

Market Ecologies: The Effect of Information on the Interaction and Profitability of Technical Trading Strategies

Antony Jackson* and Daniel Ladley†

January 26, 2016

Abstract

Technical trading strategies make profits by identifying and exploiting patterns in market prices—patterns generated by the interaction of market participants. Using a model market populated by individuals using a range of trading rules we show that the presence of technical traders may be beneficial, in some cases reducing volatility and increasing price efficiency. In particular, contrarian traders who base their decisions on high frequency data have the largest positive effect. It is also found that if technical traders condition their actions using ‘real time’ information, they partially emulate arbitrageurs and make positive profits.

Keywords: Technical trading rules, artificial market, market ecology

JEL Classification: C63, G12

1 Introduction

There is a long-standing debate surrounding the ability of rational arbitrageurs to drive noise traders out of financial markets. While prominent work has forcibly argued the case for noise traders being driven out of markets e.g., Friedman (1953) and Fama (1965), others have identified conditions under which they may persist, e.g. De Long et al. (1990). In this paper, we analyze market ecologies, combinations of traders using different strategies, in order to understand the situations in which technical traders may make long term positive returns. We conclude that, while in many circumstances fundamentalist traders make the highest returns, both trend following and contrarian traders may make positive profits, dependent on the market mechanism.

We consider a market populated with homogeneous fundamentalists and heterogeneous technical traders. Fundamentalists trade on departures of the realized price from its un-

*School of Economics, University of East Anglia, United Kingdom. Email address: antony.jackson@uea.ac.uk. Telephone: +44 (0)1603 59 3876.

†Department of Economics, University of Leicester, United Kingdom. Email address: d.ladley@leicester.ac.uk. Telephone: +44 (0)116 252 5285.

derlying value, whilst chartists¹ trade on patterns in historical prices. Trader demands—channeled through a Walrasian auctioneer—determine the realized market price, which in turn determines the flow of wealth between different strategy types. In implementing the technical trading rules, we use the notion of an acceptable risk-adjusted level of profitability. This enables buy/sell signals to be mapped to price expectations, thus enabling the construction of an ecology of diverse, realistic trading rules. By allowing the inclusion of a large variety of trading strategies, we are able to assess the impact of a complex range of strategy interactions and draw conclusions about the likelihood of their survival.

Fundamentalists and technical traders observe the same freely available fundamental value. Technical traders, however, pay no attention to this value, instead relying on a variety of trading rules commonly investigated in the empirical literature and employed by real-world traders. These include the ‘trend following’ momentum, moving average and channel rules. Natural counterparts to these trend following rules are contrarian rules, which we model by reversing trend following signals from ‘buy’ to ‘sell’ and from ‘sell’ to ‘buy’. Each rule has a single parameter, which is drawn from a set of values that correspond to common data window sizes, such as weekly, monthly, quarterly and semi-annual windows. Together with an ever-present fundamentalist strategy, we run a simulation for each permutation of parameter values for the trend following and contrarian strategy types.

Markets clear through a Walrasian auctioneer, who is tasked with finding an equilibrium price in each period. It has been common practice to construct traders’ demand schedules using prices up to—but not including—the current price. This has the convenience of producing a smooth aggregate demand curve and a unique market clearing price. However, what looks at first glance like a fairly innocuous assumption, creates a market conducive to the profitability of positive-feedback strategies. We contrast this setting with one in which traders condition their behavior on all prices up to the current period, and henceforth we refer to ‘lagged prices’ and ‘real time prices’ Walrasian auction protocols.

We first present results showing the volatility and efficiency of the markets. We find that under the lagged information setup, volatility is reduced in the presence of trend followers, while the difference between the fundamental price and the market price is lower when there are more technical traders basing their decisions on short price windows. If markets are dominated by contrarian traders using up to date information, there is a risk of disconnection between the market price and the fundamental. The market may experience low volatility, but a higher tracking error. In contrast, if trend following traders are present in these markets, volatility is increased, but the disconnection is decreased.

We go on to analyze the effect of the market ecology on the profitability of trading rules. We find that some initial trading rule endowments are more conducive to the profitability of individual strategies than others, which we illustrate through the use of sensitivity tables which capture average interaction effects. There is a clear message: under the assumption of conditioning with respect to lagged prices, both fundamentalists and trend followers accrue profits—at the expense of contrarians. In contrast, when using the most up-to-date price information, not only do contrarian strategies outperform trend followers, but they also accumulate long-run profits. Crucially, the overall level of profits

¹We use the terms ‘technical trader’ and ‘chartist’ interchangeably. It is also common to refer to fundamentalists as arbitrageurs.

and losses is much reduced in the real time protocol: the market seems better suited to its primary purpose of providing prices that properly reflect fundamental value.

2 Related Literature

The Efficient Markets Hypothesis (Fama, 1965) and the related Rational Expectations Hypothesis (Muth, 1961; Lucas, 1972) formalized the Friedman (1953) notion that irrational speculators will eventually be driven out by rational arbitrageurs. Critics of this approach, however, argue that there is too high a computational burden imposed on agents in order for them to calculate the rational expectations equilibrium. Real world examples demonstrate this. For instance, investors seemingly failed to learn the lessons of the NASDAQ crash: barely a decade later a new bubble developed, this time in the United States real estate market, with subsequent dire consequences for the global financial system.

The idea that agents operate within ‘bounds of rationality’, or that they use rules of thumb instead of complex calculations, can be traced back to Simon (1957) and Tversky and Kahneman (1974). Sargent (1993) and Evans and Honkapohja (2001) provide more recent examples of models in which heterogeneous agents learn using boundedly rational rules. Within a simple cobweb model, Brock and Hommes (1997) describe an ‘adaptively rational equilibrium’ in which heterogeneous agents choose between a costly rational expectations forecast or a costless naive forecast. The composition of the population evolves according to feedback on the success of each strategy. The responsiveness of agents to this feedback—the ‘intensity of choice’ parameter—determines whether there is a stable rational expectations equilibrium. The Brock and Hommes (1998) adaptation of this evolutionary learning mechanism to a simple asset pricing model is one contribution to a large literature on simulation-based approaches to heterogeneous agent modeling. There are a number of recent surveys of this literature, including Hommes (2006), LeBaron (2006), Chiarella et al. (2009) and Hommes and Wagener (2009).

A common feature of many models is the role played by technical analysis. A number of influential survey articles have pointed to the emphasis placed by foreign exchange professionals on this approach. Allen and Taylor (1990) and Taylor and Allen (1992) present survey evidence from the London foreign exchange market. They find that 90 percent of correspondents use technical analysis in forming exchange rate expectations, with a particular emphasis on its use at short forecasting horizons. They suggest that chartists may have an intuitive handle on nonlinearities in short-term exchange rate dynamics, but that linear fundamental models prevail in the long run. For a recent discussion of heterogeneity in exchange rate expectations, see Menkhoff et al. (2009). Early empirical studies focusing on foreign exchange include Sweeney (1986) and Levich and Thomas (1993), while Menkhoff and Taylor (2007) survey the technical strategies used in these markets.² Qi and Wu (2006) examine the common stock market rules of Brock et al. (1992) and Sullivan et al. (1999) and apply them to foreign exchange markets.

