of Initial Investment in Hedged Portfolios over 29 Quarters

**Table 13** shows the cumulative decile performance for the ANN and MGL predictors, followed by their differences. We will want to analyze what contributes to this difference in performance. We will see that the two predictors have many similarities that contribute to both their successes and their failures, but in varying proportion. In quantifying these proportions we can gain some insight into the phenomenon of change of regime and/or phase.

Portfolio	Excess Ann'l Return	Portfolio	Excess Ann'l Return	Difference
ANN H 10	50.2%	MGL H 10	-10.2%	60.4%
ANN H 20	18.1%	MGL H 20	-10.0%	28.1%
ANN H 30	19.2%	MGL H 30	-9.6%	28.8%
ANN H 40	13.2%	MGL H 40	-10.4%	23.5%
ANN H 50	11.3%	MGL H 50	-10.6%	21.9%
ANN H 60	10.9%	MGL H 60	-10.6%	21.5%
ANN H 70	8.7%	MGL H 70	-10.8%	19.5%
ANN H 80	11.0%	MGL H 80	-10.7%	21.8%

ANN H 90	11.5%	MGL H 90	-10.8%	22.3%
ANN H 100	10.7%	MGL H 100	-11.1%	21.8%

As noted before, the MGL method does reasonably well in selecting equities that simply *continue to rise*—especially as part of a universe of equities experiencing a general rise. But the method does not appear capable of differentiating those that will significantly rise (relatively) from those that will significantly fall, a weakness that it shares with the VL predictor than whose final results it does scarcely any better. We will see that the ANN shares this same weakness, but in lesser degree. This particular kind of failure appears in two guises:

1. The failure appears cross-sectionally within any given prediction set in the form of a series of equities predicted to make large price changes in one direction that actually make large price changes in the opposite direction (giving rise to the anti-persistent behavior discussed above). Typically, such mistaken predictions are admixed with a large number of correct predictions. This gives to the overall scatter plot of rank-orderings versus actual percentage price changes a butterfly shape from which quantifiable information will be extracted for the ANN predictor, the MGL control and their differences. The weighted proportion of successes and failures in the top 10, 20, ..., 100 portfolios versus the matching bottom portfolios determines the magnitude of the success or failure of the respective hedged outcome for any given quarter.
2. The failure appears intermittently and abruptly through time as the balance between overall correct and incorrect predictions shifts. As we will see, this shifting balance produces large changes in outcome from quarter to quarter generating the impression of a change in regime that is especially pronounced in the MGL model. The final results are a consequence of the accumulated successes and failures over time.

#### **A10.6.4 T/B Portfolios aggregated by global success or failure of the predictor(s)**

##### *A10.6.4.1 T/B 10 = H10 Portfolios*

One gains an intuitive impression that the universe of equities tracked by the model(s) undergoes sudden changes in performance partially captured in the “whipsaw” behavior of the predictors generated by the models. The challenge is to devise and extract a simple measure that objectively characterizes this phenomenon, if present.

We have taken the following approach: All ~1500 equities have been assigned out-of-sample predicted ranks for all 29 quarters both by the ANN and the MGL predictors. All equities likewise have their known in-sample actual percentage price-changes for every quarter. We thus start with two ~1500 X 29 tables showing the rank (by row number, with the best predicted rank being 1, the worst ~1500), one table for the ANN and the other for the MGL. From these tables we keep only the top 100 rows and the bottom 100 rows. We create an identical set of tables this time keeping only the top 90 rows and the bottom 90 rows. Again for the top and bottom 80, 70, ..., 10 rows. Within each table, we identify those quarters where counting all ~1500 equities, the mean market change is positive and where it is negative.

Table 14 shows a hypothetical example for the T100 and B100 price changes sorted by predicted rank: Gray columns represent quarters where the mean price change for the portfolio—long on the top 100, short on the bottom 100—is negative. White columns represent positive quarters.

**Table 14:** A data table with (hypothetical) actual price changes arranged by predicted rank and quarter for top and bottom 100 equities.

Rank/Quarter →	Q1	Q2	Q3	...	Q29
<b>1</b>	+0.001	+0.002	−0.004	...	−0.014
<b>2</b>	−0.017	+0.018	−0.050	...	+0.005
...	...	...	...	...	...
<b>100</b>	−0.003	−0.009	−0.011	...	−0.014
<b>1353</b>	+0.020	+0.017	+0.021	...	−0.014
...	...	...	...	...	−0.003
<b>1451</b>	+0.008	+0.015	+0.034	...	+0.001
<b>1452</b>	+0.045	−0.003	−0.016	...	+0.022
<b>Mean Hedged Ret.</b>	<b>−0.043</b>	<b>+0.032</b>	<b>−0.004</b>	...	<b>+0.007</b>

For each data table (2 tables for each of the size 100, 90, ..., 10 data sets), we then segregate the successful and unsuccessful quarters (+ or − net or mean gain for the hedged portfolio) as in **Table 15** and **Table 16**.

**Table 15:** Successful quarters only

Rank/Quarter →	Q2	...	Q29
<b>1</b>	+0.002	...	−0.014
<b>2</b>	+0.018	...	+0.005
...	...	...	...
<b>100</b>	−0.009	...	−0.014
<b>1353</b>	+0.017	...	−0.014
...	...	...	−0.003
<b>1451</b>	+0.015	...	+0.001
<b>1452</b>	−0.003	...	+0.022
<b>Mean Hedged Ret.</b>	<b>+0.032</b>	...	<b>+0.007</b>


**Table 16:** Unsuccessful quarters only

Rank/Quarter →	Q1	Q3	...
<b>1</b>	+0.001	−0.004	...
<b>2</b>	−0.017	−0.050	...
...	...	...	...
<b>100</b>	−0.003	−0.011	...
<b>1353</b>	+0.020	+0.021	...
...	...	...	...
<b>1451</b>	+0.008	+0.034	...
<b>1452</b>	+0.045	−0.016	...
<b>Mean Hedged Ret.</b>	<b>−0.043</b>	<b>−0.004</b>	...

