

A Discrete Model Gauging Market Efficiency

John Mount*

September 8, 2009

Abstract

We describe a discrete market model appropriate for quantifying certain desirable and un-desirable features of financial markets. This model allows direct exploration of the impact of different market structures on efficiency and fairness. We conclude by demonstrating that a single trader with a large budget can generate profit while making the market not profitable for smaller traders.

1 Introduction

Stochastic calculus techniques[KS01] (such as Brownian Motion, Levy Processes[App04], Wiener Processes or the Ito Calculus[Ste03b, Ste03a]) are not the only abstraction useful in thinking about financial markets. Real markets do not meet the typical assumptions of the above systems (infinitely divisible time, no trade costs, no long-term memory and no large actors) and routinely fail goodness of fit tests against such models[LM01, Lo05]. In fact there is a simple arbitrage argument that markets would have summary statistics identical to Ito processes even if they are not such processes.[Sha04] When studying which features make a market fair or efficient we can not rely on mathematical tools that assume and depend on fair and efficient markets.

To build the tools for our study we follow up on some of the ideas of Hasanhodzic, Lo and Viola [HLV09] and propose a specific discrete market model (as distinguished from more traditional continuous mathematics as in [Mer99]) that allows us to effectively apply ideas from game theory[NNV07] and theoretical computer science. We show how to solve for optimal trading strategies in this market model and conclude with an illustration of how a single trader can dominate a market by merely exercising a larger budget.

2 Outline

We will proceed as follows:

- Define our market model
- Solve for optimal trading strategies in our market model
- Perform the experiment of adding a single large trader to our model
- Draw conclusions
- Suggest further research.

3 The Market Model

Our goal is to investigate if even perfect traders are vulnerable to an additional trader that has a larger budget. To do this we must have a market model where at least:

- We can solve for the optimal trading strategy
- There is a reason to trade (profits are available).

We propose such a market model below.

*email: <mailto:jmount@win-vector.com> company: <http://www.win-vector.com/>

3.1 The Market

To simplify the description of traders (and to minimize the amount of state we have to carry) we propose a market model that abstracts out price and many other features.

Our market model is represented as an ordered sequence of the symbols “+”, “0” and “-”. A “+” represents a recent price increase, a “0” represents no change and a “-” represents a recent price decrease. We are deliberately avoiding direct representation of real market quantities such as absolute price, volume, inventory, bid/ask books, margin and elasticity. Time is represented by regular “ticks” or the simple advance to the next symbol in the market sequence. We will describe how the next symbol in the market sequence is determined after we have described trades.

3.1.1 Type 1 Trades

The first type of trade we allow in this market is a “round trip.” A round trip is one of the two following trades:

- “a long round trip”
An immediate buy in the current time tick followed by an automatic (forced) sell on the next time tick. This trade is considered profitable if the next market symbol is a + as the sell then happens at a higher price than the initial buy, yielding a profit.
- “a short round trip”
An immediate sell in the current time tick followed by an automatic (forced) buy on the next time tick. This trade is considered profitable if the next market symbol is a - as the buy then happens at a lower price than the initial sell, yielding a profit.

The forced nature of these round trip trades allow us to avoid modeling inventory and margin. Round trip trades are meant to abstract some of the aspects of high-frequency trading strategies.

3.1.2 Type 2 Trades

The second type of trade we allow is a “simple buy” or “simple sell” on the next time tick. This type of trade is meant to abstract some of the properties of a trader who is not so close to the market and has market-external interests (like inventory, customers, margin, fundamental knowledge ...).

3.1.3 Market Evolution

The market model evolves forward as follows. The second half of each type 1 trade (the sell in the long round trip and buy in the short round trip) is entered as a net impact on the upcoming time tick. So: a long round trip actually generates a sell or downward price impact on the next market tick (and a short round trip generates a buy or upward price impact on the next market tick). This “reverse impact” is in our model because we are not allowing these traders to hold inventory and in a “buy followed by a sell” pattern the initial buy impact is further in the past than the sell (so should have a lesser future impact). This is also similar to how in real markets a large net short position represents an upward influence on price as the market participants know the short position must eventually be covered.

Also each type 2 (or simple) trade is also entered directly as market impact. So: as expected simple buy trades generate upward price impact and simple sell generate downward price impact.