In early exchange rate models, prices were formed by weighting the expectations of heterogeneous agents (Frankel and Froot, 1990). Frankel and Froot (1986) proposed a model in which portfolio managers weight the forecasts of chartists and fundamentalists, using Bayesian updating to refine those weights. The model is used to construct a hypothesis

²For a broader survey, encompassing other asset classes, see Park and Irwin (2007).

for the excessive strength of the US dollar in the 1980s. De Grauwe et al. (1993) suggest that the influence of fundamentalists' expectations becomes more important the farther the exchange rate is from its equilibrium. In their model, heterogeneous expectations are used in the construction of a chaotic model of the exchange rate. These types of models have been used to examine the consequences of policy intervention (Bauer et al., 2009; Wieland and Westerhoff, 2005), the causes of excess volatility Bauer and Herz (2004), or in the case of De Grauwe and Grimaldi (2005), to examine the effect of frictions on exchange rate dynamics.

While the interaction of technical traders has been considered in foreign exchange markets (and markets for assets more broadly, see for example Westerhoff and Franke (2012); Phelps and Ng (2014)), the scope of those strategies analyzed has been limited. The majority of models have focused on limited trader types and so have simplified the complex interactions which may occur in these markets. In this paper, we analyze in detail the effect of the ecology of trading strategies, examining how the presence of different trader types may increase or decrease the profits of an investor following a different approach. This analysis allows us to comment not only on the profitability of different strategies, but also on the effect of technical traders on market behavior.

Our approach most closely follows the agent-based modeling of Brock and Hommes (1997). This work, however, only considered a single technical trading strategy interacting with a fundamentalist. In our paper, we incorporate multiple real world trading strategies in the same market in order to model the ecology of interactions. In choosing the particular strategies of our traders, we are influenced by the recent empirical results of Neely and Weller (2013), who provide evidence of persistence in the performance of back-tested adaptive technical analysis strategies in the foreign exchange market.

Building agent-based models with reference to the feedback from real-world foreign exchange strategies creates an interesting divergence from some of the more recent approaches in the agent-based literature (Cristelli et al., 2011; Sornette, 2014). In this literature, the trend has been to have individual strategies 'fade out' at the macroscopic level (see, for example, Gontis and Kononovicius (2014)). As a result the effect of the actions of a multitude of different traders is reduced to a stochastic process. This is in contrast to recent papers in the economics literature such as (Palczewski et al., 2016; Chiarella and Ladley, In Press) which model the behavior of specific technical trading strategies and in some cases e.g. Chiarella et al. (2014); Goldbaum and Zwinkels (2014) fit these strategies to market data. The two approaches are not entirely inconsistent, however, as the empirical evidence presented by Neely and Weller (2013) demonstrates individual strategies can enjoy long periods of profitability, only to 'fade out' as the market adapts, but then 'fade back in' (returning to profit) in later years. This paper attempts to bring together these views by showing how ecologies of traders may interact and produce market dynamics.

3 The Model

We model the exchange of a risky and a risk-free asset. The risk-free asset is held in domestic currency, whereas the risky asset is a foreign currency that yields a fixed rate of interest. Trade occurs daily, with all positions closed before the beginning of the following day and settled at today's price. Trading takes place between a large number of funda-

mentalist and technical traders through a single Walrasian auctioneer. At the beginning of each trading day, the underlying fundamental value of the asset is revealed (at no cost) to both the fundamentalists and the technical traders, but only the fundamentalists incorporate it into their decision-making process.

Both fundamentalists and technical traders are mean-variance optimizers, with constant absolute risk aversion (CARA) preferences. The auctioneer quotes a sequence of prices and collects and aggregates the demand schedules of all traders. With no outside supply of the risky asset, it is the auctioneer's role to set the price that induces aggregate excess demand to be as close to zero as possible.

3.1 Demand

All agents are modeled as having the exponential utility function

$$U(W) = -\exp(-\lambda W),$$

where the utility function $U(\cdot)$ is defined over wealth W and λ is the coefficient of absolute risk aversion. If it is assumed that conditional wealth at time $t+1$ is distributed normally with mean μ and variance σ^2 , then it can be shown that to maximize expected utility is equivalent to maximizing the value of

$$\mu - \frac{\lambda\sigma^2}{2}.$$

The one-period-ahead wealth equation is specified as

$$W_{t+1} = (1 + r^*) S_{t+1} Q_t + (1 + r)(W_t - S_t Q_t), \quad (1)$$

where r^* and r are the rates of return on the risky and risk-free assets, respectively. The interpretation of r^* is a fixed rate of interest held in a foreign currency. We abstract from sovereign default risk, so the volatility of the risky asset is due entirely to exchange rate movements.³ Q_t is the quantity of the risky asset purchased by the agent at time t and S_{t+1} is the spot price realized at time $t+1$. Solving the agent's maximization problem, subject to this wealth constraint, gives the optimal quantity of risky asset demanded as

$$Q_t^* = \frac{(1 + r^*)\mathbb{E}_t[S_{t+1}] - (1 + r)S_t}{\lambda(1 + r^*)^2\mathbb{V}_t[S_{t+1}]}, \quad (2)$$

where $\mathbb{E}_t[S_{t+1}]$ and $\mathbb{V}_t[S_{t+1}]$ are conditional expectations and variances of the future spot price.

3.2 Expectations

Equation 2 gives the demand for the optimal quantity of the risky asset. The numerator is interpreted as the expected excess return beyond the risk-free rate. This is composed

³Alternatively this asset could be interpreted as a stock, with the market being a standard asset market. In this case r^* would represent the (constant) dividend-price ratio for a stochastic stream of dividend payments. Given the prevalence of technical trading in the FX markets, as set out above, we define the model in these terms, however, the latter specification is also valid.

of dividends (or foreign interest earned) and capital gain from price movements. Traders have heterogeneous expectations of the price, with fundamentalists expecting a constant rate of convergence to a fundamental value and technical traders expecting an acceptable risk-adjusted level of profitability. Technical traders differ in their strategies and through the use of a diverse range of sample sizes of historical prices.

3.2.1 Fundamentalists

Fundamentalists' expectations are driven by privileged access to the underlying fundamental value of the asset. For simplicity, we assume that they have free access to this value⁴. We model the underlying fundamental value as a geometric Brownian motion. After applying Ito's Lemma, this gives the following process for the log of the value in discrete time:

$$\ln Y_t - \ln Y_{t-1} = \mu - \frac{1}{2}\sigma^2 + z_t\sigma,$$

where $Y(t)$ is fundamental value, z_t is an i.i.d. draw from a standard normal distribution and μ and σ are constants. Under the assumption that there is no drift term, $\mu = 0$, the equation for the evolution of the fundamental value is

$$\ln(Y_t/Y_{t-1}) = z_t\sigma - \frac{1}{2}\sigma^2. \quad (3)$$

Fundamentalists buy the risky asset when it is under-priced relative to its fundamental value, and sell it when it is overpriced. They do so with increased confidence the farther the asset is trading from its true level. They form expectations according to the rule

$$\mathbb{E}_t[S_{t+1}] = S_t + \gamma(Y_t - S_t), \quad \gamma \in [0, 1] \quad (4)$$

where Y_t is the realization of fundamental value at time t and γ is the rate at which fundamentalists expect the asset price to converge to its true level.

3.2.2 Technical Traders

A difficulty with the types of technical rules commonly examined in the empirical literature is that signals of the type 'buy', 'sell' or 'hold' are generated with no reference to the quantity to be traded. A key contribution of this paper is to propose that simple trading rules can be mapped into price expectations by using the notion of an acceptable risk-adjusted level of profitability. Risk-adjusted performance metrics (RAPMs) in the financial industry are dominated by the use of the Sharpe ratio (Sharpe, 1964), with variations such as the Sortino ratio (Sortino and van der Meer, 1991).⁵

Under the assumption that returns are i.i.d., the mean and standard deviation of the daily returns of a strategy or portfolio are translated into an annualized Sharpe ratio using the formula

$$SR_{\text{annual}} = \frac{252 \cdot (\mu_{\text{daily}} - r_{\text{daily}})}{\sqrt{252} \cdot \sigma_{\text{daily}}}, \quad (5)$$

⁴The impact of a cost to acquiring this information would be an interesting extension.