Within each portfolio H10 – H100, but now segregated into all successful quarters and all unsuccessful quarters, we examine the top component of the portfolio and the bottom component (and do so for both the ANN and MGL predictors). We quantify the performance *across* all successful quarters of the rank 1 stock, the rank 2 stock,....through the rank 10, 20,....,100 stock (depending on the size of the portfolio we are studying). Then we do the same with the bottom groups. We likewise quantify the performance in the same way across all unsuccessful quarters. This will allow us to examine the structure of the each predictors' assignment of rank not quarter-by-quarter, but treating successful and unsuccessful quarters (as we did in sections 5.2 and 5.) but in the aggregate instead, and each aggregation separately.

The simplest way to thus aggregate the data across rows is to use their mean and standard deviation (which we will examine shortly). But this would fail to account for the difference in the number of winning and losing quarters (and the difference in these numbers between the ANN and MGL predictors) which after all is a key consideration. The measure employed here is therefore an artificial metric adapted from the “natural” features of the financial domain: Pseudo-compounding of the numbers across a row (plus one each) as though all winning quarters occurred in sequence following an initial investment of 1; all losing quarters similarly. This method is meant to place all quarters on an independent footing while properly weighting the geometric effect of an imbalance in the number of winning and losing quarters. The resulting numbers will therefore scale somewhat like the actual results but are not “real”: Both the winning and the losing quarters, for both MGL and ANN tables have simply been compounded in sequence from the actual percent changes. Unrealistically sized final gains and losses for each row are then used in place of a mean, but the relative relations among points are properly preserved. while being pushed apart around 1.

**Figure 42** shows the results of the MGL predictor for the Top and Bottom 10 portfolios under this transformation. Each mark represents a pseudo-compounded return for a quarter with blue marks for the aggregated successful quarters (“+”); red marks for the unsuccessful ones (“-”). Keeping in mind the discussion in sections 5.1 and 5.2, we see something similar: The blue marks all together (from both top and bottom ends of the ranking, joined together at the midline, as it were) create a negative ArcTanh-like distribution, the red ones a positive ArcTanh. The blue distribution thus mimics a successful quarter when the predictor is “in phase” so to speak with the distribution of price changes; the red distribution mimics an unsuccessful quarter when the predictor is “out of phase” with the distribution of price changes. But in this case we have aggregated the data from all quarters, collating the “in phase” and “out of phase” results separately.



**Figure 42:** Pseudo-compounded returns across all quarters, for Top and Bottom 10 MGL predicted stocks, segregated by the success (“+”, blue, “in phase”) or lack of success (“-”, red, “out of phase”) of the MGL predictor. The x-axis shows the rank (top 1 through 10) of the equity to the left side of the vertical line, and the rank (bottom 1 through 10) on the right side; the y axis the pseudo-compounded returns.

There are four components to **Figure 42**, each with ten (colored) elements:

- ▲ T10 MGL portfolio pseudo-compounded return during a (+) ”phase” of the predictor
- ▲ T10 MGL portfolio pseudo-compounded return during a (-) ”phase” of the predictor
- B10 MGL portfolio pseudo-compounded return during a (+) ”phase” of the predictor
- B10 MGL portfolio pseudo-compounded return during a (-) ”phase” of the predictor

The left hand axis indicates the multiple of an initial return of 1.0 at the end of all winning (losing) quarters, were returns to be compounded back to back. (“Pseudo-compounded”. Later, we will combine all winning and losing quarters to show that these figures are not unrealistic even though they appear to be when segregated in this way).

Note that blue marks represent data aggregated (compounded as described above) from winning quarters; red marks represent data aggregated from losing quarters. The left half of the horizontal axis represents equities ranked T1 through T10 across all quarters (from left to right, i.e., ranked 1 through 10), the right half equities B10 through B1 (from left to right, i.e., ranked 1443 through 1452). The leftmost blue and red marks are in the same horizontal position, likewise for all remaining nineteen successive rightward blue and red marks. The blue leftmost mark represents the pseudo-compounded return for just the number 1 ranked equity during all quarters with the predictor in phase with the market, the red leftmost mark for just the number 1 ranked equity during all quarters with the predictor out of phase with the market; the blue second-to-leftmost mark represents the

pseudo-compounded return for just the number 2 ranked equity during all quarters with the predictor in phase with the market, and so on.

Concentrating first on just the blue marks (successful, “in phase” quarters, “ $\phi^+$ ”), note that the marks to left of the midpoint are triangular and represent from left to right the aggregated pseudo-compounded returns respectively for the top 1 through top 10 ranked equities; marks to the right of the midpoint are circular and represent from left to right equities ranked 1443 through 1452 (the bottom ten in order). Each triangle represents the pseudo-compounded final return over 18 continuous quarters. These are quarters during which the MGL predictor generated positive net returns for a T10/B10 hedged portfolio. Although there is very wide scatter on individual real returns, in this aggregated transformed data, in all “in phase” ( $\phi^+$ ) quarters, all top 10 data points lie above all bottom 10 data points.


Turning now to the red marks, we find that the above relationship has been largely inverted:

For 11 “out of phase” quarters (“ $\phi^-$ ”), the aggregated data points representing equities predicted to perform as the top 10, are now found entirely below both the blue marks representing the top 10 for winning quarters and below the red circular marks representing the predicted bottom 10 performing equities. Furthermore, with one exception, each red circle lies above its corresponding blue circle. That is, during  $\phi^-$  quarters, rather than doing poorly, as predicted, the bottom 10 of the MGL model consistently outperform the bottom 10 of the MGL model during winning quarters, often by an order of magnitude or more; as well as greatly outperform the predicted top ten of the MGL predictor itself for those same quarters, and some proportion of the MGL top 10 predictor during winning quarters.

Thus, the MGL predictor—which simply uses the prior actual rank ordering of the market as a predictor—comes very close to demonstrating two distinct states that are evident at a glance—at least at its extremes. (The same phenomenon is evident, if somewhat less clearly, using charts of similarly aggregated and transformed data for the top and bottom 20, 30, ..., 100 as well as using non-transformed and non-aggregated data.) With this in view, the  $\phi^+$  quarters may be thought of as periods of “positive” predictor state in the sense that the state of the predictor is in phase with the performance of the universe of equities; Declining quarters we deem periods of “negative” predictor state  $\phi^-$ .

However, we may observe at least one significant departure from a strict inversion: During periods of negative predictor state  $\phi^-$ , the positive return of the B10 portfolio is not nearly so great as the positive return of the T10 portfolio during periods of positive state  $\phi^+$ . In consequence, the MGL predictor generates a net positive return, especially for the T10 portfolio alone (and similarly for T20-T100). We keep this asymmetry in mind as we turn to the ANN portfolio with which we compare it, at the end adding numerical quantification.