To determine the next market-symbol we sum the net impact entered against the next tick, if the net impact is positive the symbol is a +, if it is zero the symbol is 0 and if it is negative the symbol is a -. This differs both from the market model in [HLV09] (where price is additive) and from real markets (where elasticity of price with respect to trades is very complicated).

For example: if three traders choose “long round trip” (betting the market will go up in the short term) and one trader chooses “simple buy” (betting the market will go up long term) then the net impact on the next tick is $(-1) + (-1) + (-1) + (+1) = -2$ and the next symbol is -. The long round trip traders lose money and the simple buy trader is has an unrealized loss.¹ Just as we settled on a standard unit for trade size we will use a standard unit for profit and arbitrarily say all traders with realized loss lost one unit per share.

This market model is deliberately simple, but just as symbolic dynamics offers insights to continuous dynamical systems [TBS91] this market model serves as a platform for analyzing aspects of real markets.

¹We do not enforce any sort of “conservation of money” (that the amount of profit earned by the short trader should equal the amount of money lost by the long traders). In the real market there is an aspect of conservation of money in trades, but there is not a conservation of money in a single time period if the traders have net holdings.

3.2 Type 1 Traders

We have described a very simple and very limited market. We will now describe some of the traders. Our first set of traders we call “Type 1 Traders” and they are meant to represent high-frequency quantitative or technical traders. Type 1 traders perform only type 1 trades (long round trip or short round trip) or abstain from trading. For now we are restricting each type 1 trader to trade a single unit either in a long round trip, a short round trip, or to not trade.

We will model these traders as having no internal state and a limited window of memory of the market. We allow the traders to use probabilistic strategies (so they do not get caught always performing the exact same trade in a repeating situation). Under these limits we can write each trader as a simple table representing a map from $\{+, 0, -\}^k$ (the sequences of symbols length k , i.e. what the trader is modeled as remembering) to pairs $(p_{\text{long}}, p_{\text{short}})$ where p_{long} is the trader’s chosen probability of making a long round trip in this situation and p_{short} is the trader’s chosen probability of making a short round in this situation.²

We place no limit on how much effort the Type 1 Traders make in pre-computing their strategy tables. One important point is: since the traders are allowed to use probabilistic tables we can assume (in the limit) that the optimal trading strategy is the same for all type 1 traders. This is because if a type 1 trader is losing money to other type 1 traders who are themselves making a profit then the original type 1 trader can “cannibalize their own business” by copying a bit of the strategy they are vulnerable to into their own strategy. For example if a trader is losing money to profitable short round trippers they can fix this by trading short round trips a bit more often.³ When we can use the assumption that all the type 1 traders have identical strategy tables we can then solve for this table and immediately demonstrate the characteristic of the market formed by these optimal traders.

The market model evolves as follows: if there are m type 1 traders with a common memory window size of k then the market symbol at time- t is:

$$\text{market}(t) = \text{sign} \left(\sum_{i=1}^m \chi_i(\text{market}(t-1), \dots, \text{market}(t-k)) \right)$$

where $\chi_i()$ is the random variable associated with the i -th type 1 trader.

Already we can show: if the market is only populated by type 1 traders then the optimal trading strategy is to set $p_{\text{long}} = p_{\text{short}} = 0$ (to not trade) and there is in fact no market (no trades happen). This follows because for every time no more than half of the active type 1 traders can be on the profitable side, so at best the type 1 traders break even as a group and not trading is a dominant strategy.

To model another important aspect of markets (and to give the type 1 traders a reason to trade) we introduce type 2 traders.

3.3 Type 2 Traders

Type 2 traders are completely oblivious to the market. Type 2 traders trade only type 2 trades (simple buy and simple sell). Type 2 traders trade, but do not look at or remember the market sequence. Oddly enough the type 2 traders abstract both the idea of completely informed traders (traders that know something about the future, so do not need to use the market past) and completely uniformed traders (traders trading due to some external to the market pressure like a need to recover liquid assets). For now we are restricting each type 2 trader to trade a single unit either in a simple buy or a simple sell.

We assume one family of type 2 traders that operate as follows: assume a simple sequence of “+” and “-” generated by the Markov Chain in Figure 1. This Markov Chain emits a sequence of symbols where the same symbol follows the last with probability p (and the symbol changes with probability $1 - p$).

²So $p_{\text{long}} \geq 0$, $p_{\text{short}} \geq 0$ and $p_{\text{long}} + p_{\text{short}} \leq 1$.

