THE EUROPEAN
PHYSICAL JOURNAL B

Regular Article

Illusory versus genuine control in agent-based games

J.B. Satinover[a] and D. Sornette[b]

Dept. of Management, Technology and Economics, ETH Zurich, 8032 Zurich, Switzerland

Abstract. In the Minority, Majority and Dollar Games (MG, MAJG, $G) agents compete for rewards, acting in accord with the previously best-performing of their strategies. Different aspects/kinds of real-world markets are modelled by these games. In the MG, agents compete for scarce resources; in the MAJG agents imitate the group to exploit a trend; in the $G agents attempt to predict and benefit both from trends and changes in the direction of a market. It has been previously shown that in the MG for a reasonable number of preliminary time steps preceding equilibrium (Time Horizon MG, THMG), agents' attempt to optimize their gains by active strategy selection is "illusory": the hypothetical gains of their strategies is greater on average than agents' actual average gains. Furthermore, if a small proportion of agents deliberately choose and act in accord with their seemingly worst performing strategy, these outperform all other agents on average, and even attain mean positive gain, otherwise rare for agents in the MG. This latter phenomenon raises the question as to how well the optimization procedure works in the THMAJG and TH$G. We demonstrate that the illusion of control is absent in THMAJG and TH$G. This provides further clarification of the kinds of situations subject to genuine control, and those not, in set-ups a priori defined to emphasize the importance of optimization.

PACS. 89.75.-k Complex systems – 89.65.Gh Economics; econophysics, financial markets, business and management – 02.50.Le Decision theory and game theory

1 Introduction

"Illusion of control" [1] 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. We have previously shown [2] that in one of the most extensively-studied agent-based game-models of markets, the minority game (MG), agents' control is illusory in the following sense: the mean actual performance of all agents averaged over many different initial conditions is poorer than the mean hypothetical performance of all their given strategies. The finding is striking because at each time-step agents deploy that strategy with the best hypothetical performance. This finding is most generally true under the following conditions: the initial state is iterated for a "reasonable" number of time-steps ($\leqslant 2000$ say) short of equilibrium ($\geqslant 5000$ say) at which point a rolling window of cumulative strategy scores is maintained, which window distinguishes the Time Horizon MG (THMG) from the MG proper. In the THMG, the illusion is observed for all m (m is the number of bits of binary history in agents' memory and an "m-bit history" consists of a string of 1's and 0's of length m. 1 indicates

that a minority of agents have adopted action "-1", 0 indicates that they have adopted action "$+1$"). The finding is less generally true in the MG strictly speaking, in which the system is always allowed to run to equilibrium. In that case, it is true for all m, except at the so-called critical point $\alpha_c = \frac{2^{m_c}}{N} = 0.34$, where $\alpha \equiv \frac{2^m}{N}$ in general and N is the number of agents. And even at α_c the illusion of control is just barely overturned. For example with $N = 31$, $3 < m_c = 3.4 < 4$ (therefore integer $m_c \equiv 4$). However, equilibrium in the MG proper is only reached after $\sim 100 \times 2^{m+1}$ iterations away from α_c and orders or magnitude more near α_c, therefore for $N = 31$ at $m_c = 4$, $t_{eq} \gg 3200$ iterations. Arguably, no real-world market is sufficiently stationary to attain such an equilibrium state.

Another important finding presented in [2] is the surprising fact that, if a small proportion of agents ($\leqslant 0.15$, e.g., 3 of 31) deploy their previously *worst* performing strategy, these on average outperform all others. In fact they relatively consistently attain average net positive gain, otherwise rare in the MG as the mean gain for both strategies and agents is in general negative due to the minority rule. Note that such an inversion of the selection rule is a symmetric alteration with no privileging nor increase in agent computational capacity (see Appendix).

a e-mail: jsatinover@ethz.ch
b e-mail: dsornette@ethz.ch

The success of "anti-optimizing" agents is a further marker of the "illusory" nature of the standard optimization rule. It also raises the following question: is an inverted rule equivalent to agents playing a Majority Game (MAJG)? As they are not optimizing to be in the minority, it seems that agents 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. Does the MAJG thus also demonstrate an "illusion of control"? Since agent gain in the $G also follows a majority rule, albeit time-lagged, would "illusion of control" be found in it, too?

The goal of the present paper is to clarify these questions and 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 a follow-up to this paper [3], 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.

Our main result with respect to the THMG is detailed in [2] and Section 1. above: strategies' mean hypothetical performance exceeds agents' mean actual performance. Agents should thus behave in risk-averse fashion by switching randomly between strategies rather than optimizing. Two conditions must hold for the statement to be false: (1) $m \geqslant m_c$; (2) the system must be allowed to reach equilibrium ($t \geqslant \tau_{eq}$). Condition (2) requires an exceedingly large number of preliminary steps before agent selection begins, and orders of magnitude more steps if $m \approx m_c$). In the MG proper the situation is more complex and is discussed at length in [2].

In the MAJG and $G the reverse holds: optimization greatly enhances agent over strategy gain, both being positive.

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.

2 Minority, Majority and $ Games

2.1 Definition and overview of Minority Games (MG) and Time-Horizon MG (THMG)

In the MG, those agents gain whose actions are in the minority subset at each time-step. Agents in the complementary (majority) subset lose. As discussed in [2], the MG is an agent-based model of financial markets that violates the "rational expectations" assumption of standard economic theory [4]: when every agent behaves in accord with its most rational expectation, the collective action of all agents makes the expectation false. Common expectations cancel, leading to anti-persistent behavior both for the collective and for individuals [5,6].