**Figure 43** shows the comparable chart for the ANN predictor and the resulting Top and Bottom 10 portfolios:



**Figure 43:** Pseudo-compounded returns across all quarters, for Top and Bottom 10 ANN predicted stocks, segregated by the success (“+”, blue, “in phase”) or lack of success (“-”, red, “out of phase”) of the ANN predictor. The x-axis shows the rank (top 1 through 10) of the equity to the left side of the vertical line, and the rank (bottom 1 through 10) on the right side; the y axis the pseudo-compounded returns.

There are likewise four components to the above chart each with ten (colored) elements:

- ▲ T10 ANN portfolio pseudo-compounded return during a (+) ”phase” of the predictor
- ▲ T10 ANN portfolio pseudo-compounded return during a (-) ”phase” of the predictor
- B10 ANN portfolio pseudo-compounded return during a (+) ”phase” of the predictor
- B10 ANN portfolio pseudo-compounded return during a (-) ”phase” of the predictor


These results from the ANN predictor clearly have many general features in common with the MGL predictor. However, there are a number of important differences. Chief among these is the fact that in the right half of the chart (representing the selection of stocks predicted to fall or do poorly, thus suitable for shorting), the circular marks representing the B10 aggregated data from both the positive and negative phases are relatively closely clustered and interpenetrating rather than separated. This is not an artifact of the semi-log representation: In contrast to the MGL predictor, four of ten ANN negative phase B10 lie above their positive phase counterparts. Thus, as a time-averaged statement, the ANN predictor only inverts its prediction for the **top** 10 – 100 equities during periods of lack of success; it simply loses any predictive ability for the bottom during these phases, which leaves long and hedged strategies intact.

Also, while the ANN predictor does not make as many exceptionally large correct predictions on the positive side, it makes far fewer mistakes in general, especially when selecting equities to sell short (generally a much more difficult task, especially in a rising market).

We may understand the ANN’s success vis à vis the MGL predictor as follows. First, we remind ourselves that by “success” we at times simply mean a lesser degree of failure (which for purposes of investment may call forth a high premium). Second, we note that the VL system, the MGL predictor and the top half of the ANN predictor are all subject to an inversion of their behavior; and place this observation in context of the above observations: The possibility that when they fail, the simpler models are falling prey to a change occurring in the market as a whole which they cannot anticipate or adjust for.

Third, in both the MGL and ANN predictors, inversions most often occur *delayed* relative to changes in the universe of equities, but this happens less often with the ANN. (Hence the anti-persistence evident in both as discussed above, but the greater anti-persistence for the MGL.) This makes sense, of course, as it explains why a large number of losing quarters arise—e.g., 9 of 29 for the H10 ANN predictor, 11 of 29 for the H10 MGL predictor, in the statistics provided above.

Studying this problem carefully requires a more rigorous definition of what constitutes a “state” than we have constructed, and what constitutes an inversion. However, a quick impression can be obtained from **Figure 44**:



**Figure 44:** ▲ Top 10 ● Bottom 10  
Quarter 16 T10 and B10 equities ranked by the ANN within full rank spectrum of 1443 stocks (of 1452 possible).

**Figure 44** shows the percent returns for the top and bottom 10 equities ranked by the ANN just for quarter 16 showing their actual ranks within all stocks available during that

quarter, during which the universe of (during that quarter) 1443 (out of the possible 1452) equities experienced a mean decline of 21.8%. (There are usually fewer stocks to rank within a given quarter because of the listing and delisting of corporations.) The preceding quarter saw a rise of 2.7% and the following quarter a rise of 24.3%. During this quarter, for this portfolio, the MGL predictor persisted in maintaining the “usual” ranking structure which had a (weakly) inverse relationship to the actual market during this quarter. While all four of the T10, B10 MGL and ANN *individual* portfolios lost—T10 MGL, B10 MGL, T10 ANN and B10 ANN—the T10 MGL did *worse* than the B10 MGL for a net H10 MGL loss during this one quarter of almost 4%.

On the other hand, while the ANN predictor does not resist the overall downward trend of the market and even—consistent with its overall volatility—amplifies it, the ANN H10 portfolio makes a significant gain. The ANN has created a winning rank-ordering structure for a hedged portfolio almost entirely from negative returns, in this case, in the face of (or composed out of) a declining market (following a series of quarters in which the overall market had risen, but less so, in each successive quarter). Note that in the long run, successful results are obtained for the T10-T100 portfolios on an absolute basis.

In the additional two quarters where the ANN succeeds and the MGL predictor fails, the market as a whole was similarly experiencing an overall decline (as in the one example just given). The successful adaptation in the structure of the predictor occurs “simultaneously” with the change in the equity market as a whole (within the discrete time unit of the procedure—one quarter).

The ANN predictor resists simply following the prior state, and rather seems to adapt when called for (with mixed success), with less of a delay. Hence, there are quarters when instead of failing because it does not keep up with a changing market, it partially adapts. The ANN predictor is much more frequently able than the MGL predictor to generate a ranking structure that is useful during generally adverse periods.

#### **Table 17 and**

**Table 18** present descriptive statistics comparing the MGL and ANN predictors for the T10 and B10 combined portfolios using the above aggregated measures. The figures in **Table 17** present *idealized* mean quarterly returns that would yield the pseudo-compounded returns in

**Table 18**. Note that whereas the idealized pseudo-compounded ANN returns are approximately equal to the actual compounded returns in the simulation, the idealized MGL returns are much larger. Thus fortunately, this comparison has tended conservatively (in this instance) to favor the MGL predictor.

The different results arise as follows: The product of 1 plus each specific quarterly difference between T10 and B10 is not equal to the product of 1 plus their average difference. I.e., in general, “pseudo-compounded” returns calculated as

$$\left(1 + \frac{1}{n} \sum_{i=1}^n x_i\right)^n = (1 + \bar{x})^n \neq \prod_{i=1}^n (1 + x_i),$$

the last expression being the actual cumulative returns. For example, if one of the specific factors  $x_j = -1$ , the final product

$\prod_{i=1}^n (1+x_i) = 0$ , whereas the second product would equal zero only if  $\frac{1}{n} \sum_{i=1}^n x_i = \bar{x} = -1$ .