³This “traders can imitate each other” is a “linearity of expectation argument” [AS92] and is a common argument technique in game theory.




Figure 1: Simple Markov Chain

We will call this sequence “the hidden symbol” as only our type 2 traders can see it (the type 1 traders can not). Each of our type 2 traders looks at the current hidden symbol and independently does the following: with probability q they enter a simple buy or simple sell for the next time tick betting in the direction of the hidden symbol and with probability $1 - q$ they enter a simple buy or simple sell for the next time tick betting in the direction opposite to the hidden symbol. For now we will assume all type 2 traders share the same q . The type 2 traders do not perform round trip trades, but instead hold inventory. Thus a type 2 trader’s long bet is modeled as adding a net upward impact to the next time period.

The market model now evolves as follows. If there are n type 2 traders then the market symbol at time- t is:

$$\text{market}(t) = \text{sign} \left(\sum_{i=1}^m \chi_i(\text{market}(t-1), \dots, \text{market}(t-k)) + \sum_{i=1}^n \Upsilon_i(\text{hidden}(t-1)) \right)$$

where $\Upsilon_i()$ is the random variable associated with the i -th type 2 trader.

As is often the case in mathematics what the abstract model means can change if we add different interpretations. If the hidden sequence that all of the type 2 traders simultaneously observe is thought to represent some important hidden value like the true value of the company underlying the equity being traded, then we consider the type 2 traders to be informed and consider their knowledge to be an advantage. If we consider the shared sequence to be irrelevant noise then we see these traders as some loose coalition whose value comes only from the fact their trades correlate with each other. If $q = 0.5$ then we have truly uninformed (and uncorrelated) traders who are indeed doing nothing. Many real market properties that are attributed as being consequences of non-arbitrage are in fact consequences of conventions no more meaningful than the one given here (for example: closed end funds).

The interesting point is if p is not too near 0.5 and q is not too near 0.5 then the type 2 traders have a serial correlation (a correlation over time) that the type 1 traders can learn and exploit for profit. Or, from another point of view, the type 1 traders can profit by supplying liquidity to the type 2 traders.

4 Solving for the Optimal Strategy

Our market was designed to allow a very succinct description. With only type 1 traders and one uniform family of type 2 traders our market is completely specified if we know:

- m : The number of type 1 traders in the market
- k : The memory length of type 1 traders
- n : The number of type 2 traders in the market
- p : The symbol stability odds on the hidden sequence watched by type 2 traders
- q : the faithfulness of type 2 traders in trading the hidden symbol.

Given these parameters there is a unique shared optimal strategy for the type 1 traders, and we can efficiently solve for this strategy (without resorting to approximate or simulation results).

The entire state of the market at a given time can be written as a tuple $s = (x, y)$ where x is the sequence of the k most recent result symbols from the market sequence $(+, 0, -)$ and y is the most recent symbol from the hidden sequence $(+, -)$. So there are only $2 * 3^k$ possible states for the market. Any posited type 1 strategy (along with the above parameters) completely determines the transition odds between each of these detailed market states. Figure 2 illustrates the states that make up a $k = 1$ market.

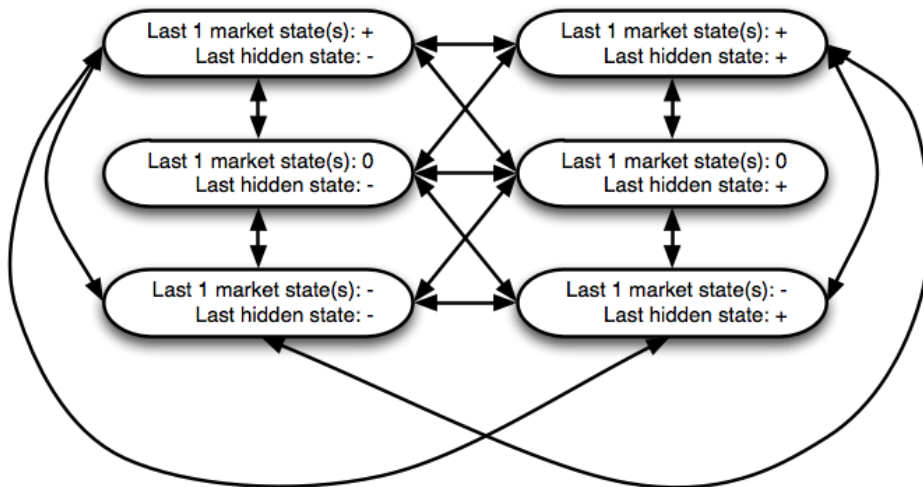


Figure 2: Detailed Market Markov Chain for $k = 1$

Once the transition odds are known between all states it is a simple matter of linear algebra to solve exactly for the stationary distribution and expected value of the market (for type 1 traders).[KS76] Global optimization techniques can be used to identify the optimal strategies and we can then characterize how these market models behave when populated with optimal traders.⁴