In brief, the MG proceeds as follows: at each time step t, each agent $i \in \{1, 2, \ldots, N\}$ chooses one of two alternative actions $a_i(t) \in \{-1, 1\}$. This yields a value $A(t) = \sum_{i=1}^{N} a_i(t)$, and binary series $D(0) \ldots D(t-1), D(t)$ with $D(t) = 1/2 \left[\mathrm{Sgn}\left(2A - N\right) + 1 \right] \in \{0, 1\}$. Each agent is randomly endowed with S strategies (selected from 2^{2^m} possible strategies) assigned for always at the beginning of the game (quenched disorder state $\hat{\Omega}$). Agents choose their actions at each t based on a history $\mu(t)$ consisting of the most recent m digits of $D(t)$ in accord with one of its S strategies $s_i^r, r \in \{1, 2, \ldots S\}$: a strategy s_i^r maps $\mu(t) \to a_i(t)$. Strategies accumulate $+1$ hypothetical point each time they win, unboundedly, and -1 hypothetical point each time they lose. Agents whose strategies lead to a win accumulate $+1$ real point, and those that lose accumulate -1 point. A running score table is kept for every strategy and agent. The strategy chosen by each agent at t is the one which would have been in the minority most frequently until then. This is the key optimization step. If an agent's best strategy predicts $D(t) = 1$ (resp. 0), she will take the action -1 (resp. $+1$) to try to be in the minority. Ties are broken by a fair coin toss. If hypothetical points are summed over a rolling window of bounded length τ up to the last information available at t, the game is the "Time Horizon MG". For $\tau \geqslant \tau_{eq}$, the THMG is equivalent to the MG, where τ_{eq} is the equilibrium time for the MG.

The distinguishing characteristic of the MG is the corresponding instantaneous payoff of agent I, given by $-a_i(t)A(t)$ – the minus sign encodes the minority rule (and similarly for each strategy for which it is added to the $\tau - 1$ previous payoffs. The payoff may also be $-\mathrm{Sgn}[a_i(t)A(t)] = -a_i(t)D(t)$. This does not alter the dynamics of the game. If $-a_i(t)A(t) < 0 \Rightarrow -\mathrm{Sgn}[a_i(t)A(t)] < 0$, agent i is rewarded. Agents in the MG and THMG thus try to be anti-imitative. 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, assuming that the system is first allowed to reach equilibrium [5]. The phenomenon discussed in [2] 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 nonetheless underperform the mean of their own measured strategy performance and do so for reasonable lengths of τ at all m. (In the MG, 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 m_c (where $\alpha_c \equiv 2^m/N$), but in this region the number of steps to equilibrium is extremely large [5].)

More specifically, for any given realization, a minority of agents outperform their strategies and the majority of other agents (some may also achieve net positive gain, if rarely). In the MG proper, however, τ is

unbounded and a stationary state is reached at some very large $\tau_{eq} \geqslant 2^m \times 200$. At equilibrium, the largest possible number of agents "freeze" their choice of strategy: one virtual strategy score attains a permanently higher value than any other (the fraction of agents susceptible of freezing is not only determined by N, m and S, but by the specific initial allocation of strategies, i.e., the quenched disorder tensor $\hat{\Omega}$). 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 exactly as well as their best. In [2] and here we focus on the THMG with real-world markets in mind. Such time series are non-stationary, and trading strategies are weakened if they incorporate an unbounded history of prior strategic success or failure. Furthermore, agents do not take account of the impact of their choice on the probable minority state. We refer to such agents as "standard".

2.2 Definition and overview of the Majority Game (MAJG)

A MAJG 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 $g_i^{min}(t) = -\text{Sgn}[a_i(t)A(t)]$ whereas $g_i^{maj}(t) = +a_i(t)A(t)$ or $g_i^{maj}(t) = +\text{Sgn}[a_i(t)A(t)]$. The plus sign means that agents are rewarded when they act with the majority of agents. 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 $\bar{G}^{maj}(t) = \frac{1}{N} \sum_{i=1}^{N} g_i^{maj}(t) > 0$. Cumulative wealth tends to increase over time.

In the MG, $A(t)$ is typically anti-persistent (as is $D(t)$), paralleling the anti-imitative behavior of agents. In the MAJG, $A(t)$ is typically persistent, paralleling the imitative behavior of agents.

2.3 Definition and overview of the Dollar Game ($G)

The $G [8] captures more accurately the changing actions of traders, but in a framework as close as possible to the

MG. We wrote that agents in the MG are "contrarian" in trying to do what most others are not. But this is not what true contrarians do: They attempt to be in the minority when a market falls, but also at a turning point. And they attempt to be in the majority when the market rises. This is what non-contrarian traders are also attempting. Indeed every trader tries to do this. Contrarians differ from conformists in their *reasoning* about 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. A similar correction to "conformist" can be made.

An agent with "real world" behaviors is one that tries to join the majority when the market begins a rise; He tries to join the minority when the market begins a decline – i.e., at convex or concave inflection points respectively. This behavior is captured by the following change in the agent gain: $g_i^\$(t) = +a_i(t-1)A(t)$ or $g_i^\$(t) = +\text{Sgn}[a_i(t-1)A(t)]$. Action at $t-1$ is interpreted as a judgment about $\text{Sgn}[A(t)]$, and determines whether an agent gains or loses. Mean agent gain retains the same form: $\bar{G}^\$(t) = \frac{1}{N} \sum_{i=1}^{N} g_i^\$(t)$. We anticipate that $\bar{G}^\$(t) > 0$ 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 reference [8].