The data in

**Table 18** present the data represented in **Figure 43**. Together they quantify the preceding discussion.

**Table 17:** Mean quarterly returns segregated by “ $\phi$ ”

<i>Quarterly Mean</i>	$\phi^+$	$\phi^-$
<i>No. Quarters</i>	20	9
ANN T10	0.20	-0.106
ANN B10	0.01	0.052
<b>ANN H10</b>	<b>0.19</b>	<b>-0.16</b>
<i>No. Quarters</i>	18	11
MGL T10	0.26	-0.10
MGL B10	0.01	0.11
<b>MGL H10</b>	<b>0.25</b>	<b>-0.21</b>

**Table 18:** Pseudo-compounded quarterly returns segregated by “ $\phi$ ”

<i>Ps. Compounded</i>	$\phi^+$	$\phi^-$	<i>Cumulative</i>	<i>Ann. ret.</i>
<i>No. Quarters</i>	20	9		
ANN T10	38.55	0.37		
ANN B10	1.22	1.58		
<b>ANN H10</b>	<b>32.72</b>	<b>0.21</b>	<b>6.98</b>	<b>30.7%</b>
<i>No. Quarters</i>	18	11		
MGL T10	67.21	0.31		
MGL B10	1.19	3.16		
<b>MGL H10</b>	<b>58.64</b>	<b>0.07</b>	<b>4.29</b>	<b>22.3%</b>

During 20 “positive” regime quarters (“ $\phi^+$ ”: predictor in phase with the universe of equities), the ANN predictor generated average returns of 0.191 for the H10 portfolio (combined T10 and B10). For 9 “negative” regime quarters (“ $\phi^-$ ”: predictor out of phase with the universe of equities), the ANN predictor generated average returns of  $-0.158$  for the H10. By contrast, the MGL predictor generated larger mean positive regime returns of 0.254 for the H10, but for fewer quarters (18) and also generated larger negative regime returns of  $-0.212$  for more quarters (11).


**Table 18** illustrates the equivalent pseudo-compounded results with 30.7% annualized return for the ANN predictor (close to actual) and 22.3% for the MGL (significantly higher than actual).

The most important point to be noted is that the relative success of the network predictor—and relative failure of the MGL predictor—is caused by the relatively large values in the negative phase columns, MGL B10 rows, compared to the ANN B10 rows, as these values are contrasted to their respective T10 rows.


*A10.6.4.2 T/B 100 = H100 Portfolios*

The above discussion can be extended throughout all the portfolios from 10 through 100 with (almost) uniform results. **Figure 45** and **Figure 46** are comparable to **Figure 42** and **Figure 43**; **Table 19** and **Table 20** are comparable to **Table 17** and

**Table 18**, presenting results for the T100, B100 and H100 portfolios, both ANN and MGL.



**Figure 45:** ANN T/B = H100 Portfolios



**Figure 46:** MGL T/B = H100 portfolios

**Table 19:** Mean quarterly returns segregated by phase  $\phi$

<i>Quarterly Mean</i>	$\phi^+$	$\phi^-$
<i>No. Quarters</i>	16	13
ANN T100	0.20	0.01
ANN B100	0.03	0.10
<b>ANN H100</b>	<b>0.17</b>	<b>-0.09</b>
<i>No. Quarters</i>	15	14
MGL T100	0.19	0.036
MGL B100	0.01	0.163
<b>MGL H100</b>	<b>0.18</b>	<b>-0.13</b>

**Table 20:** Pseudo-compounded quarterly returns segregated by phase  $\phi$ 

<i>Ps.Compounded</i>	$\phi^+$	$\phi^-$	<i>Cum. ret.</i>	<i>Ann. ret.</i>
<i>No. Quarters</i>	16	13		
ANN T100	14.91	1.11		
ANN B100	1.58	3.90		
<b>ANN H100</b>	<b>12.50</b>	<b>0.28</b>	<b>3.49</b>	<b>18.8%</b>
<i>No. Quarters</i>	15	14		
MGL T100	11.88	1.68		
MGL B100	1.10	8.80		
<b>MGL H100</b>	<b>11.56</b>	<b>0.15</b>	<b>1.72</b>	<b>7.8%</b>

Here again the same general pattern prevails, if somewhat attenuated. The superior relative and absolute performance of the ANN predictor is attributable primarily to its capacity to occasionally adapt rank structure in phase with shifts in the market. It therefore has a larger number of winning quarters. Consistent with this capacity, it tends to “resist” an apparent inversion at the bottom end of the ranking during losing quarters, thus incurring smaller losses than the MGL predictor.

## **A10.7 Discussion**

The prediction method outlined above has many complex features only a few of which have been treated. We discuss here a small number of key points, including problems.

### **A10.7.1 Data Cleaning and Errors**

As discussed previously, the original VL data has an extremely large number of errors. It was not possible to obtain meaningful results using uncleaned data and the process of cleaning data to an acceptable level is very resource-intensive. Because of this, it has not yet been possible to run simulations on all possible monthly, weekly and daily cycles within a quarter. Randomly selected and targeted runs performed on partial sets show insignificant differences with the results presented here. Identical simulations performed at shorter data intervals (monthly, weekly, daily) should yield similar results with more robust statistics. But as the interval shortens, the data problem worsens. Furthermore, with respect to earnings, the quarterly time period is the only natural time step.

### **A10.7.2 Why Does the ANN Predictor Work?**

There are a sufficient number of occasions when the ANN predictor seems able to adapt and anticipate an inversion in the market, i.e., a change in the structure of the VL universe of equities it is attempting to rank, and so undergoes a change in the rank-ordering it establishes. Rather than mimic the rank-ordering of the immediately prior quarter (as does the MGL predictor), the ANN has the option of looking back at the rankings by change in price and by change in earnings for all 1452 equities over ten quarters. For each equity it then assesses the impact of each of these twenty parameters on the most recent change in price. But every other input is itself an earlier period’s change in price (merely transformed into a rank), that was itself once an output. Not every quarter most closely mimics the prior quarter. There are in fact many quarters when the average change in price over the entire universe of equities more closely mimics what happened two quarters prior, or three. (A MGL predictor based on the rank ordering of two quarters prior still outperforms a random ranking.) To some extent, the ANN is apparently able to

properly weight the changing relative contribution of these prior quarters, including the contribution of changes in earnings: An identical predictor that excludes earnings inputs still outperforms the MGL predictor, but underperforms this ANN predictor and is much more volatile. In particular, because the ANN tacitly develops a nonlinear relationship among the input variables, it may detect as a “pattern”, a pattern in the “change-of-state” of the market, as it were, if such a pattern exists or seems to exist.