⁵The Sortino ratio belongs to a general class of statistics which only consider the volatility of returns below some target rate of return. Another example is the semi-variance measure, which only includes those returns below the sample mean.

where μ_{daily} is the mean daily return, r_{daily} is the risk-free rate and σ_{daily} is the daily volatility of returns. Lyons (2001) points to the use of the Sharpe ratio as a performance metric in banks and hedge funds. Grinold and Kahn (2000) and Menkhoff and Taylor (2007) suggest that the benchmark used for committing capital is a Sharpe ratio of 0.5. The daily expected return implied by a prior Sharpe ratio of 0.5 and Equation 5 corresponds to an expected appreciation or depreciation of the asset price. The direction depends on the current trading signal, leading to an expression for technical trader expectations of

$$\mathbb{E}_t[S_{t+1}] = \left(1 + r_{\text{daily}} + \frac{I_t \cdot \text{SR}_{\text{annual}} \cdot \sigma_{\text{daily}}}{\sqrt{252}} \right) S_t, \quad (6)$$

where I_t is the value of the trading signal $\in \{-1, 0, 1\}$.

3.3 Trading Rules

We select three classes of technical trading rules that tend to dominate the empirical literature. As discussed above, the three rules do not offer obvious ways of using extrapolation to derive point forecasts. Instead, they have been designed with a view to signaling possible changes of trend, rather than being used as forecasting methods. We consider trend following versions of these rules— as set out below—and contrarian versions, which use the same rule, but simply reverse the sign of the output.

3.3.1 Momentum

In this study, momentum is defined as a simple threshold rule in which a buy signal is given if the rate of change of prices is positive over a defined period, a sell signal is given if the rate of change is negative, and in the rare case where there is no rate of change, a neutral signal. Rate of change is measured with respect to the endpoints of the sample only; no attempt is made to model the rate of change by incorporating all of the sample observations. Much of the literature on momentum is equity-based (Chan et al., 2000; Conrad and Kaul, 1998), but the general principle is the same across all asset types— strong price movements are expected to follow strong price movements. The rule measures whether the rate of change is positive or negative over a time period n , which includes today’s price S_t

$$I_t(n) = \begin{cases} +1 & \text{if } S_t > S_{t-n} \\ 0 & \text{if } S_t = S_{t-n} \\ -1 & \text{if } S_t < S_{t-n} \end{cases}$$

3.3.2 Moving Average

Moving average rules vary with respect to the weighting scheme used (e.g. exponential and simple) as well as with respect to the method of implementation, which can involve several averages in combination. Menkhoff and Taylor (2007) discuss the method in which a long position is taken upon a short period moving average crossing a long period moving average from below, and a short position when the short period moving average crosses the long period from above. Moving average combinations are also considered in Brock

et al. (1992), Sullivan et al. (1999) and Qi and Wu (2006), amongst others. We use the simple moving average

$$\text{SMA}_t = \sum_{i=t-n}^t S_i / (n + 1)$$

The indicator function for the simple moving average rule is

$$I_t(n) = \begin{cases} +1 & \text{if } S_t > \text{SMA}_t \\ 0 & \text{if } S_t = \text{SMA}_t \\ -1 & \text{if } S_t < \text{SMA}_t \end{cases}$$

The simple moving average is similar to the momentum rule except that the test condition relates to the average of all the prices in the sample, rather than just the start and end points.

3.3.3 Trading Range Breakout

Also known as a support and resistance or channel method, this trading style is event-driven, in that the breakout is intended to signal a change in market conditions. Early investigations of channel rules include Irwin and Uhrig (1984), Lukac et al. (1988) and Taylor (1994). Several variants appear in the modern bootstrap-based studies. Brock et al. (1992) consider channels of 50, 150 and 200 days, where a channel is defined as those prices contained by a local maximum and minimum. It is argued that at these extremes many investors are willing to buy or sell, with the minimum acting as ‘support’ and the maximum acting as ‘resistance’; the eventual ‘breakout’ is deemed to be a significant event. We use an n -day breakout indicator, where the indicator takes a positive value if the most recent close is greater than the highest high of the previous n prices.

$$I_t(n) = \begin{cases} +1 & \text{if } S_t > \max(S_{t-1}, S_{t-2}, \dots, S_{t-n}) \\ -1 & \text{if } S_t < \min(S_{t-1}, S_{t-2}, \dots, S_{t-n}) \\ I_{t-1} & \text{otherwise} \end{cases}$$

At the start of the simulation, the indicator will be zero until the first breakout. The indicator then maintains this value until a breakout occurs in the opposite direction.

3.4 Walrasian Auction Protocols

Trade is conducted through a Walrasian auctioneer. In the literature, the standard implementation assumes traders use prices up to, but not including, the current price being quoted by the auctioneer. This is a matter of convenience, as it provides smooth individual demand functions and a unique market-clearing price. However, it is a simplified view of how technical traders behave. In this paper, we also consider an auction mechanism in which expectations are adjusted contemporaneously with the sequence of prices quoted by the auctioneer. We refer to the two auction protocols as ‘lagged’ and ‘real time’.

In both the lagged and the real time versions of the model, the current fundamental value is revealed before trade. In both versions, fundamentalists adjust their expectations contemporaneously with the auctioneer’s quote. The difference between the two mechanisms lies in the way technical traders update their rules. For example, a trader using

a lagged version of the breakout rule has already decided whether to be long or short, irrespective of the auctioneer’s quote. His quantity depends on the difference between his expectation of the price and the auctioneer’s quote. In contrast, a trader using the real time version of the same breakout rule notes that quotes lying above (below) the current maximum (minimum) channel threshold will result in a buy (sell) signal. The trading signal for the current round of trade updates at the same time as the auctioneer’s quote. These threshold values create discontinuities in the real time breakout trader’s demand schedule. Discontinuities in the aggregate demand schedule follow.

Note we are still striving for a market clearing solution in each time step, so there is a material difference between this paper’s real-time Walrasian auction and a market making model in which the current price is known to traders. The market making model perpetually trades out of equilibrium, with the market maker’s inventory taking up the slack of excess demand or supply. Under the assumption of a risk-neutral market maker, there are no bounds to the size of this inventory.

3.4.1 Lagged Price Information

[Figure 1 about here.]

With expectations formed from a combination of a trading signal $I_{t,i}$ and a prior Sharpe ratio SR_{annual} , it can be shown that technical trader aggregate demand, as a function of the current price quote S_t , is given by

$$D^T(S_t) = \sum_i \frac{I_{t,i} w_i \cdot SR_{\text{annual}}}{\lambda \sigma S_t \cdot \sqrt{252}},$$

where i indexes the technical rules and w_i is the proportion of the population using a particular rule. Similarly, fundamentalist aggregate demand is

$$D^F(S_t) = \frac{w_f \gamma (Y_t - S_t)}{\lambda \sigma^2 S_t^2},$$

where w_f is the proportion of the population employing the fundamentalist strategy. With no outside supply of foreign currency, adding D^T and D^F and setting to zero yields an expression for the market clearing rate of

$$S_t^* = \frac{w_f \gamma Y_t}{w_f \gamma - \frac{1}{\sqrt{252}} \cdot \lambda \sigma \sum_i \frac{I_{t,i} w_i \cdot SR_{\text{annual}}}{\lambda}}. \quad (7)$$

The lagged prices case treats $I_{t,i}$ as exogenous in each time step, leading to a unique market clearing price S_t^* . This approach also has the advantage of a low computational burden for the calculation of the market clearing price in each time step. Panel A of Figure 1 illustrates this case.