3 Quantitative statement and tests

3.1 Analytic calculation versus numeric simulation

We use the Markov chain formalism for the THMG [7] and extend it to both a THMAJG and a TH$G to obtain theoretical prediction for agent and strategy gains, ΔW_{Agent} and $\Delta W_{Strategy}$ [10]:

$$\langle \Delta W_{agent}^{game} \rangle = \pm \frac{1}{N} |\vec{A_D}| \cdot \vec{\mu} \tag{1}$$

$$\langle \Delta W_{strategy}^{game} \rangle = \frac{1}{2N} \left(\hat{\mathbf{s}}_\mu \cdot \vec{\kappa} \right) \cdot \vec{\mu}. \tag{2}$$

In equations (1) and (2), brackets denote a time average. The superscript "game" identifies the game type with game $\in \{M, MAJ, \$\}$. In equation (1), the minus sign is needed for the MG, otherwise not; $|\vec{A_D}|$ is a $(m+\tau)$-bit-length vector of absolute values $|A_D(\mu)|$, one for each μ. μ is a $(m+\tau)$-bit "path history" [10] (sequence of 1-bit states); $\vec{\mu}$ is the normalized steady-state probability vector for the history-dependent $(m+\tau)$ $(m+\tau)$ transition matrix $\hat{\mathbf{T}}$, where a given element $T_{\mu_t,\mu_{t-1}}$ represents the transition probability that μ_{t-1} will be followed by μ_t; $\vec{A_D}$ is a $2^{(m+\tau)}$-element vector listing the particular sum of decided values of $A(t)$ associated with each path-history; $\hat{\mathbf{s}}_\mu$

Table 1. Numeric and analytic results for a typical quenched initial disorder state in the THMG, -MAJG and -\$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	$+0.40$	$+0.06$

is the table of points accumulated by each strategy for each path-history; $\overrightarrow{\kappa}$ 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 reference [2], $\hat{\mathbf{T}}$ may be derived from $\overrightarrow{A_D}$, $\hat{\mathbf{s}}_\mu$ and $\overrightarrow{N_U}$, 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.

Agreement is excellent between numerical simulations and the analytical predictions (1) and (2) for the THMG, THMAJG and TH\$G. For instance, for $\{m, S, N, \tau\} = \{2, 2, 31, 1\}$, for one identical quenched disorder state, 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.

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 $\tau \ll \tau_{eq}$ and $\tau < \tau_{eq}$ – 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.

3.2 Illusory versus genuine control for $\tau \ll \tau_{eq}$ and for $\tau < \tau_{eq}$

Almost all results that hold for multiple values of $\tau \ll \tau_{eq}$ are illustrated for $1 \leqslant \tau \leqslant 1000$. In Figure 1 through Figure 6 we present graphic representations of the ensemble average of 50 runs comparable to Table 1 but over many values of m, for all three games, at $\tau = 1$ and then at $\tau = 1000$. At 1000, τ is close to equilibrium for small $m < m_c$, and as long or longer than real-world trading "lookbacks" for all but high-frequency data (e.g., for $m = 2$, $\tau_{eq} \approx 800$).
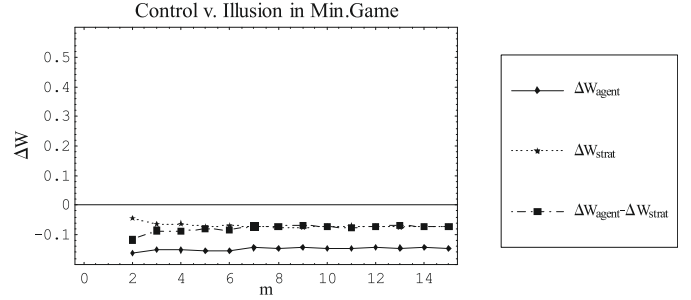

Fig. 1. Agent versus Strategy mean per-step gain in the THMG at various m with $\tau = 1$. The phase transition expected at $m = 4$ is absent with respect to the metrics presented here; strategies outperform agents at all m as indicated by the black squares: agent performance is always negative relative to strategy performance. The optimization procedure employed by agents yields worse performance than their component strategies on the basis of which agents select wich strategies to deploy at each time-step.
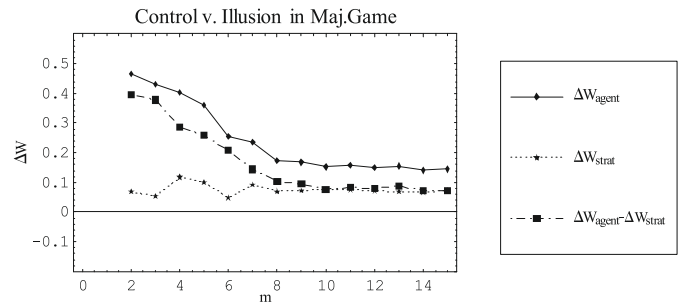

Fig. 2. Agent versus strategy mean per-step gain in the THMAJG at various m with $\tau = 1$. Agent performance is always positive and greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies.