This, of course, raises a caution: With only 29 quarters, it is not possible to differentiate between a “pattern of change of state” that actually exists and one that only seems to exist, even if the internal structures within each quarter are very robust. In other words, the ANN predictor may indeed be better at detecting global patterns of the above kind than the MGL predictor. But one can only *argue* that the markets undergoing large scale declines during the quarters captured by the ANN (and missed by the MGL) did so as part of a pattern detectable by the ANN rather than a pattern invented by it. The success of the ANN in generating large returns provides the evidence for the former.

### **A10.7.3 Relation between the ANN and the VL ranking system**

The behavior of the ANN in generating rankings shows many similarities to the VL system proper which, as noted in the introduction, has been the object of significant attention because of its apparently anomalous success. One similarity is the ability of the ANN to generate significant positive returns; another is the fact that when it fails, it, too, does not simply lose predictive ability, it produces rank-orderings that are inverted relative to the actual price changes and so creates large losses. Nonetheless, better than a MGL predictor, or the VL system proper, it is able both to anticipate the need for a different rank ordering, on occasion, and to do a better job in resisting failure when identifying stocks for short-selling. On balance, at least in this study sample, the ANN is able to improve upon the VL approach proper and generate net positive returns over the long run in excess of a buy and hold strategy and sufficient to overcome transaction costs. (These are kept to a minimum in the quarterly trading process employed). These results suggest that at least part of the power inherent in the VL approach is in wide use of rank orderings as a general method for coarse-graining financial data.

## References

1. Bandura, A., *Self-efficacy: the exercise of control*. 1997: WH Freeman New York.
2. Taylor, S.E. and J.D. Brown, *Illusion and well-being: A social psychological perspective on mental health*. Psychological Bulletin, 1988. **103**(2): p. 193-210.
3. Langer, E., *The illusion of control*. Journal of Personality and Social Psychology, 1975. **7**: p. 185—208.
4. Challet, D. and Y.C. Zhang, *Emergence of cooperation and organization in an evolutionary game*. Physica A, 1997. **246**(3): p. 407.
5. Challet, D., M. Marsili, and Y.C. Zhang, *Minority Games*. 2005, New York: Oxford University Press.
6. Parrondo, J.M.R., *EEC HC&M Network on Complexity and Chaos (ERBCHRX-CT940546)*. ISI (Torino, Italy), 1996.
7. Harmer, G.P. and D. Abbott, *Parrondo's Paradox*. Statistical Science, 1999. **14**(2): p. 206-213.
8. Bastiat, F., *Economic Sophisms*, ed. A. Goddard. 1966 (1845), Irvington-on-Hudson, NY: The Foundation for Economic Education, Inc.
9. Gray, J., *Hayek on Liberty*. 1998, London: Routledge.
10. Hayek, F.A., “*The Use of Knowledge in Society*,”. American Economic Review, 1945. **35**(S 519): p. 530.
11. Read, L., *I, Pencil*. Irvington-on-Hudson, NY: Foundation for Economic Education, 1998.
12. Von Neumann, J., *Von Neumann's Self-Reproducing Automata*, in *Essays on Cellular Automata*, A. Burks, Editor. 1970. p. 4-65.
13. Flake, G.W., *The computational beauty of nature*. 1998: MIT Press Cambridge, MA, USA.
14. Gell-Mann, M., *Complex Adaptive Systems*, in *Complexity: Metaphors, Models, and Reality*. 1999, Santa Fe Institute: Santa Fe.
15. Satinover, J., *The Quantum Brain*. 2000, New York: Wiley, Inc.

16. Hopfield, J.J., *Neural Networks and Physical Systems with Emergent Collective Computational Abilities*. Proceedings of the National Academy of Sciences of the USA, 1982. **79**: p. 2554-2558.
17. Widrow, B. and M.A. Lehr, *30 years of adaptive neural networks: perceptron, Madaline, and backpropagation*. Proceedings of the IEEE, 1990. **78**(9): p. 1415-1442.
18. Hoppensteadt, F.C. and E.M. Izhikevich, *Weakly Connected Neural Networks*. Applied Mathematical Sciences, ed. J.E. Marsden, L. Sirovich, and F. John. Vol. 126. 1997, New York: Springer.
19. Hoppensteadt, F. and E. Izhikevich, *Canonical models in mathematical Neuroscience*. Documenta Mathematica, 1998. **ICM III**: p. 593-599.
20. Hoppensteadt, F.C. and E.M. Izhikevich, *Oscillatory Neurocomputers with Dynamic Connectivity*. Phys. Rev. Lett., 1999. **82**(14): p. 2983-2986.
21. Merton, R.K., *The Unanticipated Consequences of Purposive Social Action*. American Sociological Review, 1936. **1**(6): p. 894-904.
22. Robins, P.K., *A Comparison of the Labor Supply Findings from the Four Negative Income Tax Experiments*. The Journal of Human Resources, 1985. **20**(4): p. 567-582.
23. Groeneveld, L.P., N.B. Tuma, and M.T. Hannan, *The Effects of Negative Income Tax Programs on Marital Dissolution*. The Journal of Human Resources, 1980. **15**(4): p. 654-674.
24. Keeley, M.C., *The Effects of Negative Income Tax Programs on Fertility*. The Journal of Human Resources, 1980. **15**(4): p. 675-694.
25. Keeley, M.C., *The Effect of a Negative Income Tax on Migration*. The Journal of Human Resources, 1980. **15**(4): p. 695-706.
26. Murray, C.A., *Losing Ground: American Social Policy, 1950-1980*. 1984: Basic Books.
27. Cassar, G. and B. Gibson, *Budgets, Internal Reports and Manager Forecast Accuracy*. 2007, SSRN eLibrary <http://ssrn.com/paper=939332>.
28. Cassar, G., *Information Acquisition and the Rationality of Entrepreneur Expectations*. Babson College Working Papers, 2007.