3.4.2 Real-time Price Information

The setting of prices is more complex if traders condition their demand on the current trade price. In the real time setting, $I_{t,i}$ now becomes an endogenous variable of Equation 8:

$$S_t^* = \frac{w_f \gamma Y_t}{w_f \gamma - \frac{1}{\sqrt{252}} \cdot \lambda \sigma \sum_i \frac{I_{t,i}(S_t^*) w_i \cdot \text{SR}_{\text{annual}}}{\lambda}}. \quad (8)$$

Equation 8 makes explicit the joint determination of the market clearing price and trading signals. This interaction may produce discontinuities in the demand curve which are a result of discontinuities in the trading rules, i.e., points at which traders change the nature of their demand. The result of this is that there are three qualitatively different scenarios which do not fit with the concept of a single market clearing price. The first, in Panel B of Figure 1, shows an example in which the discontinuities do not preclude the existence of an equilibrium or market clearing exchange rate. The curve crosses the $D(S_t) = 0$ line at a unique point. The next possible outcome (Panel C) shows the non-existence of an equilibrium crossing point. This raises economic as well as computational questions. If the market cannot clear at equilibrium, where does the market trade? And if the market trades out of equilibrium, how is excess demand dealt with? In this example, just as the demand curve is about to intersect the zero line, certain strategies adjust their signals so as to create gaps both before and after the likely crossing point. Finally, by updating their information sets to include the current quote, traders can generate multiple equilibria (Panel D) leading to the question of which price should be chosen.

Additional assumptions need to be made about how the auctioneer goes about calculating the traded price for each potential scenario. Whilst in some cases it may be possible to use a technique such as the bisection method, in general the potential multiple or zero crossing points mean this is not possible. It is therefore necessary to use a computationally demanding search. We define the set of prices considered as: $\{S_L, S_L + \delta, \dots, S_U - \delta, S_U\}$, where S_L marks the lower limit of the search range, δ is the price increment, and S_U is the upper limit. In all cases, $\delta = 0.001$ or 10 basis points, which we judge to be a fair trade-off between precision and speed⁶. The boundary conditions are obtained by computing the market clearing price that would occur if all technical traders were buyers or if they were all sellers. This is achieved by setting all trading signals, $I_{t,i}$, to either +1 or -1 in Equation 8. In each case, we select the price with minimum excess demand. Since the set of search prices is discrete, the excess demand will never be zero. The auctioneer deals with this by rationing members of the larger group in proportion to individual demands. In essence, the auctioneer operates as an order matcher. An alternative would be for them to store excess quantity in a similar manner to a market maker. However, this would potentially add or remove wealth from the trading environment, and complicate the dynamics.

4 Results

We consider markets containing combinations of the six classes of trading strategies described above (momentum, momentum contrarian, moving average, moving average contrarian, breakout and breakout contrarian) using four different historical sample sizes: $n \in \{5, 21, 63, 126\}$. These values were chosen to represent psychologically important periods of a week, a month, a quarter and half a year. In total, this gives $6 \times 4 = 24$

⁶Our simulations are intended to represent trading at the daily frequency, so it may be more appropriate to use a finer precision for simulations which are intended to capture activity at higher frequencies.

different technical trading strategies. Assuming the fundamentalist strategy is present within all simulations, this gives $2^{24} = 16,777,216$ possible strategy combinations. Each strategy within a simulation is allocated an equal share of the population i.e., $1/n$.⁷ For each combination of trading strategies, we perform 1000 repetitions of the simulation, with different underlying time series of fundamental value. The same underlying price series are used for all trader combinations, ensuring that variations are due to agent interactions. Each simulation runs for 3024 days (12 years), with each day consisting of a single market-clearing time step. Data is only recorded for the final 2520 days (10 years) in order to allow the rules with longer windows to base their behavior on market data. The volatility of log returns in the underlying geometric Brownian motion is 1% and we assume that all traders use this value as their estimate of future volatility. Hence, heterogeneous expectations of future volatility are not driving realized volatility⁸. All risk aversion parameters are set to 1, and all prior Sharpe ratios are set to 0.5. Note, as long as they are identical across individuals, neither of these parameters qualitatively affect results as they merely scale the size of investment that traders make. The fundamentalists use a convergence parameter of 0.01: they expect 1% of the difference between the actual price and the underlying value to be eradicated each day. The dimensionality of the simulation problem is reduced by holding the fundamentalists' parameter fixed at this value. These parameters are summarized in Table 1.

[Table 1 about here.]

4.1 Market Statistics

We first present results showing the effect of the composition of the population of traders on market behavior. We focus on two aspects: efficiency—as measured by the mean squared error between the market prices and the underlying fundamental price—and price volatility. Statistics are calculated for each permutation of the 2^{24} technical strategies. Results are presented for each trading rule averaged over all simulations in which they are present. The first column contains the average volatility over all simulations for each type of trader. The remaining columns give the volatility, as a percentage of the base volatility, for the case when traders using a particular window size for the rule are present. Note since these values are percentages relative to the mean and trend followers and contrarians take opposite decisions at all times the values are symmetric around 1.

[Table 2 about here.]

Table 2 presents the results for price volatility. In line with previous findings (Chiarella et al., 2009; Hommes and Wagener, 2009), the technical traders have a significant effect on the dynamics of the market price. There is, however, a surprising difference between the

⁷This decision was partially made to bound the analysis to a computationally tractable number of simulations. Experiments were, however, performed for a limited number of cases with non equal population sizes. The results of interactions were found to be qualitatively similar to those presented below.

⁸Endogenous estimation of volatility is an interesting avenue for future research. In the same manner in which there are many ways to estimate prices, there are also many ways to estimate volatility, and potentially many complex ways in which they could interact. We therefore leave this aspect to future work, in order to focus on the dynamics of price rules.

real time and lagged specifications. Volatility is approximately four times greater under the lagged specification than the real time equivalent. This indicates that if traders use real time information, the destabilizing effect of technical traders may be lower than previous estimates suggest. Notably, however, the real time price volatility is itself approximately three times higher than the underlying fundamental volatility. The use of real time information allows traders to avoid some trading mistakes which destabilize prices. For instance, a trader may avoid buying into a trend which was previously present, but which does not persist in the current period.

The effect of different types of traders on volatility is not uniform. In the real time setting, those markets which contain trend followers have higher volatility than those containing contrarians. The contrarian traders act to dampen the market, resisting trends in price changes. In contrast, trend followers actively move the price away from its current value. In the lagged setting, the pattern is reversed. The contrarians, using out of date information, destabilize the price. In many cases this is because their efforts ignore those of the fundamentalist trader. For instance, if a trend appeared in period $t - 1$ away from the fundamental, the fundamentalist would take a position which would act to correct this. Contrarians would take a similar position (in reverse of the existing trend) making the reversion larger. In contrast, trend followers would oppose the fundamentalist, reducing the size of the reversion and consequently the overall volatility.

[Table 3 about here.]