We see that the illusion of control in the THMG persists at all values of m and whether $\tau = 1$ or $\tau = 1000$. Simulations at various intermediate values of τ confirm a smooth transition from Figures 1–3 to Figures 4–6 respectively. Incidentally we note that the phase transition to be expected at $m = 4$ is strongly suppressed for $\tau = 1$ 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, 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
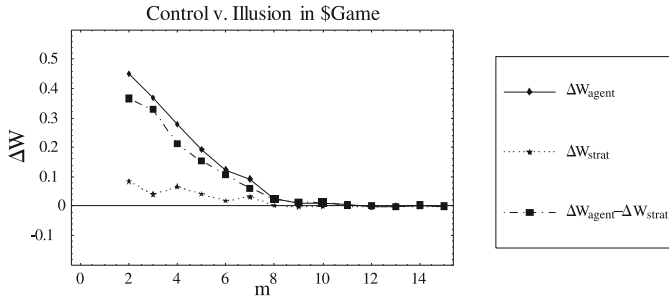
Fig. 3. Agent versus strategy mean per-step gain in the TH$G at various m with $\tau = 1$. Agent performance is always greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies, but the gain becomes asymptotically small for large m.
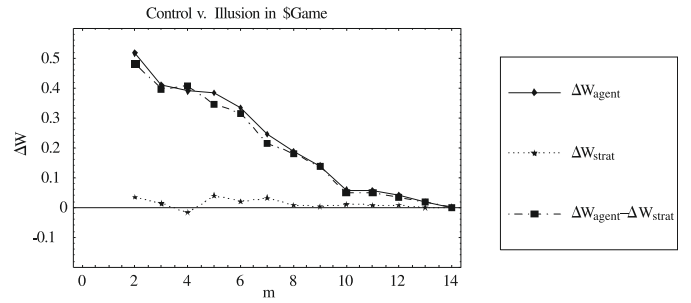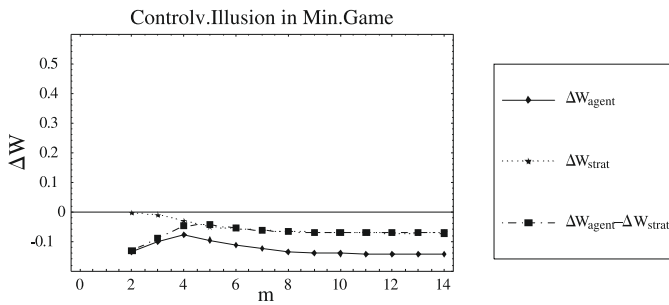

Fig. 4. Agent versus strategy mean gain per-step in the THMG at various m with $\tau = 1000$. The phase transition expected at $m = 4$ is clearly visible. Even with a very long lookback of historical data, the optimization procedure employed by agents yields worse performance than their component strategies.
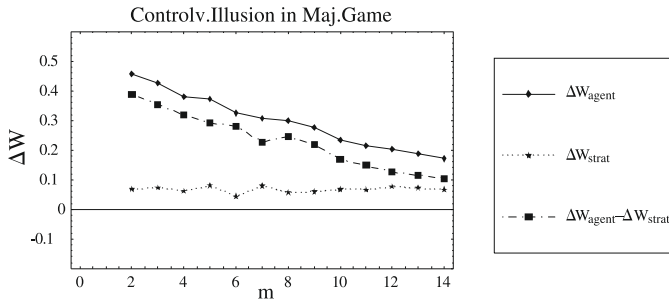

Fig. 5. Agent versus strategy mean gain per-step in the THMAJG at various m with $\tau = 1000$. Agent performance is always positive and greater than strategy performance. The optimization procedure employed by agents yields better performance than their component strategies.

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 τ 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 [12,13].

Fig. 6. 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.

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 [8].

In the MG proper, where τ grows without bound and agent and strategy performance begins to be measured only after τ_{eq} steps, agent performance will exceed strategy performance – optimization succeeds – but only for $m \geqslant m_c$, as reported in the supplementary material for [2]. 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). In the MAJG and $G, equilibrium is reached much sooner and so there is no significant departures from the results shown above for $\tau = 1000$.

4 Interpretations: crowding-out, anti-optimizing agents and persistence

4.1 Illusion of control and the crowding-out mechanism

Illusion-of-control effects in the THMG result from the fact that there is a larger probability for a strategy that has performed well in the past to be chosen by an increasing number of agents, which inevitably leads to its failing, i.e., it becomes crowded out due to the minority rule for winning [2]. Optimizing agents tend on average to adapt to the past but not the present. They choose an action $a(t)$ that is on average out-of-phase with the collective action $A(t)$ or $D(t)$. In contrast, non-optimizing (random- or fixed-choice) agents average over all regimes for which their strategy may be good or bad, and do not face the crowding-out effect. The crowding-out effect also explains simply why anti-optimizing agents over-perform [2]: choosing their worst strategy ensures that it will likely be less used by other agents in the next time step, which makes it more probable that they will be in the minority. In reference [2], we consider generalizations of the MG in which more clever agents use this insight to develop strategies based on their awareness of prior-level effects: when agents

use a boundless recursion scheme to learn and optimize their strategy, the game converges to an equilibrium with fully symmetric mixed strategies where agents randomize their choice at each time step with unbiased coin tosses. 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 to [2] for further details.)

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-delayed factor: on their own, strategies show effectively neither gain nor loss. Gains 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.

4.2 Illusion of control and "anti-optimizing" agents

We 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 τ, rather than the strategy with the most points as do s-agents. In [2], we studied this phenomenon in the "crowded regime" ($m < m_c, \alpha < \alpha_c$) where crowd and anti-crowd formation is likely and the minority mechanism causes the larger-sized crowds to lose. In this regime, c-agents that choose their worst performing strategy consistently outperform s-agents that choose their best. Here 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 $\tau = 100$. τ is long enough so that the phase transition in the MG is not suppressed at $m_c(=4$ for $N = 31)$.