29. Cassar, G., *Are Individuals Entering Self-Employment Overly-Optimistic? An Empirical Test of Plans and Projections on Nascent Entrepreneur Expectations*. Babson College Working Papers, 2008. <http://ssrn.com/paper=945206>
30. Cassar, G. and B. Gibson, *Forecast Rationality in Small Firms*. Journal of Small Business Management, 2007. **45**(3): p. 283-302.
31. Gaylord, R.J. and L.J. D'Andria, *Simulating Society: A Mathematica Toolkit for Modeling Socioeconomic Behavior*. 1998: Springer.
32. De Martino, A., M. Marsili, and R. Mulet, *Adaptive drivers in a model of urban traffic*. Europhysics Letters, 2004. **65**(2): p. 283-289.
33. Coile, C., et al., *Fiscal Effects of Social Security Reform in the United States*. 2003: Center for Retirement Research at Boston College.
34. Fiske, A.P., *A proxy theory of emotions as motives that represent the value of social relationships*. Unpublished paper, 2007.
35. Uchitelle, L., *The richest of the rich, proud of a new gilded age*, in *New York Times*. 2007: New York. p. 1,20.
36. Solomon, K., *Ecclesiastes 9:11*, in *The Bible, King James Version*, V.h. committees, Editor. unknown.
37. Malevergne, Y., A.I. Saichev, and D. Sornette, *Zipf's Law for Firms: Relevance of Birth and Death Processes*. Journal of Economic Theory, SSRN eLibrary: <http://ssrn.com/paper=1083962> 2008.
38. Mackay, C., *Extraordinary Popular Delusions and the Madness of Crowds*. 2000: Templeton Foundation Press.
39. Andersen, J.V. and D. Sornette, *The \$-game*. Eur. Phys. J.B, 2003. **31**: p. 141-145.
40. Satinover, J.B. and D. Sornette, *Illusion of Control in a Brownian Game*. Arxiv preprint physics/0703048, 2007.
41. Hart, M., et al., *Crowd-anticrowd theory of multi-agent market games*. Eur. Phys. J. B, 2001. **20**(4): p. 547-550.
42. Doran, J.S. and C. Wright, *What Really Matters When Buying and Selling Stocks?*, in *Working paper of Florida State University*. 2007.
43. Hart, M.L., P. Jefferies, and N.F. Johnson, *Dynamics of the time horizon minority game*. Physica A, 2002. **311**(1): p. 275-290.

44. D'Hulst, R. and G.J. Rodgers, *The Hamming distance in the minority game*. Physica A, 1999. **270**(3): p. 514-525.
45. Cavagna, A., *Irrelevance of memory in the minority game*. Physical Review E, 1999. **59**(4): p. 3783-3786.
46. Challet, D. and M. Marsili, *Relevance of memory in minority games*. Physical Review. E, 2000. **62**(2 Pt A): p. 1862-1868.
47. Ho, K.H., et al., *Memory is relevant in the symmetric phase of the minority game*. Phys. Rev. E., 2005. **71**: p. 066120.
48. Lee, C.Y., *Is memory in the minority game irrelevant?* Phys. Rev. E, 2001. **64**(1 Pt 2): p. 015102.
49. Ho, K.H., et al., *Memory is relevant in the symmetric phase of the minority game*. Physical Review E, 2005. **71**(6): p. 66120.
50. Li, Y., R. Riolo, and R. Savit, *Evolution in minority games (II). Games with variable strategy spaces*. Physica A, 2000. **276**(1): p. 265-283.
51. Andrecut, M. and M.K. Ali, *Q learning in the minority game*. Phys. Rev. E, 2001. **64**(6): p. 67103.
52. Araujo, R.M. and L.C. Lamb, *Towards Understanding the Role of Learning Models in the Dynamics of the Minority Game*. Proceedings of the 16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'04), 2004. **00**: p. 727-731.
53. Chau, H.F. and F.K. Chow, *How To Attain Maximum Profit In Minority Game?* Arxiv preprint nlin.AO/0112049, 2001.
54. Sysi-Aho, M., A. Chakraborti, and K. Kaski, *Intelligent minority game with genetic crossover strategies*. Eur. Phys. J. B, 2003. **34**(3): p. 373-377.
55. Sysi-Aho, M., A. Chakraborti, and K. Kaski, *Searching for good strategies in adaptive minority games*. Phys. Rev. E, 2004. **69**(3): p. 36125.
56. Menche, J. and J.R.L. de Almeida, *It is worth thinking twice or Improving the performance of minority games*. Arxiv preprint cond-mat/0308181, 2003.
57. Challet, D., M. Marsili, and Y.-C. Zhang, *Modeling market mechanism with minority game*. Physica A, 2000. **276**(1-2): p. 284-315.
58. Marsili, M. and D. Challet, *Trading behavior and excess volatility in toy markets*. Arxiv preprint cond-mat/0004376, 2000.

59. Challet, D., M. Marsili, and R. Zecchina, *Statistical mechanics of systems with heterogeneous agents: minority games*. Phys. Rev. Lett., 2000. **84**(8): p. 1824-1827.
60. Duffy, J. and E. Hopkins, *Learning, Information and Sorting in Market Entry Games: Theory and Evidence*. forthcoming, Games and Economic Behavior, 2004.
61. Parrondo, J.M.R., *How to cheat a bad mathematician*, in *EEC HC&M Network on Complexity and Chaos (ERBCHRX-CT940546)*. 1996, ISI, Torino, Italy: Torino.
62. Smoluchowski, M., « *Experimentell nachweisbare der üblichen Thermodynamik widersprechende Molekularphänomene* », Phys. Z., 1912. **13**: p. 1069-1080.
63. Magnasco, M.O., *Chemical Kinetics is Turing Universal*. Physical Review Letters, 1997. **78**(6): p. 1190-1193.
64. Pyke, R., *on random walks and diffusions related to Parrondo's games*, in *Mathematical Statistics and Applications: Festschrift for Constance Van Eeden*. 2003.
65. Kinderlehrer, D. and M. Kowalczyk, *Diffusion-Mediated Transport and the Flashing Ratchet*. Archive for Rational Mechanics and Analysis, 2002. **161**(2): p. 149-179.
66. Moraal, H., *Counterintuitive behaviour in games based on spin models*. J.Phys.A: Math.Gen, 2000. **33**(16): p. L203-L206.
67. Harmer, G.P., et al., *Brownian ratchets and Parrondo's games*. Chaos, 2001. **11**(3): p. 705-714.
68. Maslov, S. and Y.C. Zhang, *Optimal Investment for Risky Assets*. International Journal of Theoretical and Applied Finance, 1998. **1**(3): p. 377-387.
69. Boman, M., S.J. Johansson, and D. Lyback, *Parrondo strategies for artificial traders*. *Intelligent Agent Technology* (eds. N. Zhong, J. Liu, S. Ohsuga & J. Bradshaw), 150–159. World Scientific. Also available at cs/0, 2001. **204**: p. 051.
70. Almgren, W.S. and M. Boman, *An Active Agent Portfolio Management Algorithm*, in *Artificial Intelligence and Computer Science*, S. Shannon, Editor. 2005, Nova Science Publishers. p. 123-134.
71. Satinover, J.B. and D. Sornette, *Illusion of control in a Brownian game*. Physica A, 2007. **386**(1): p. 339-344.