The mean squared error between the market price and the fundamental price gives a measure of the efficiency of the financial market: its ability to price an asset fairly. Table 3 shows that this value is higher for the lagged case than the real time setting. This is unsurprising, as the real time protocol uses more up-to-date information in the calculation of prices, and therefore it should be expected that the fundamental price and the market price will be closer.

Under the lagged protocol, short-period momentum and moving average traders decrease tracking error, whilst those with longer periods increase it (and vice versa for contrarians). In this setting, it is those traders who are most volatile in their opinions who do most to make the market price accurate. This is because changes in the fundamental value often take time to be incorporated into the price. For instance, a fundamentalist corrects mispricings by 1% per day, meaning there is some persistence in deviations. Short period trend followers may inadvertently capture the trend associated with these deviations, speeding up the mean reversion. Longer period traders, in contrast, may capture deviations which have already been corrected, and so move the price away from the fundamental.

In the real time case, when contrarian traders are present, the tracking error is increased. However, as Table 2 shows, this is accompanied by lower volatility. Whilst the market may have been relatively calm, there is a disconnection between the market price and fundamental value. In contrast, the higher volatility evidenced in this setting when trend followers are present is associated with lower tracking error. The trend followers, whilst causing volatility, must be moving the price back towards the fundamental price.

4.2 Profits

In this section, we consider the average profits of each trading strategy across all possible market ecologies. We collect results for fundamentalists, trend followers and contrarians, aggregating across time windows. We calculate average annual profits for each permutation of the 2^{24} technical strategies. Given that the number of strategies varies in each ecology, we present our results as profits per trader.⁹ Total profit is calculated as the sum of daily profits i.e. the profit to the trader from taking their desired position in day t and liquidating that position on day $t + 1$.¹⁰

[Figure 2 about here.]

The left column of Figure 2 presents the results for lagged information. The losses of contrarian strategies largely feed the profits of the fundamentalists, who make positive returns in all scenarios. There is also evidence that trend-followers gain from some of these losses, as they benefit from positive-feedback effects. In this setting, there is little difference between the trend-following momentum and moving average rules: both distributions are centered around zero and have long right tails. In contrast, the breakout rule exhibits a tighter, symmetrical distribution. This is because the breakout rule, only being concerned with the high or low of a range of prices, is insensitive to erratic price movements within those bounds. The endpoints of the momentum rule and the value of the moving average change with every new price—resulting in higher profit volatility, which in some circumstances can result in very high gains.

In line with the price deviation results of the previous section, the right column of Figure 2 shows that the absolute levels of profits and losses are much reduced when technical traders update their orders contemporaneously with prices. We have accordingly adjusted the range of the horizontal axes in the right column of Figure 2 to improve clarity. This change in the way traders update their orders means that in each time step each trader is behaving in a way they believe is the optimal manner—trading at a price consistent with their beliefs. In the lagged price scenario, a trader’s actions are in essence out of date. The direction of their trade may be the opposite of what would be desired if the current market price were included. This seemingly innocuous change in the way traders process price information limits the number of profit opportunities.

An examination of the composition of winners and losers reveals the driving force behind this increased market efficiency. All of the trend-following strategies now accumulate losses, with the fortunes of contrarians also reversing. In particular, the distribution of the moving average contrarian strategies almost exactly matches that of the fundamentalists, suggesting that technical strategies can mimic arbitrage strategies if information is processed in a timely manner.

It is clear that the ability of trend-following traders to push prices for their own benefit is largely removed when the population as a whole updates prices in real time. In the absence of positive-feedback effects, temporary deviations from fundamental value are quickly eradicated by fundamentalists, who now face fierce competition from contrarians

⁹Calculated as gross profit divided by population share.

¹⁰Traders frequently do not change their position day to day. They could therefore be viewed as closing and opening a position of equal size at the same price on day $t + 1$. As a result this is equivalent to the sum of the profits and losses of trading over the length of the simulation.

for these limited wealth transfers. The limits to arbitrage have been overcome by those members of the noise trading population who push prices back to equilibrium. This finding also relates to the ‘hipster effect’ (Touboul, 2014). In our model contrarians are able to synchronize and make profits in real time by opposing the trend in the market. However, under the lagged market mechanism the increased delay leads to the same traders missing opportunities and losing money.

4.3 Strategy Interaction

In this section, we consider the interaction of trading strategies, identifying those ecologies that are most conducive to profit or loss-making for each type. Strategies are classified both by rule type and by the size of the data window they use. The first column of each table presents the unconditional mean of annual profits by strategy type across all simulation in which that type of trader is present. White cells represent profitable strategies, whilst gray cell represent loss making approaches. Subsequent columns measure the impact of each alternative strategy type on the row type’s profit. These values are averaged across all simulated ecologies where the column type and row type are both present (potentially along with other types). For the columns we aggregate across window lengths to aid comparison. Cells are shaded gray if the inclusion of that particular strategy decreases profitability, and white if their inclusion is beneficial. For example, in Table 4 fundamentalist profits have an unconditional mean of 122.28. The inclusion of trend following rules within the population diminishes these profits, whereas the inclusion of contrarian rules—in particular moving average rules—increases profitability. For those cells on the diagonal the traders own type are removed from the comparison as they are always present. For example for breakout trend following traders with a window of 5 the breakout trend following column contains the effect of breakout trend following traders with all other window lengths. We have highlighted those strategies that have the largest impact on profitability in bold-type.

Table 5 confirms the reduction in the scale of wealth transfers when traders use up-to-date price information. Otherwise, the shading of the table matches that of Table 4, but with one major difference: fundamentalists accrue greater profits when the rest of the population is dominated by trend-following rules. The histograms demonstrate a similarity between the distributions of fundamentalists and contrarians—Table 5 adds further evidence that contrarians trade in the same direction as fundamentalists. Even though the contrarians can be regarded as noise traders, they behave much the same as fundamentalists, alleviating the limits to arbitrage evident in the lagged prices protocol.

Fundamentalists and trend followers benefit from the inclusion of contrarian rules, but suffer from the inclusion of trend followers. Sensitivities are the opposite for contrarians, who benefit from the presence of trend followers, but suffer in the presence of other contrarians. The inclusion of traders of an opposite class provides strategies with individuals willing to take opposite positions—leading to potentially greater volumes. However, traders of the same type serve as competition. This leads to the symmetry of effects seen in Table 4. For example for breakout traders with a time window of 5, trend following moving average traders decrease profits by 2.05 whilst the presence of contrarian moving average traders increase profits by the same amount. As the two groups always take the

opposite actions the effects are equal and opposite.¹¹ Table 5 shows a similar effect except that there is some noise in the pairs of values. This is due to the equilibrium price selection process. Random selection among multiple possible equilibrium prices inducing a stochastic element to the payoffs. In general this effect is quite small, however, in cases where the effect of a trading strategy on another is very small - such as for trend following break out traders it can change the sign of the effect.

Within each class, it is evident that in the majority of cases strategies are most sensitive to the inclusion of moving average rules, closely followed by momentum rules and, to a much lesser extent, breakout rules. This reflects the frequency with which indicator functions change sign. As highlighted above, the breakout indicator is measured with reference to a range, which can remain unaltered for long periods. The indicators of the other rules, however, are measured against values that change with each subsequent price. The pattern observed for the longest window trend following momentum traders in Table 4 is worthy of note. These traders do not show the same pattern as other trend following momentum traders – they make significant losses. This is because these traders are the slowest to respond to price changes - particularly when there have been large previous price movements. The trading signal is based solely on the current price and that 126 day's ago. Whilst this price changes every day it can take a long time for it to reflect large price changes. As a result the trader may buy into trends for long periods that do not exist and miss new trends, resulting in all other traders exploiting them to a lesser or greater extent. This effect will be particularly prevalent around times the market changes direction - leading to the positive interaction with trend following break out traders. A similar effect is seen for the contrarian momentum traders with window 126 except that the larger average losses for contrarian traders in general mean that this effect is not sufficient to cause a change in sign.