Figure 7 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 τ, agents underperform strategies, we see that the opposite is true for c-agents: in
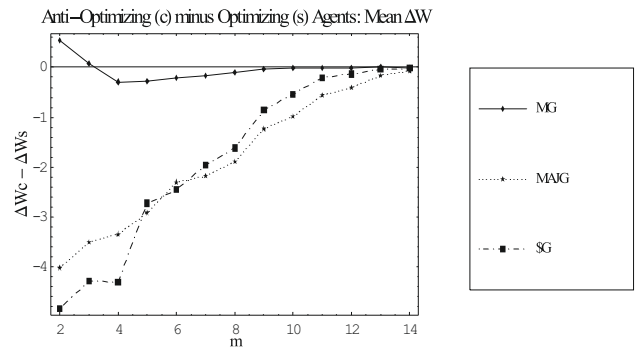

Fig. 7. Difference between c-agent and s-agent mean per-step change in wealth for 3 of 31 c-agents, averaged over 100 days and 200 runs with $\tau = 400$ in the THMG, THMAJG and TH$G.

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.

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 Figures 5 and 6.

4.3 Persistence versus anti-persistence in the THMG, THMAJG and TH$G

As discussed in [2], in the MG and THMG, the degree to which agents underperform their own strategies varies with the phase as parameterized by α. As noted in [7], in the crowded phase ($\alpha < \alpha_c, m < m_c$), the "crowd" of such 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, when a strategy is used, it is probably used by more than one agent, often by many agents. If used by 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 $\overrightarrow{A_D}$ therefore does not show trends (or persistence), but rather anti-persistence.

Anti-persistence is not equivalent to "random" and is scale-dependent. Consider a binary time-series with an m-bit $\mu(t)$ defined in the same way as we have in the MG or THMG: $\mu(t)$ is a sliding window of 1-bit states each of length m: $s(t - m + 1), \ldots s(t)$. A *perfectly* anti-persistent binary series at scale $m = 2$, for example, is
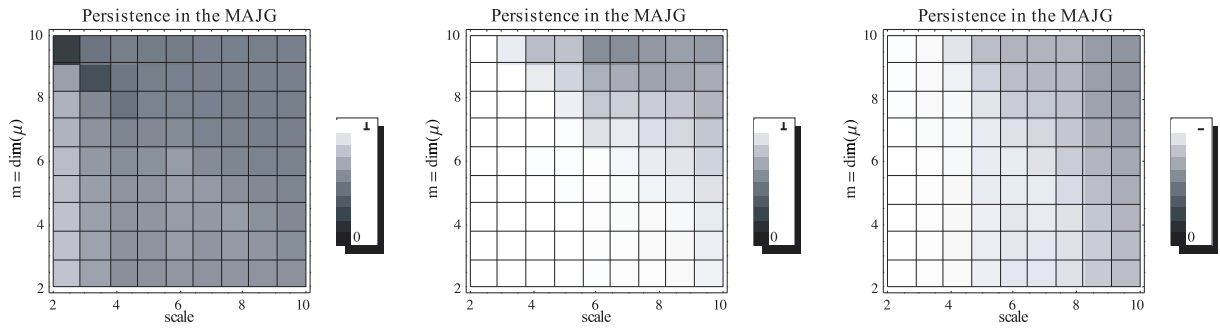
Fig. 8. Persistence (white)/Anti-Persistence (black) at various scales and memory lengths in the MG, MAJG and \$G. The grey scale between 0 and 1 given to the right of the checkerboards encodes the degree of persistence at the chosen scale (abscissa) for different m values (ordinate), calculated as described in the text, using game-generated binary histories of length 1000 over 100 different runs for each game type.

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. 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 [14]. These tables are identical to strategy pairs at the maximum Hamming distance in the MG. No matter which of the $2^{m+1} = 8$ possible strategies is used for the first table, and regardless of which of the $2^m = 4$ possible $\mu(t)$ are used to initiate it, the time series generated by these tables will rapidly settle into perfect anti-persistence.

The "persistence" $\mathcal{P}$ of a given series at scale (defined as) m_s is thus simply the proportion of times that histories of length $m_s + 1$ end in bits 00 or 11; anti-persistence (i.e., $1 - \mathcal{P}$) is the proportion they end in 01 or 10 (One must examine strings of length $m_s + 1$ to determine the persistence as scale m_s).

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 [15,16] for an illuminating study on how to determine the optimal coarse-grained scale in simple cellular automata.

However, for m or τ 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 8 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 $m \in \{2, 3, ..., 10\}$, scale $\in \{2, 3, ...10\}$. 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

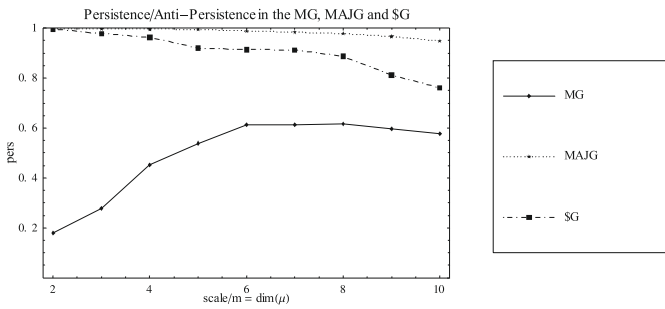Persistence/Anti−Persistence in the MG, MAJG and $G

Fig. 9. Persistence/anti-persistence at various scales = memory lengths in the MG, MAJG and $G.

persistence and anti-persistence of a random sequence), those in the THMAJG and TH$G with persistence ($>0.5$).

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 9 charts the degree of persistence along this diagonal for all three games.

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.

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 reference [2], 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). We see in the preceding analysis of persistence, 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.

The fact that the illusion of control occurs in at least one kind of market model under a wide range of conditions is consistent with the tacitly adopted "rule" implied by the fact that sophisticated academic investors have learned *not* to employ widely-supported optimization

theories. I.e., Doran and Wright 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 *un*important in the decision of whether to buy or sell a specific stock (in particular, the CAPM, APT and Fama and French and Carhart models) [9]. Our analyses of genuine and illusory control may be considered a small step toward formalizing the conditions under which optimization is likely to succeed and when not, conditions which many sophisticated investors may learn to detect from sheer experience and without necessarily being able to reduce to quantitative form.