72. Parrondo, J.M.R., G.P. Harmer, and D. Abbott, *New Paradoxical Games Based on Brownian Ratchets*. Physical Review Letters, 2000. **85**(24): p. 5226-5229.
73. Kay, R.J. and N.F. Johnson, *Winning combinations of history-dependent games*. Phys. Rev. E, 2003. **67**(5): p. 56128.
74. Dinis, L., *Optimal strategies in collective Parrondo games*. Europhysics Letters, 2003. **63**(3): p. 319-325.
75. Parrondo, J.M.R., et al., *Paradoxical games, ratchets, and related phenomena*, in *Advances in Condensed Matter and Statistical Mechanics*, E. Korutcheva and R. Cuerno, Editors. 2003, Nova Science Publishers.
76. Huber, J., M. Kirchler, and M. Sutter, *Is more information always better? Experimental financial markets with cumulative information*. Journal of Economic Behavior and Organization, forthcoming, 2006.
77. Eurich, C.W. and K. Pawelzik, *Optimal Control Yields Power Law Behavior*, in *Artificial Neural Networks: Formal Models and Their Applications - ICANN 2005*. 2005, Springer: Berlin. p. 0302-974.
78. Berczuk, S.P. and B. Appleton, in *Software Configuration Management Patterns: Effective Teamwork, Practical Integration*. 2002, Addison-Wesley.
79. Malkiel, B.G., *A Random Walk Down Wall Street.: the Time-tested Strategy for Successful Investing*. 2003: WW Norton & Company.
80. Shiller, R.J., *Market Volatility*. 1992: MIT Press.
81. Grandin, T. and C. Johnson, *Animals in Translation: Using the Mysteries of Autism to Decode Animal Behavior*. 2005: Scribner.
82. Kozłowski, P. and M. Marsili, *Statistical mechanics of the majority game*. Journal of Physics A, 2003. **36**: p. 11725-11737.
83. Satinover, J.B. and D. Sornette, *Cycles, determinism and persistence in agent-based games and financial time-series*. submitted to Quantitative Finance, 2008.
84. Satinover, J.B. and D. Sornette, *"Illusion of control" in Time-Horizon Minority and Parrondo Games*. Eur. Phys. J. B, 2007. **60**(3): p. 369-384.
85. Pring, M.J., *Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points*. 2002: McGraw-Hill.
86. Fama, E.F., *Efficient Capital Markets: II*. The Journal of Finance, 1991. **46**(5): p. 1575-1617.

87. Fama, E.F., *Efficient Capital Markets: A Review of Theory and Empirical Work*. The Journal of Finance, 1970. **25**(2): p. 383-417.
88. Satinover, J., Sornette, D., *Illusion of Control in Time Horizon Minority and Parrondo Games*. Physica A, 2007. **386**: p. 339-344.
89. Satinover, J., Sornette, D., *Illusory and genuine control in agent-based games*. In preparation, 2008.
90. Jefferies, P., M.L. Hart, and N.F. Johnson, *Deterministic dynamics in the minority game*. Phys. Rev. E, 2002. **65**(1 Pt 2): p. 016105.
91. Cvitanovic, P., I. Percival, and A. Wirzba, *Chaos focus issue on periodic orbit theory*. Chaos, 1992. **2**(1).
92. Eckhardt, B., *Periodic Orbit Theory, to appear in Proceedings of the International School of Physics "Enrico Fermi" 1991, Course CXIX "Quantum Chaos"*, G. Casati, I. Guarneri, and U. Smilansky (eds), 1992.
93. Metzler, R., *Antipersistent binary time series*. J.Phys.A, 2002. **35**(25): p. 721-730.
94. Israeli, N. and N. Goldenfeld, *Computational Irreducibility and the Predictability of Complex Physical Systems*. Physical Review Letters, 2004. **92**(7): p. 74105.
95. Israeli, N. and N. Goldenfeld, *Coarse-graining of cellular automata, emergence, and the predictability of complex systems*. Physical Review E, 2006. **73**(2): p. 26203.
96. Ryden, T., T. Terasvirta, and S. Asbrink, *Stylized Facts of Daily Return Series and the Hidden Markov Model*. Journal of Applied Econometrics, 1998. **13**(3): p. 217-244.
97. Knab, B., et al., *Model-based clustering with Hidden Markov Models and its application to financial time-series data*. Between Data Science and Applied Data Analysis, M. Schader, W. Gaul, and M. Vichi, eds., Springer, 2003: p. 561-569.
98. Andersen, J.V. and D. Sornette, *A mechanism for pockets of predictability in complex adaptive systems*. Europhysics Letters, 2005. **70**(5): p. 697-703.
99. Lamper, D., S.D. Howison, and N.F. Johnson, *Predictability of Large Future Changes in a Competitive Evolving Population*. Physical Review Letters, 2001. **88**(1): p. 17902.
100. Gou, C., *Predictability of Shanghai Stock Market by Agent-based Mix-game Model*. Proceedings of the International Conference on Neural Networks and Brain, 2005 (ICNN&B'05), 2005. **3**.