[Table 4 about here.]

¹¹It is worth highlighting that for those values on the diagonal this rule does not hold as the table excludes the effect on traders of their own type. For instance for trend following break out traders with a window of 5, the column giving the effect of other trend following break out traders does not have the same value as that for contrarian break out traders. This is because the first column records the effect of trend following break out traders with windows of 21, 63 and 126 whilst the second also includes the effect of contrarian break out traders with a window of 5.

[Table 5 about here.]

5 Conclusions

This paper has considered an environment in which ecologies of traders, consisting of multiple varieties of heterogeneous technical strategies, interact with fundamentalist traders. We have utilized two different market protocols, one of which assumes that market participants include the currently-being-determined price in their decision making. Fundamentalists submit demands based on an assumed rate of convergence to a known fundamental level of a risky asset. Technical traders make decisions using historical windows of market prices and by invoking the notion of the annual risk/reward profile of an ‘acceptable’ trading strategy.

We find that in many settings the presence of a technical trader reduces market volatility or increases the accuracy with which the market price tracks the underlying fundamental value. This is particularly true for trend following strategies based on short term price patterns. In terms of the long-run profitability of strategies, when technical traders use a lagged information set, trend following traders induce positive-feedback effects—effects that benefit themselves, but cause contrarians to incur losses. The levels of profits and losses in the market are substantial, indicating that market prices are doing a poor job of instantaneously adjusting to the arrival of new information. The consequences of the ‘real time’ setting, in which technical traders update their rules using the most up-to-date information, are profound. Not only are the absolute levels of profits and losses significantly reduced—as arbitrage opportunities are quickly eradicated—but contrarian traders now join fundamentalists in profiting from this return to fundamental value. We also show that the inclusion of competing strategies within an ecology can help or hinder the profitability of individual strategies.

Our approach invites further development. The set of trading strategies considered in this paper was constrained to a number which could be analyzed effectively. In reality, there are many different classes of strategy which we do not consider, but whose presence could radically affect the behavior of the market. This paper also focuses primarily on profits, but an important next step is to analyze market dynamics over time in the presence of changing population shares. If individuals adapt their strategy based on recent profits, then the profile of profits in a particular ecology gives the instantaneous derivative of population change. A detailed analysis of this type of dynamic would potentially allow the identification of groups of strategies which lead to fixed populations, and those which lead to chaotic mixes.

References

- Allen, H., Taylor, M. P., 1990. Charts, noise and fundamentals in the london foreign exchange market. *Economic Journal* 100 (400), 49–59.
- Bauer, C., De Grauwe, P., Reitz, S., 2009. Exchange rate dynamics in a target zone—a heterogeneous expectations approach. *Journal of Economic Dynamics and Control* 33 (2), 329–344.

- Bauer, C., Herz, B., 2004. Technical trading and the volatility of exchange rates. *Quantitative Finance* 4 (4), 399–415.
- Brock, W. A., Hommes, C., 1997. A rational route to randomness. *Econometrica* 65 (5), 1059–1096.
- Brock, W. A., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22 (8-9), 1235–1274.
- Brock, W. A., Lakonishok, J., Lebaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance* 47 (5), 1731–64.
- Chan, K., Hameed, A., Tong, Wilson, H. S., 2000. Profitability of momentum strategies in the international equity markets. *Journal of Financial and Quantitative Analysis* 35 (2), 153–172.
- Chiarella, C., Dieci, R., He, X., 2009. Heterogeneity, market mechanisms, and asset price dynamics. In: Hens, T., Schenk-Hoppe, K. R. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier, pp. 277–344.
- Chiarella, C., He, X.-Z., Zwinkels, R. C., 2014. Heterogeneous expectations in asset pricing: Empirical evidence from the S&P500. *Journal of Economic Behavior & Organization* 105 (C), 1–16.
- Chiarella, C., Ladley, D., In Press. Chasing trends at the micro-level: The effect of technical trading on order book dynamics. *Journal of Banking and Finance*.
- Conrad, J., Kaul, G., 1998. An anatomy of trading strategies. *Review of Financial Studies* 11 (3), 489–519.
- Cristelli, M., Pietronero, L., Zaccaria, A., 2011. Critical overview of agent-based models for economics.
URL <http://arxiv.org/pdf/1101.1847.pdf>
- De Grauwe, P., Dewachter, H., Embrechts, M., 1993. *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets*. Blackwell Publishers, Oxford, UK.
- De Grauwe, P., Grimaldi, M., 2005. Heterogeneity of agents, transactions cost and the exchange rate. *The Journal of Economic Dynamics and Control* 29, 691–719.
- De Grauwe, P., Grimaldi, M., 2006. *The Exchange Rate in a Behavioral Finance Framework*. Princeton University Press, New Jersey.
- De Long, J. B., Shleifer, A., Summers, L. H., Waldmann, R., 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98 (4), 703–38.
- Evans, G. W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- Fama, E. F., 1965. The behavior of stock-market prices. *Journal of Business* 38, 34–105.

- Frankel, J. A., Froot, K. A., 1986. Understanding the u.s. dollar in the eighties: The expectations of chartists and fundamentalists. *The Economic Record*, Supplementary Issue 62, 24–38.
- Frankel, J. A., Froot, K. A., 1990. Chartists, fundamentalists, and trading in the foreign exchange market. *American Economic Review* 80 (2), 181–85.
- Friedman, M., 1953. The case for flexible exchange rates. In: *Essays in Positive Economics*. University of Chicago Press, Chicago, IL.
- Goldbaum, D., Zwinkels, R. C., 2014. An empirical examination of heterogeneity and switching in foreign exchange markets. *Journal of Economic Behavior & Organization* 107 (PB), 667–684.
- Gontis, V., Kononovicius, A., 2014. Consentaneous agent-based and stochastic model of the financial markets. *PLoS ONE* 9 (7).
- Grinold, Richard, C., Kahn, Ronald, N., 2000. *Active Portfolio Management: A quantitative approach for producing superior retrurns and selecting superior money managers*. McGraw-Hill, New York.
- Hommes, C., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K. L. (Eds.), *Handbook of Computational Economics*, Vol. 2. North-Holland, Amsterdam, pp. 1109–1186.
- Hommes, C. H., Wagener, F. O. O., 2009. Complex evolutionary systems in behavioral finance. In: Hens, T., Schenk-Hoppe, K. R. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. Elsevier, pp. 217–276.
- Irwin, S. H., Uhrig, J. W., 1984. Do technical analysts have holes in their shoes? *Review of Research in Futures Markets* 3, 264–277.
- LeBaron, B., 2006. Agent-based computational finance. In: Tesfatsion, L., Judd, K. L. (Eds.), *Handbook of Computational Economics*, Vol. 2. North-Holland, Amsterdam, pp. 1187–1233.
- Levich, R. M., Thomas, L. R., 1993. The significance of technical trading rule profits in the foreign exchange market: A bootstrap approach. *Journal of International Money and Finance* 12, 451–474.
- Lucas, R. E., 1972. Econometric testing of the natural rate hypothesis. In: *The Econometrics of Price Determination Conference*. Board of Governors of the Federal Reserve System and Social Science Research Council, Washington DC, pp. 50–59.
- Lukac, L. P., Brorsen, B. W., Irwin, S. H., 1988. A test of futures market disequilibrium using twelve different trading systems. *Applied Economics* 20, 623–639.
- Lyons, R. K., 2001. *The Microstructure Approach to Exchange Rates*. MIT Press, Cambridge.