In another paper [3], we extend the preceding 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, identifying the circumstances under which such predictors succeed and when they fail.

Appendix A: Analytic methods for the THMG, THMAJG and TH$G for choosing the best strategy

To emphasize the relation of the THMG, THMAJG and TH$G to market-games and to either genuine or illusory optimization, we have transformed the fundamental results of the games from statements on the properties of δ^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 use the simplest possible formulation – if an agent's actual (or strategy's hypothetical) vote places it in the winning group (minority in the THMG, majority in the THMAJG and TH$G), 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:

$$a_i(t) \in \{1, 0\}, \forall i, t. \tag{A.1}$$

In the simplest version of the MG, MAJG and $G with N agents, every agent has $S = 2$ strategies and $m = 2$. In the THMG, THMAJG and TH$G, 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 τ (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 $\tau \geqslant t_{st}$. The

present formulation follows closely that presented in Appendix (A.1) of [2], applying it simultaneously to all three games, the THMG, THMAJG and TH\$G.

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:

$$A(t) = 2\sum_{i=1}^{N} a_i(t) - N = N_1 - N_0 \qquad (A.2)$$

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 $(N_1 < N_0)$. In the MG the minority choice is the "winning" decision for t. In both the MAJG and \$G the majority choice is the "winning" decision for t. This choice is mapped back to $\{0,1\}$:

$$D_{sys}(t) = -\text{Sgn}[A(t)] \therefore D_{sys}(t) \in \{-1,+1\} \to \{0,1\}. \qquad (A.3)$$

For the non-TH version all three games, binary strings of length m form histories $\mu(t)$, with $m = \dim[\mu(t)]$. For the TH games, 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 μ_t as a path in the THMG. Note that for memory m, there are 2^{2^m} possible strategies from which agents select S at random. However as first detailed in reference [11], 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 reference [10], we may represent this quenched disorder in the allocation of strategies among agents (from the RSS) by a $\dim = \prod_{s=1}^{S} 2^{m+1}$ tensor, $\hat{\Omega}$ (or from the full strategy space by a $\dim = \prod_{s=1}^{S} 2^{2^m}$ 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 [10], 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 8×8 matrix $\hat{\Omega}$ and its symmetrized equivalent $\hat{\Psi} = \frac{1}{2}\left(\hat{\Omega} + \hat{\Omega}^T\right)$. Positions along all S edges of $\hat{\Omega}$ represent an ordered listing of all available strategies. The numerical values $\Omega_{ij\ldots}$ in $\hat{\Omega}$ indicate the number of times a specific strategy-tuple has been selected (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 strategies in an agent has no meaning. The example (A.4) is a typical such

tensor $\hat{\Omega}$ for $S = 2$, $N = 31$.

$$\hat{\Omega} = \begin{pmatrix} 1 & 2 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 3 & 3 & 1 & 1 \\ 0 & 0 & 2 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 2 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \qquad (A.4)$$

Actions are drawn from a reduced strategy space (RSS) of dimension $r = 2^{m+1}$. Each action is associated with a strategy k and a history $\mu(t)$. There are 2^m histories. Together, all actions by history and strategy (converted for convenience from $\{0,1\}$) can be represented in table form as a $\dim(\text{RSS}) \times 2^m$ binary matrix i.e., $a_k^{\mu_t} \in \{-1,+1\}$. In this case, the table reads:

$$\hat{a} \equiv \begin{pmatrix} -1 & -1 & -1 & -1 \\ -1 & -1 & +1 & +1 \\ -1 & +1 & -1 & +1 \\ -1 & +1 & +1 & -1 \\ +1 & -1 & -1 & +1 \\ +1 & -1 & +1 & -1 \\ +1 & +1 & -1 & -1 \\ +1 & +1 & +1 & +1 \end{pmatrix}. \qquad (A.5)$$

The change in wealth (point gain or loss) associated with each of the $r = 8$ strategies for the 8 allowed transitions among the 4 histories) at any time t for each of the three games is then:

$$\delta \overrightarrow{S}_{\mu(t),\mu(t-1)}^{\min} = +\left(\hat{a}^T\right)_{\mu(t)} \times \{2\text{Mod}\,[\mu(t-1),2]-1\} \qquad (A.6)$$

$$\delta \overrightarrow{S}_{\mu(t),\mu(t-1)}^{maj} = -\left(\hat{a}^T\right)_{\mu(t)} \times \{2\text{Mod}\,[\mu(t-1),2]-1\} \qquad (A.7)$$

$$\delta \overrightarrow{S}_{\mu(t),\mu(t-1)}^{\$} = -\left(\hat{a}^T\right)_{\mu(t-1)} \times \{2\text{Mod}\,[\mu(t-1),2]-1\}. \qquad (A.8)$$

$\text{Mod}[x,y]$ 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 $\{1,2,3,4\}$. Equations (A.6), (A.7) and (A.8) pick out from (A.5) 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 along all the τ histories (i.e., along the path μ_t, accounting for the last τ time steps used to score the strategies) is:

$$\overrightarrow{s}_{\mu_t}^{game} = \sum_{i=0}^{\tau-1} \delta \overrightarrow{S}_{\mu(t-i),\mu(t-i-1)}^{game} \qquad (A.9)$$