For concreteness we show a piece of the computation for the $m = 1, k = 1, n = 2, p = 0.8, q = 0.9$ market model. If the market's last symbol was + and the last hidden state was + then the odds of moving from this state to this same detailed state (both a new hidden symbol of + and a new market symbol of +) for the next time is given by:

$$P(\text{hidden}_{\text{next}} = + | \text{hidden} = +) P(\chi_1(+) + \Upsilon_1(+) + \Upsilon_2(+) > 0)$$

(where $\chi_1()$ is random variable representing the trade of the type 1 trader and $\Upsilon_1(), \Upsilon_2()$ are the random variables representing the trades of the type 2 traders).

Using nothing more complicated than knowledge of the binomial distribution we can compute the complete transition matrix for the detailed Market Markov Chain. For example: assume our type 1 traders trade the most recent market symbol (except 0) with 0.7 probability (and makes no trade otherwise). Now label our states as:

Last Market Symbol	Hidden Symbol	State ID Number
+	+	1
+	-	2
0	+	3
0	-	4
-	+	5
-	-	6

then it is merely a matter of detailed arithmetic to derive the state to state transition probability matrix⁵:

$$P = \begin{pmatrix} 0.648 & 0.162 & 0.648 & 0.162 & 0.7488 & 0.1872 \\ 0.002 & 0.008 & 0.002 & 0.008 & 0.0272 & 0.1088 \\ 0.0432 & 0.0108 & 0.144 & 0.036 & 0.0432 & 0.0108 \\ 0.0108 & 0.0432 & 0.036 & 0.144 & 0.0108 & 0.0432 \\ 0.1088 & 0.0272 & 0.008 & 0.002 & 0.008 & 0.002 \\ 0.1872 & 0.7488 & 0.162 & 0.648 & 0.162 & 0.648 \end{pmatrix}.$$

⁴ The optimization problem has some easy aspects. At the optimum we can assume all the type 1 traders are identical (so we solve for one trader of magnitude n instead of solving for a population of n traders) and we can use automatic differentiation techniques[RC96] to get gradients as we work.

⁵ We are being a little non-standard here in that we are writing P as an operator on the left, so if x is the state-vector of probabilities at a given time tick then Px is the state-vector of probabilities at the next time tick. This is not the convention in the Markov Chain literature, but more compatible with other topics in linear algebra.

Solving for the stationary distribution is, as promised, quite easy. We want to find a vector x such that $(P - I)x = 0$ and $1 \cdot x = 1$ (I denoting the identity matrix). Under very general conditions this will be a set of $s + 1$ equations over s variables with rank s (so will have a unique solution and we don't need to add any sign constraints).

This solution gives us the stationary odds of the market (how likely we are to see the market in any state at a random observation time):

$$x = \begin{pmatrix} 0.420497 \\ 0.048611 \\ 0.030892 \\ 0.030892 \\ 0.048611 \\ 0.420497 \end{pmatrix}.$$

Once we know this it is a matter of arithmetic to determine the expected value of the market for the type 1 trader.⁶ The trading strategy we imposed was not optimal but does have the positive value of 0.13 units expected profit per time tick. We can completely characterize these markets for moderate values of k and arbitrary values of m and n .

Already we can confirm some features we would expect to see in this model. For example the type 1 traders have a “tragedy of the commons” situation in that they are using up the correlations that the type 2 traders introduce. If there are too many technical traders trying to follow the type 2 traders then the market becomes anti-correlated and oscillates in a way that is not profitable for these traders (until they adjust their strategies). For example raising m to 2 in our example makes the “follow the market 70%” of the time an unprofitable strategy that loses money at a rate of 0.12 units per time tick. However, with $m = 2, n = 3$ this same strategy is profitable at a rate of 0.08 units per time tick. The $m = 2, n = 2$ market can be made to be profitable if both of the technical traders act “superrationally”⁷ and lower their trade rate from following the market 70% of the time to something lower like 20% of the time. Figure 3 shows the expected value of the market $m = 2, k = 1, n = 2, p = 0.8, q = 0.9$ for the type 1 traders as the type 1 traders odds of “following the last symbol” are moved from 0 to 1 (and, as earlier, refrain from trading in all other cases).