- Menkhoff, L., Rebitzky, R. R., Schroder, M., 2009. Heterogeneity in exchange rate expectations: Evidence on the chartist-fundamentalist approach. *Journal of Economic Behavior & Organization* 70 (1-2), 241–252.
- Menkhoff, L., Taylor, M. P., 2007. The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature* 45 (4), 936–972.
- Muth, John, F., 1961. Rational expectations and the theory of price movements. *Econometrica* 29, 315–335.
- Neely, C. J., Weller, P. A., 2013. Lessons from the evolution of foreign exchange trading strategies. *Journal of Banking & Finance* 37, 3783–3798.
- Palczewski, J., Schenk-Hoppé, K. R., Wang, T., 2016. Itchy feet vs cool heads: Flow of funds in an agent-based financial market. *Journal of Economic Dynamics and Control* 63, 53–68.
- Park, C.-H., Irwin, S. H., 2007. What do we know about the profitability of technical analysis? *Journal of Economic Surveys* 21 (4), 786–826.
- Phelps, S., Ng, W. L., 2014. A simulation analysis of herding and unifractal scaling behaviour. *Intelligent Systems in Accounting, Finance and Management* 21 (1), 39–58.
- Qi, M., Wu, Y., 2006. Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange market. *Journal of Money, Credit, and Banking* 38 (8), 2135–2158.
- Sargent, T., 1993. *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures*. Oxford University Press.
- Sharpe, William, F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19, 425–442.
- Simon, Herbert, A., 1957. *Models of man; social and rational*. Wiley, Oxford, England.
- Sornette, D., 2014. Physics and financial economics (1776-2014): Puzzles, ising and agent-based models. *Reports on Progress in Physics* 77 (6), 062001.
- Sortino, Frank, A., van der Meer, R., 1991. Downside risk. *Journal of Portfolio Management* 17, 27–31.
- Sullivan, R., Timmermann, A., White, H., 1999. Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance* 54 (5), 1647–1691.
- Sweeney, R. J., 1986. Beating the foreign exchange market. *Journal of Finance* 41, 163–182.
- Taylor, M. P., Allen, H., 1992. The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance* 11 (3), 304–314.
- Taylor, S. J., 1994. Trading futures using a channel rule: a study of the predictive ability of technical analysis with currency examples. *Journal of Futures Markets* 14, 215–235.

- Touboul, J., 2014. The hipster effect: When anticonformists all look the same.
URL <http://arxiv.org/abs/1410.8001v1>
- Tversky, A., Kahneman, D., 1974. Judgement under uncertainty: Heuristics and biases. *Science* 185, 1124–1131.
- Westerhoff, F., Franke, R., 2012. Converse trading strategies, intrinsic noise and the stylized facts of financial markets. *Quantitative Finance* 12 (3), 425–436.
- Wieland, C., Westerhoff, F., 2005. Exchange rate dynamics, central bank interventions and chaos control methods. *Journal of Economic Behavior & Organization* 58 (1), 117–132.

Parameter	Description	Value
λ	Risk aversion	1.00
σ	Daily volatility	1%
SR_{annual}	Target risk-adjusted return	0.50
γ	Fundamental convergence	1%
n	Number of strategies	[1, 25]
w_i	Population shares	1/n

Table 1: The table documents those parameters common to all simulations. Traders share the same utility function, with coefficient of absolute risk aversion λ . The volatility of underlying fundamental value and traders' expected conditional volatility of returns is given by σ . Chartists' prior notion of an acceptable strategy is SR_{annual} and fundamentalists expect the exchange rate to converge to fundamental value at the rate γ . Populations are allocated equal fractions of the arbitrage strategy and [1, 24] technical strategies. Each simulation runs for 12 years.

		Mean	Trend Following				Contrarian			
			5	20	63	126	5	20	63	126
Lagged	Breakout	0.1201	0.930	0.965	0.981	0.986	1.070	1.035	1.019	1.014
	Momentum	-	0.855	0.869	0.895	0.940	1.145	1.131	1.105	1.060
	Moving Average	-	0.833	0.845	0.872	0.898	1.167	1.155	1.128	1.102
Real Time	Breakout	0.0288	1.075	1.053	1.041	1.027	0.925	0.947	0.959	0.973
	Momentum	-	1.093	1.078	1.061	1.050	0.907	0.922	0.939	0.950
	Moving Average	-	1.113	1.095	1.073	1.060	0.887	0.905	0.927	0.940

Table 2: Table presenting the average price volatility of markets in which a particular trader participates. The first column contains the average volatility of all simulations for the lagged and real time cases. The remaining columns detail the volatility relative to the base case when the specified trader type is present in the market.

		Mean	Trend Following				Contrarian			
			5	20	63	126	5	20	63	126
Lagged	Breakout	0.0150	1.030	1.042	1.028	1.025	0.970	0.958	0.972	0.975
	Momentum	-	0.972	0.998	1.011	1.043	1.028	1.002	0.989	0.957
	Moving Average	-	0.960	0.992	1.017	1.031	1.040	1.008	0.983	0.969
Real Time	Breakout	0.0013	0.997	0.991	0.990	0.992	1.003	1.009	1.010	1.008
	Momentum	-	0.995	0.994	0.992	0.994	1.005	1.006	1.008	1.006
	Moving Average	-	0.980	0.994	0.993	0.994	1.020	1.006	1.007	1.006

Table 3: Table presenting the average mean squared error between the market price and the fundamental for markets in which a particular trader participates. The first column contains the average mean squared error of all simulations for the lagged and real time cases. The remaining columns detail the mean squared error relative to the base case when the specified trader type is present in the market.

		Trend Follower				Contrarian		
		Mean	Breakout	Momentum	Moving Average	Breakout	Momentum	Moving Average
Fundamental	-	122.28	-7.51	-28.97	-36.38	7.51	28.97	36.38
Breakout	5	1.51	0.27	-1.66	-2.05	0.82	1.66	2.05
	21	-0.29	-0.20	-0.71	-0.89	0.36	0.71	0.89
	63	-0.05	-0.16	-0.31	-0.39	0.18	0.31	0.39
	126	0.08	-0.14	-0.18	-0.23	0.12	0.18	0.23
Momentum	5	11.09	-0.64	-3.96	-7.11	0.64	6.97	7.11
	21	10.67	-0.54	-5.00	-6.97	0.54	6.97	6.97
	63	10.49	-0.71	-5.34	-6.50	0.71	6.53	6.50
	126	-0.95	0.21	-0.68	-1.39	-0.21	1.48	1.39
Moving Average	5	15.05	-0.63	-6.15	-7.06	0.63	6.15	8.90
	21	13.91	-0.99	-6.43	-8.63	0.99	6.43	9.36
	63	12.38	-1.16	-6.17	-8.66	1.16	6.17	8.92
	126	10.88	-1.18	-5.69	-7.99	1.18	5.69	8.07
Breakout	5	-12.60	1.87	2.93	3.67	-0.16	-2.93	-3.67
Contrarian	21	-2.40	0.84	1.12	1.33	-0.51	-1.12	-1.33
	63	-0.69	0.36	0.43	0.49	-0.30	-0.43	-0.49
	126	-0.36	0.19	0.22	0.26	-0.19	-0.22	-0.26
	Momentum	5	-48.71	2.13	12.05	13.96	-2.13	-8.86
Contrarian	21	-43.46	2.33	12.42	13.67	-2.33	-9.91	-13.67
	63	-37.78	2.45	11.60	12.63	-2.45	-9.75	-12.63
	126	-16.33	1.10	4.83	4.58	-1.10	-1.97	-4.58
	Moving Average	5	-51.04	1.76	10.91	14.76	-1.76	-10.91
21		-43.88	2.28	11.23	15.47	-2.28	-11.23	-14.48
63		-37.98	2.57	10.72	14.86	-2.57	-10.72	-14.50
126		-33.16	2.58	9.94	13.52	-2.58	-9.94	-13.37

Table 4: Table presenting the average profits of traders. The first column contains the average yearly profits of each trader type over all simulations in which they are present. The remaining columns detail the effect on the average profits of the row trader of the presence of traders of the column type. Traders are grouped by their rule type. Results are calculated for markets using lagged price information.