(A.9) accounts for the change in points along path μ_t by summing them over all transitions on the path, with game $\in \{\min, \text{maj}, \$\}$ for example, for $m = 2$ and $\tau = 1$,

$$A_D^{min}\left(\mu_t\right) = \left(\sum_{r=1}^{8}\left\{\left[\left(1-\mathrm{Sgn}\left[\overrightarrow{s}_{\mu_t}^{\,min}\ominus\overrightarrow{s}_{\mu_t}^{\,min}\right]\right)\circ\hat{\boldsymbol{\Psi}}\right]\bullet\hat{a}\right\}_r\right)_{\left(\mathrm{Mod}[\mu_t-1,2^m]+1\right)} \tag{A.16}$$

$$A_D^{maj}\left(\mu_t\right) = \left(\sum_{r=1}^{8}\left\{\left[\left(1-\mathrm{Sgn}\left[\overrightarrow{s}_{\mu_t}^{\,maj}\ominus\overrightarrow{s}_{\mu_t}^{\,maj}\right]\right)\circ\hat{\boldsymbol{\Psi}}\right]\bullet(-\hat{a})\right\}_r\right)_{\left(\mathrm{Mod}[\mu_t-1,2^m]+1\right)} \tag{A.17}$$

$$A_D^{\$}\left(\mu_t\right) = \left(\sum_{r=1}^{8}\left\{\left[\left(1-\mathrm{Sgn}\left[\overrightarrow{s}_{\mu_t}^{\,\$}\ominus\overrightarrow{s}_{\mu_t}^{\,\$}\right]\right)\circ\hat{\boldsymbol{\Psi}}\right]\bullet(-\hat{a})\right\}_r\right)_{\left(\mathrm{Mod}[\mu_{t-1}-1,2^m]+1\right)}. \tag{A.18}$$

the strategy scores are kept for only a single time-step and $\tau-1=0$ so the sum vanishes. (A.9) in matrix form therefore reduces to the score:

$$\overrightarrow{s}_{\mu_t}^{\,game} = \delta\overrightarrow{S}_{\mu(t),\mu(t-1)}^{\,game} \tag{A.10}$$

or, listing the results for all 8 path histories:

$$\hat{\mathbf{s}}_\mu^{game} = \delta\hat{\mathbf{S}}^{game} \tag{A.11}$$

$\delta\hat{\mathbf{S}}^{game}$ is an 8×8 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 μ_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:

$$A^{game}(\mu_t) = A_D^{game}(\mu_t) + A_U^{game}(\mu_t). \tag{A.12}$$

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}. This number can be explicated (using an extension to the method employed in [10] for the THMG) as:

$$N_U^{min}\left(\mu_t\right) =$$
$$\left\{\left(1-\left[\left(\hat{a}^T\right)_{(\mathrm{Mod}[\mu_t-1,4]+1)}\otimes_\delta\left(\hat{a}^T\right)_{(\mathrm{Mod}[\mu_t-1,4]+1)}\right]\right)\right.$$
$$\left.\circ\left(\overrightarrow{s}_{\mu_t}^{\,min}\otimes_\delta\overrightarrow{s}_{\mu_t}^{\,min}\right)\circ\hat{\boldsymbol{\Omega}}\right\}_{(\mathrm{Mod}[\mu_t-1,2^m]+1)} \tag{A.13}$$

$$N_U^{maj}\left(\mu_t\right) =$$
$$\left\{\left(1-\left[\left(\hat{a}^T\right)_{(\mathrm{Mod}[\mu_t-1,4]+1)}\otimes_\delta\left(\hat{a}^T\right)_{(\mathrm{Mod}[\mu_t-1,4]+1)}\right]\right)\right.$$
$$\left.\circ\left(\overrightarrow{s}_{\mu_t}^{\,maj}\otimes_\delta\overrightarrow{s}_{\mu_t}^{\,maj}\right)\circ\hat{\boldsymbol{\Omega}}\right\}_{(\mathrm{Mod}[\mu_t-1,2^m]+1)} \tag{A.14}$$

$$N_U^{\$}\left(\mu_t\right) =$$
$$\left\{\left(1-\left[\left(\hat{a}^T\right)_{(\mathrm{Mod}[\mu_{t-1}-1,4]+1)}\otimes_\delta\left(\hat{a}^T\right)_{(\mathrm{Mod}[\mu_{t-1}-1,4]+1)}\right]\right)\right.$$
$$\left.\circ\left(\overrightarrow{s}_{\mu_t}^{\,\$}\otimes_\delta\overrightarrow{s}_{\mu_t}^{\,\$}\right)\circ\hat{\boldsymbol{\Omega}}\right\}_{(\mathrm{Mod}[\mu_{t-1}-1,2^m]+1)} \tag{A.15}$$

"$\otimes_\delta$" is a generalized outer product, with the product being the Kronecker delta. $\overrightarrow{N}_U$ constitutes a vector of such values. The summed value of all undetermined decisions for a given μ_t is distributed binomially. Note that (A.13) and (A.14) are structurally identical while (A.15) differs from these in that the indices on $\left(\hat{a}^T\right)$ and on the entire expression reference path-histories μ_{t-1} rather than μ_t, reflecting the one-step time-lag in the payoff for the \$G. Similarly:

see equations (A.16–A.18) above

Details of the derivation as applying to the THMG may also be found in reference [17]. We define $\overrightarrow{A}_D$ as a vector of the determined contributions to $A(t)$ for each path μ_t. In expressions (A.11) and (A.12) μ_t numbers paths from 1 to 8 and is therefore here an index. $\overrightarrow{s}_{\mu_t}$ is the "μ_t^{th}" vector of net point gains or losses for each strategy when at t the system has traversed the path μ_t (i.e., it is the "μ_t^{th}" element of the matrix $\hat{\mathbf{s}}_\mu = \delta\hat{\mathbf{S}}$ in (A.11)). "$\ominus$" is a generalized outer product of two vectors with subtraction as the product. The two vectors in this instance are the same, i.e., $\overrightarrow{s}_{\mu_t}$. "$\circ$" is Hadamard (element-by-element) multiplication and "$\bullet$" the standard inner product. The index r refers to strategies in the RSS. Summation over r transforms the base-ten code for μ_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 (A.13). This expression yields an index number, i.e., selection takes place $1 + $ Modulo 4 with respect to the value of (μ_t-1) for the THMG and with respect to the value of $(\mu_{t-1}-1)$ for the THMAJG and TH\$G.