⁶Some care has to be taken that in computing the value of a strategy as we need access to some several additional transition matrices (each conditioned on knowing the proposed trade of the type 1 trader we are studying).

⁷That is each type 2 trader must dial down their trading activity to account for the number of other type 2 traders present. Douglas Hofstadter called such behavior “superrational” [Hof85]. Traders with small budgets who can not collaborate are actually likely to do this- because while they are trading at too high a rate they lose money. However, a trader that can work at higher volume or tolerate larger losses can outwait the others and have the market for themselves.




Figure 3: Strategy values for $m = 2, k = 1, n = 2, p = 0.8, q = 0.9$

5 The Experiment

Now that we have set up a market and described how to evaluate and solve for the optimal trading strategies we are ready to run an experiment. The experiment is the introduction of a large trader that trades at a much larger size than other type 1 traders. This large trader will act like a type 1 trader but it is allowed larger trade sizes and a small informational advantage over the other type 1 traders. This informational advantage is the ability to remember if their own last trade was one of three possible strategies (so it is not really extending the windows size, and this extension would not help the smaller type 1 traders against this strategy).

To illustrate we assume a market where $m = 0, k = 1, n$ is large, $q > 1/2, p > 3/4$ and $n * (q - 1/2) * (p - 3/4)$ is large (and all known to the large trader).

The large trader trades as follows: define three states to remember the large trader's last state "odd time tick", "even time tick following bluff" and "even time tick following non bluff." We illustrate the large strategy in Figure 4. On odd time ticks the large trader either bluffs (trades to flip the market symbol and takes a forced loss) or trades normally (allows the market to evolve under the influence of the type 2 traders and takes an expected profit). On even time ticks the large trader's behavior depends if the last odd tick was a bluff (and the type 2 traders' influence on the market is masked) or the last odd tick was not a bluff (and the type 2 traders' influence on the market is visible). These two different states are marked in Figure 4 and the large trader abstains from trading after a bluff or trades for expected profit after a non-bluff.

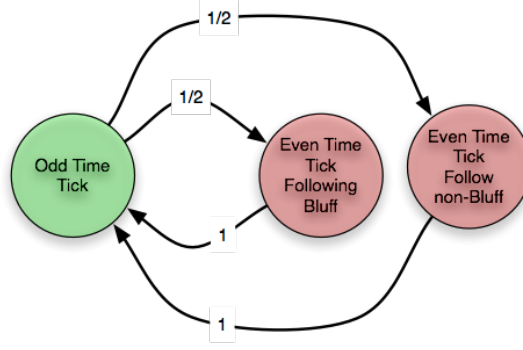


Figure 4: Large Type 1 Trader States

The large trader’s strategy yields an augmented Markov chain that reflects the large trader’s state, the last symbol seen in the market and the last symbol of the hidden sequence. This Markov chain is shown in Figure 5 (with links from even time states to odd time states and links to and from unlikely states suppressed for clarity). We will describe the large trader’s strategy in detail below, but there are some simplifying points to keep in mind. Since $n(q - 1/2)$ is large we are assuming that on odd time ticks and for even time ticks following non-bluffs the states where the market symbol and the hidden symbol disagree are very rare (and we will omit them from the analysis).

Stepping through the large trader strategy (see Figure 5): on the odd time periods the large trader assumes that the market symbol equals the hidden symbol (i.e. the n type 2 traders successfully copied the hidden symbol to the market without interference). The large trader then flips a fair coin and with 50% chance “bluffs” (forcing the market to the symbol opposite the hidden symbol by trading a little more than n units in the appropriate direction) or on the other 50% of the time trades slightly less than n units to try and profit off the obvious tick to tick correlation in the market. On the even time ticks the large trader trades to profit if the previous trade was not a bluff or otherwise abstains from trading. The expected value of the sum of contributions of the type 2 traders is $\text{hidden_symbol} * (q * n - (1 - q) * n)$ which has an absolute value of $(2q - 1)n$. Let $Q = (2q - 1)n$. A bluff costs the large trader $Q + o(n)$ ⁸ units as they enter a trade in large enough to overwhelm the type 2 traders with high probability. A trade for profit (either on a non-bluff odd time tick or a even time tick following a non-bluff) has a maximum expected value of $(Q - o(n)) * (p * (1) + (1 - p) * (-1)) = Q(2p - 1) - o(n)$ as the large trader must not overwhelm the expected effect of the type 2 traders. Every two time ticks the large trader either bluffs then abstains (with probability 1/2) or makes two profitable trade attempts in a row (with probability 1/2). So every 2 time ticks the expected return is:

$$0.5 * (-Q - o(n)) + 0.5 * 2 * (Q(2p - 1) - o(n)) = (Q/2)(p - 3/4) - o(n).$$

Or $(q - 1/2)(p - 3/4)n/2 - o(n)$ expected units return per time tick.

⁸ The $o(n)$ is an “order-of” notation meant to denote a quantity that increases more slowly than n as n gets large. An example $o(n)$ quantity would be \sqrt{n} . This notation (when used properly) greatly speeds up calculation by suppressing irrelevant details.

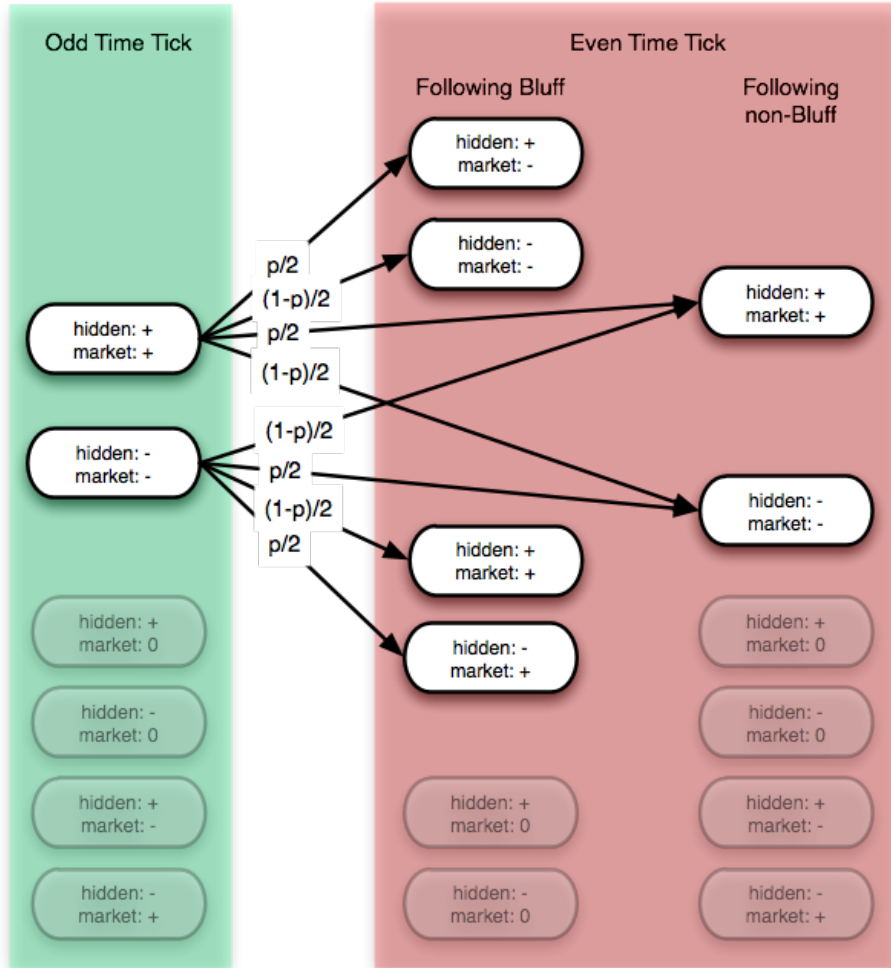


Figure 5: Example Strategy for Large Type 1 Trader

(for clarity transitions between unlikely states and from even time ticks to odd time ticks are not shown)

This large trader strategy is for illustration, and is in no sense optimal⁹. The important result is that when looking at the sequence of market symbols with a window of length 2 (the length of window that would be useful in defining a trading strategy for a $k = 1$ type 1 opposing trader) all the zero free market symbol sequences of length 2 come up with the same probability: $1/4$. To a limited memory type 1 opponent (or one who has to encode their strategy with limited memory) the market looks like a fair coin with no serial correlation. Thus, if we start with $m = 0$ (i.e. no other type 1 traders) a single large trader can take over the market and when we later increase m the new type 1 traders will compute an optimal strategy of not trading (i.e. they will see no method to profitably enter the market).