		Trend Follower				Contrarian		
		Mean	Breakout	Momentum	Moving Average	Breakout	Momentum	Moving Average
Fundamental	-	4.95	0.34	0.58	0.73	-0.32	-0.56	-0.71
Breakout	5	-2.72	-0.04	-0.17	-0.18	0.12	0.16	0.17
	21	-0.84	-0.02	-0.03	-0.02	0.05	0.03	0.02
	63	-0.26	-0.00	0.00	0.01	0.02	-0.00	-0.01
	126	-0.11	-0.00	0.01	0.02	0.01	-0.01	-0.02
Momentum	5	-3.70	-0.14	-0.14	-0.24	0.12	0.36	0.22
	21	-2.77	-0.10	-0.12	-0.16	0.08	0.27	0.15
	63	-1.88	-0.05	-0.07	-0.09	0.05	0.16	0.08
	126	-1.40	-0.04	-0.05	-0.05	0.03	0.12	0.04
Moving Average	5	-5.39	-0.18	-0.39	-0.38	0.16	0.37	0.59
	21	-4.39	-0.18	-0.31	-0.35	0.16	0.29	0.45
	63	-3.10	-0.10	-0.18	-0.21	0.09	0.17	0.27
	126	-2.36	-0.06	-0.12	-0.13	0.05	0.11	0.18
Breakout	5	1.93	0.14	0.27	0.34	-0.05	-0.25	-0.33
Contrarian	21	0.59	0.07	0.12	0.15	-0.06	-0.11	-0.14
	63	0.16	0.04	0.06	0.08	-0.04	-0.06	-0.08
	126	0.04	0.04	0.05	0.06	-0.04	-0.05	-0.06
	Momentum	5	2.17	0.13	0.29	0.45	-0.12	-0.24
Contraian	21	1.74	0.11	0.26	0.38	-0.10	-0.24	-0.37
	63	1.25	0.08	0.20	0.30	-0.07	-0.21	-0.29
	126	0.91	0.07	0.17	0.24	-0.06	-0.18	-0.23
	Moving Average	5	3.42	0.17	0.42	0.55	-0.15	-0.40
21		2.93	0.16	0.39	0.53	-0.15	-0.38	-0.49
63		2.18	0.11	0.29	0.42	-0.10	-0.28	-0.42
126		1.69	0.08	0.24	0.35	-0.07	-0.23	-0.37

Table 5: Table presenting the average profits of traders. The first column contains the average annual profits of each trader type over all simulations in which they are present. The remaining columns detail the effect on the average profit of the row trader of the presence of traders of the column type. Traders are grouped by their rule type. Results are calculated for markets using up to date price information.

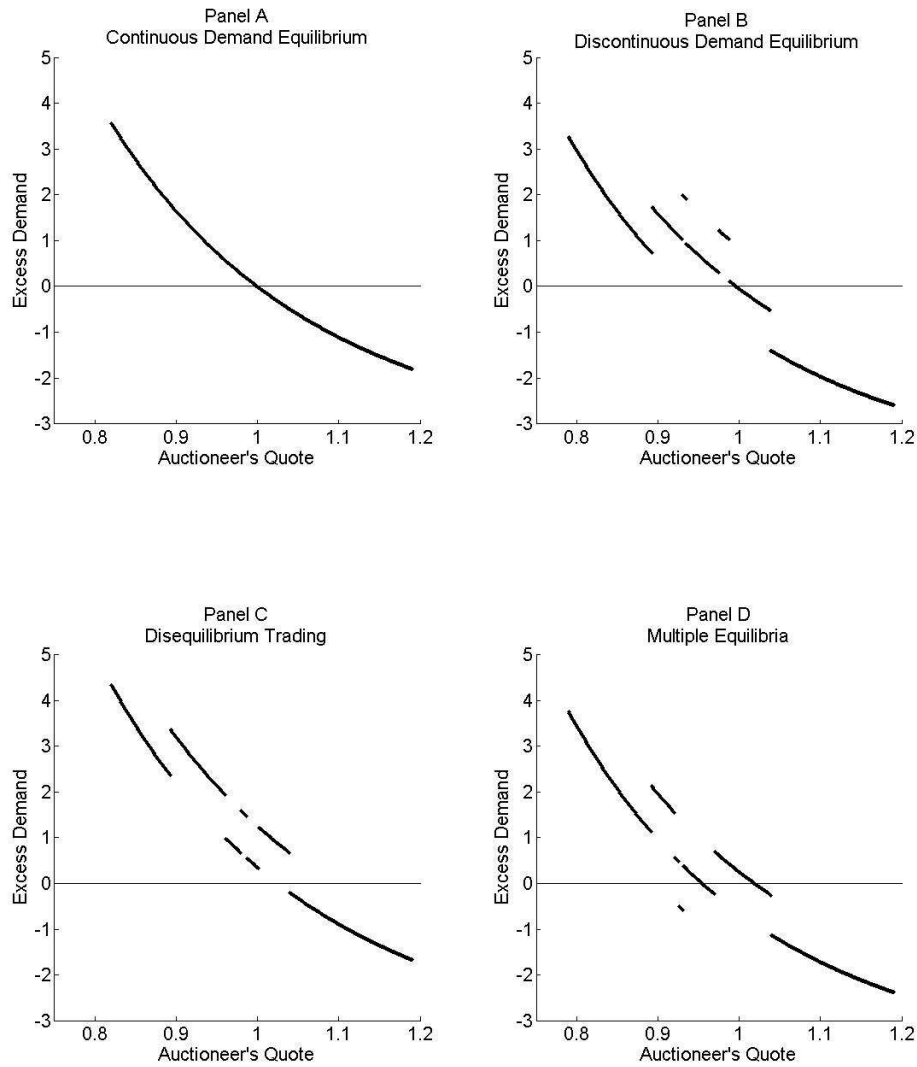


Figure 1: Market Clearing Scenarios. When technical traders condition their rules on lagged prices, aggregate demand is smooth, yielding a unique market clearing price (panel A). But when trading rule signals are synchronized with the auctioneer's quote, discontinuities in individual demand curves lead to discontinuities in aggregate demand. A unique clearing price is still possible (panel B), but there is also the possibility of disequilibrium trading (panel C) and multiple equilibria (panel D).




Figure 2: Histograms of average annual profits for traders using lagged information (Left) and real time information (Right). Top: Breakout rules, Middle: Momentum rules and Bottom: Moving average rules. In each case the dashed line is for fundamentalists, the solid line is for trend followers and the dotted line is for contrarians.