101. Papageorgiou, C.P. *Proceedings of the IEEE/IAFE. in Computational Intelligence for Financial Engineering (CIFEr)*. 1997: unpublished.
102. Zhang, Y.C., *Toward a theory of marginally efficient markets*. Physica A, 1999. **269**(1): p. 30-44.
103. Lo, T.S., P.M. Hui, and N.F. Johnson, *Theory of the evolutionary minority game*. Phys. Rev. E, 2000. **62**(3): p. 4393-4396.
104. Hod, S. and E. Nakar, *Self-Segregation versus Clustering in the Evolutionary Minority Game*. Physical Review Letters, 2002. **88**(23): p. 238702.
105. Chen, K., B.H. Wang, and B. Yuan, *Adiabatic theory for the population distribution in the evolutionary minority game*. Physical Review E, 2004. **69**(2): p. 25102.
106. Jefferies, P., et al., *Mixed population Minority Game with generalized strategies*. J. Phys. A, 2000. **33**(43): p. L409-L414.
107. Johnson, N.F., et al., *Enhanced winnings in a mixed-ability population playing a minority game*. J. Phys. A, 1999. **32**(38): p. L427-L431.
108. D'Hulst, R. and G.J. Rodgers, *Strategy selection in the minority game*. Physica A, 2000. **278**(3): p. 579-587.
109. Zhong, L.X., et al., *Effects of contrarians in the minority game*. Phys. Rev. E, 2005. **72**(2 Pt 2): p. 026134.
110. Hart, M., et al., *Generalized strategies in the minority game*. Phys. Rev. E., 2001. **63**(1 Pt 2): p. 017102.
111. Dert, C. and B. Oldenkamp, *Optimal Guaranteed Return Portfolios and the Casino Effect*. Operations Research, 2000. **48**(5): p. 768-775.
112. Chun, S.H. and Y.J. Park, *Dynamic adaptive ensemble case-based reasoning: application to stock market prediction*. Expert Systems With Applications, 2005. **28**(3): p. 435-443.
113. Kalpazidou, S. and J.E. Cohen, *Orthogonal cycle transforms of stochastic matrices*. Circuits, systems, and signal processing, 1997. **16**(3): p. 363-374.
114. Minping, Q. and Q. Min, *Circulation for recurrent markov chains*. Probability Theory and Related Fields, 1982. **59**(2): p. 203-210.
115. Kalpazidou, S.E. and N.E. Kassimatis, *Markov chains in Banach spaces on cycles*. Circuits, Systems, and Signal Processing, 1998. **17**(5): p. 637-652.

116. Kalpazidou, S., *On the Weak Convergence of Sequences of Circuit Processes: A Probabilistic Approach*. Journal of Applied Probability, 1992. **29**(2): p. 374-383.
117. Bernhard, A., *How to Use the Value Line Investment Survey: A Subscriber's Guide*. 1984: Value Line.
118. Shelton, J.P., *The value line contest: A test of the predictability of stock price changes*. J. Bus., 1967. **40**(3): p. 254–269.
119. Black, F., *Yes, Virginia, There Is Hope: Tests of the Value Line Ranking System*. Financial Analysts Journal, 1973. **29**(5): p. 10–14.
120. Feingold, J.H., *A reevaluation of the value line investment strategy: can active common stock portfolio management produce superior returns?*, in *Sloan School of Management*. 1979, Massachusetts Institute of Technology: Cambridge, MA. p. 158.
121. Kaplan, R.S. and R.L. Weil, *Risk and the value line contest*. Financial Anal. J., 1973. **29**(4): p. 56-61.
122. Holloway, C., *A Note on Testing an Aggressive Investment Strategy Using Value Line Ranks*. Journal of Finance, 1981. **36**(3): p. 711-719.
123. Stickel, S.E., *The Effect of Value Line Investment Survey Rank Changes on Common Stock Prices*. Journal of Financial Economics, 1985. **14**(1): p. 121-143.
124. Huberman, G. and S. Kandel, *Value line rank and firm size*. J. Bus., 1987. **60**(4).
125. Affleck-Graves, J. and R. Mendenhall, *The relation between the Value Line enigma and post-earnings-announcement drift*. Journal of Financial Economics, 1992. **31**(1): p. 75-96.
126. Choi, J., *The Value Line Enigma: The Sum of Known Parts*. Journal of Financial and Quantitative Analysis, 2000. **35**(3): p. 485-498.
127. Szakmary, A.C., C.M. Conover, and C. Lancaster, *An examination of Value Line's long-term projections*. Journal of Banking and Finance, 2008. **32**(5): p. 820-833.
128. Philbrick, D.R. and W.E. Ricks, *Using Value Line and IBES Analyst Forecasts in Accounting Research*. Journal of Accounting Research, 1991. **29**(2): p. 397-417.
129. Yemma, J., *A Former Drama Critic Who Plots the Action on Wall Street.*, in *Christian Science Monitor, Business Section*. 1983. p. 14.

130. Brown, L.D., *Earnings Surprise Research: Synthesis and Perspectives*. Financial Analysts Journal, 1997. **53**(2): p. 13–19.
  131. Battalio, R.H. and R.R. Mendenhall, *Earnings expectations, investor trade size, and anomalous returns around earnings announcements*. Journal of Financial Economics, 2005. **77**(2): p. 289-319.
  132. Fama, E.F., *Efficient capital markets: a review of theory and empirical work*. Journal of Finance, 1970. **25**: p. 383-417.
  133. Fama, E.F., *Efficient capital markets: II* Journal of Finance, 1991. **46**(5): p. 1575-1617.
  134. Benesh, G.A. and S.B. Perfect, *An analysis of Value Line's ability to forecast long-run returns*. Journal of Financial And Strategic Decisions, 1997. **10**(2).
  135. Copeland, T. and D. Mayers, *The Value Line Enigma (1965-1978): A Case Study of Performance Evaluation Issues*. Journal of Financial Economics, 1982. **10**(3): p. 289-321.
  136. Graham, J.R., *Herding among Investment Newsletters: Theory and Evidence*. The Journal of Finance, 1999. **54**(1): p. 237-268.
  137. Cullen, L.R., *What Happened to Value Line?* Money Magazine, 2000: p. 97–101.
  138. Amin, G. and H. Kat. *Hedge Fund Performance 1990-2000: Do the "Money Machines" Really Add Value?* in *EFMA 2001 Lugano Meetings*. 2001. Lugano, Switzerland.
  139. deSouza, C. and S. Gokcan, *Hedge Fund Volatility: It's Not What You Think It Is*. AIMA Journal, 2004(Sept.).
  140. Amin, G. and H. Kat, *Hedge Fund Performance 1990-2000: Do the "Money Machines" Really Add Value*, 2003: p. 251-274.
  141. Satinover, J.B. and D. Sornette, *Illusory versus genuine control in agent-based games*. submitted to Eur. Phys. J. B., 2008.
-