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 $\Gamma_{\mu_t,\mu_{t-1}}$

$$\hat{\boldsymbol{\Gamma}} = \begin{pmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{A.19}$$

Equations (A.13), (A.14) and (A.15) yield the history-dependent $(m+\tau)$ by $(m+\tau)$ 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 μ_t and

μ_{t-1} (the game-type superscripts on A_D and N_U are understood in context):

$$T_{\mu_t,\mu_{t-1}}^{game} = \Gamma_{\mu_t,\mu_{t-1}}^{game} \sum_{x=0}^{N_U^{game}(\mu_t)} \left\{ \begin{matrix} N_U^{game}(\mu_t) \\ \mathbb{C} \\ x \end{matrix} \left(\frac{1}{2}\right)^{N_U^{game}(\mu_t)} \right.$$

$$\times \, \delta \left[\mathrm{Sgn}(A_D^{game}(\mu_t) + 2x - N_U^{game}(\mu_t)) \right.$$

$$\left. \left. + \, (2\mathrm{Mod}\{\mu_{t-1}, 2\} - 1) \right] \right\}. \quad (A.20)$$

The expression $\begin{matrix} N_U^{game}(\mu_t) \\ \mathbb{C} \\ x \end{matrix} \left(\frac{1}{2}\right)^{N_U^{game}(\mu_t)}$ in (A.20) represents the binomial distribution of undetermined outcomes under a fair coin-toss with mean $= A_D^{game}(\mu_t)$. Given a specific $\hat{\mathbf{\Omega}}$,

$$\langle A^{game}(\mu_t)\rangle = A_D^{game}(\mu_t) \forall \mu_t. \quad (A.21)$$

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), a value which is the same across all games (hence we drop the "game" superscript):

$$\vec{\kappa} \equiv \sum_{k=1}^{2^{m+\tau}} \left(\hat{\mathbf{\Omega}} + \mathbf{\Omega}^T\right)_k = 2\sum_{k=1}^{2^{m+\tau}} \hat{\mathbf{\Psi}}_k$$

$$= \{n(\sigma_1), n(\sigma_2), \ldots n(\sigma_{2^{m+\tau}})\} \quad (A.22)$$

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

$$\langle \Delta W_{Agent}\rangle = \pm\frac{1}{N} Abs\left(\vec{A_D}\right) \cdot \vec{\mu} \quad (A.23)$$

(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

$$\langle \Delta W_{Strategy}\rangle = \frac{1}{2N}\left(\hat{\mathbf{s}}_\mu \cdot \vec{\kappa}\right) \cdot \vec{\mu} \quad (A.24)$$

with $\vec{\mu}$ the normalized steady-state probability vector for $\hat{\mathbf{T}}$. Expression (A.23) 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, with a minus sign for the THMG. Expression (A.24) 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 $\hat{\mathbf{\Omega}}$) times the sum over the per-step wealth change associated with every history.

References

1. E. Langer, Journal of Personality and Social Psychology **7**, 185 (1975)
2. J. Satinover, D. Sornette, Eur. Phys. J. B **60**, 369 (2007)
3. J.B. Satinover, D. Sornette, *Cycles, determinism and persistence in agent-based games and financial time-series*, submitted (2008), preprint available at arXiv.org/abs/0805.0428
4. W.B. Arthur, *Out-of-Equilibrium Economics and Agent-Based Modeling, in the Handbook of Computational Economics, Agent-Based Computational Economics*, edited by K. Judd, L. Tesfatsion (Elsevier, North-Holland, 2005), Vol. 2
5. D. Challet, M. Marsili, Y.-C. Zhang, *Minority Games: Interacting Agents in Financial Markets*, Oxford Finance (Oxford University Press USA, 2005)
6. A.C.C. Coolen, *The Mathematical Theory of Minority Games: Statistical Mechanics of Interacting Agents*, Oxford Finance (Oxford University Press, USA, 2005)
7. M. Hart, P. Jefferies, P.M. Hui, N.F. Johnson, Eur. Phys. J. B **20**, 547 (2001)
8. J.V. Andersen, D. Sornette, Eur. Phys. J. B **31**, 141 (2003)
9. J.S. Doran, C. Wright, *What really matters when buying and selling stocks*? Working paper of Florida State University (2007) (http://ssrn.com/abstract=980291)
10. M.L. Hart, P. Jefferies, N.F. Johnson, Physica A **311**, 275 (2002)
11. D. Challet, Y.C. Zhang, Physica A **256**, 514 (1998)
12. E.F. Fama, J. Finan. **25**, 383 (1970)
13. E.F. Fama, J. Finan. **46**, 1575 (1991)
14. R. Metzler, J. Phys. A: Math. Gen. **35**, 721 (2002)
15. N. Israeli, N. Goldenfeld, Phys. Rev. Lett. **92**, 074105 (2004)
16. N. Israeli, N. Goldenfeld, Phys. Rev. E **73**, 026203 (2006)
17. M.L. Hart, P. Jefferies, N.F. Johnson, Physica A **311**, 275 (2002)