The large trader has rendered the market untradable for other type 1 traders in the strongest possible sense. Because this market model is symmetric, has no trading costs and no margin requirements, no strategy can exist that forces other adapting strategies to lose money. This is because a strategy that is forced to lose money can be adapted into a profitable strategy by reversing the long and short actions. The large trader is using slightly more memory but this is just an accounting gimmick so they know on which ticks the market has information from the type 2 traders and on which ticks are noise from their own “bluff” or “Pyrrhic” trades. The other type 1 traders have no advantage when given the equivalent gimmick.¹⁰ Also, the large trader strategy is self financing: the large trader can hold the market (make the market look purely random to outsiders) while extracting a profit.

⁹At the very least we could tune the bluff frequency and also trade (albeit with less certainty) in the after-bluff periods

¹⁰Unless they use the gimmick to collude to overcome the organized size of the large trader, but then the other type 2 traders are essentially also one large trader

6 Conclusion

We have described a combinatorial market model that is designed for simplicity. As is well known from mathematics and theoretical computer science even very simple systems can exhibit arbitrarily complex behavior when feedback, recursion or iteration are involved.

We have shown how to explicitly derive optimal trading behavior for small traders in this market model. We then demonstrated how a large trader (allowed to move more volume than the small traders) can “hold the market” in the sense they can make the market appear to be uncorrelated to outsiders while extracting a profit on their own. The ability to completely characterize our market model allows us to show that a self financing large trader is a stable solution in this market model even in the presence of optimal opponents with similar computational power.

It is beyond the scope of current techniques to show under which conditions a self-financing large trader could exist in a “fully realistic” market model. But by demonstration we have shown that we can not assume there are no self financing large traders.

7 Further Research

Interesting follow up studies, which are well within the scope of the methods demonstrated here, include:

- Larger k and heterogeneous k
- A cross-market arbitrage interpretation for the type 2 traders
- More detailed price and hidden symbol trajectories
- Non-finite strategies (strategies indexed by integers instead of a small set of symbols)
- Inventory and margin
- Trade volume controlling price change (i.e. a model of price’s elasticity with respect to trade volume).

References

- [App04] David Applebaum, *Levy processes- from probability to finance and quantum groups*, Notices of the AMS **51** (2004), no. 1336-1347, 12.
- [AS92] Noga Alon and Joel H. Spencer, *The probabilistic method*, Wiley, 1992.
- [HLV09] Jasmina Hasanhodzic, Andrew W Lo, and Emanuele Viola, *A computational view of market efficiency*, 1–14.
- [Hof85] Douglas R. Hofstadter, *Metamagical themas: Questioning for the essence of mind and pattern*, Basic Books Inc., 1985.
- [KS76] John G. Kemeny and J. Lauri Snell, *Finite markov chains*, Springer, 1976.
- [KS01] Ioannis Karatzas and Steven E. Shreve, *Methods of mathematical finance*, Springer, September 2001.
- [LM01] Andrew W Lo and A Craig MacKinlay, *A non-random walk down wall street*, Princeton University Press, 2001.
- [Lo05] Andrew W Lo, *Reconciling efficient markets with behavioral finance: The adaptive markets hypothesis*, 44.
- [Mer99] Robert C. Merton, *Continuous-time finance*, Blackwell, 1999.
- [NNV07] Eva Tardos Noam Nisan, Tim Roughgarden and Vijay V. Vazirani, *Algorithmic game theory*, Cambridge, 2007.
- [RC96] Louis B Rall and George F Corliss, *An introduction to automatic differentiation*, SIAM: Computational Differentiation: Techniques, Applications and Tools (1996), 1–18.
- [Sha04] Glenn Shafer, *Why do price series look like ito processes?*, Rutgers (2004), 43.
- [Ste03a] J Michael Steele, *Ito calculus*, Encyclopedia of Actuarial Sciences (2003), 1–12.
- [Ste03b] J. Michael Steele, *Stochastic calculus and financial applications*, Springer, June 2003.
- [TBS91] Michael Keane Tim Bedford and Caroline Series, *Ergodic theory, symbolic dynamics and hyperbolic spaces*, Oxford University Press, 1991.