Non-Fundamental Volatility in Financial Markets

Keisuke Teeple*

Job Market Paper

December 29, 2021

Abstract

I use coarse Bayesian updating to explain three stylized facts: trading anomalies around major stock market milestones, excess price volatility and trade volume, and heavy-tailed prices. When expectations are coarse and traders are heterogeneous, traders make heterogeneous mistakes. This disagreement generates substantial trade volume and I show that the ensuing price discovery process converges in distribution to an empirically relevant class of Lévy processes. Notably, I obtain this result without any exogenous shocks or processes; the only exogenous random variable is a heterogeneity parameter drawn at time zero. I then establish a perfectly non-revealing equilibrium: because traders are different, their aggregated trades generate volatile prices, yet this volatility impedes traders from learning their differences.

JEL Classification: G14; G41; D84; D91.

Keywords: Non-fundamental volatility; Behavioral finance; Coarse Bayesian updating; Price barriers

*kwteeple@ucdavis.edu, University of California Davis, One Shields Ave Davis, California 95616

1 Introduction

In his book *Irrational Exuberance*, Shiller (2015) offers a compelling yet underexplored explanation for non-fundamental price volatility, which he calls quantitative anchors. He describes these as "levels of the market that some people use as indications of whether the market is over- or underpriced and whether it is a good time to buy." As possible anchors, he suggests the nearest milestone of a prominent index such as the Dow or the nearest round number. I formalize this idea by assuming that traders use coarse Bayesian updating. I then show that my model explains three stylized facts: trading anomalies observed near major stock market milestones, excess price volatility and trade volume, and heavy-tailed price distributions. Notably, I derive my results without assuming any exogenous information arrival or i.i.d. shocks; when traders trade based on heterogeneous expectation errors, it is the ensuing *price discovery* process which generates volatility.

Coarse Bayesians select the most likely posterior from an exogenous set of posteriors. These possible posteriors can be thought of as "competing theories of the world," from which traders select the one closest to Bayes' Rule. The interpretation is of unmodeled cognitive costs of attention. In the simplest possible example, there is a continuum of prices for a single financial asset, yet a coarse Bayesian can only form expectations over a low or high price. The trader classifies all prices to the left of some cutoff with the low expected price, and similarly all prices to the right of the cutoff with the high expected price. These two "regimes" should be thought of as a bull and bear market. Heterogeneous coarse Bayesians may not agree on where the bear market ends and the bull market begins. And this is precisely the motive for trade: sellers and buyers must, by construction, observe the same price today. But a seller believes she is in a bear market and hence that the asset is *overpriced*, which is why she sells. A buyer believes he is in a bull market and hence that the asset is *underpriced*, which is why he buys. Because there is no exogenous news or fundamental value, this is a model of technical trading: traders view prices today and extrapolate prices tomorrow (with error due to their coarse beliefs). These are technical traders who use rules of thumb.

The final participant in my economy is a market maker who attempts to clear markets by raising price when there is excess demand and lowering price when there is excess supply. Notice this is a model of disequilibrium; per period, markets may not clear and the market maker would have to absorb these excess positions. It is this trial-and-error, or price discovery, process which I show to be volatile. In Figure 1, I compare the price discovery process in a calibrated economy to the S&P 500 for a one-year period.

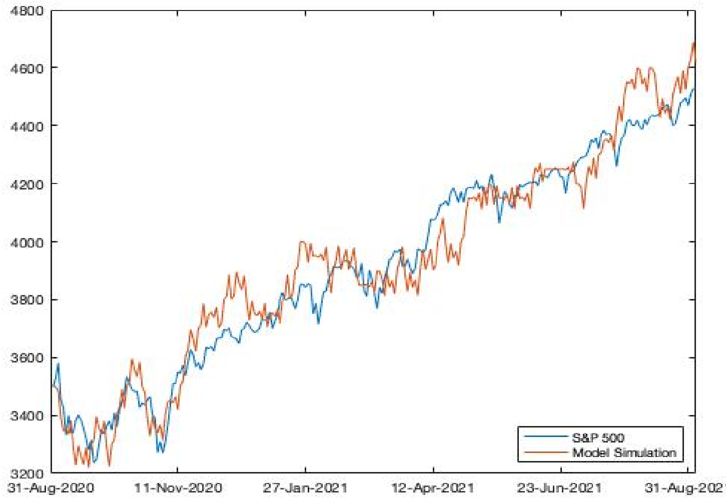

Figure 1: S&P 500 versus Model Simulation

Instead of conventional models where exogenous information arrival generates volatile prices, here traders disagree due to their coarse beliefs, this disagreement leads to large trade volumes, and the ensuing price discovery process is shown to be volatile. In my main theoretical result, I show that this price discovery process converges in distribution to a Brownian motion as coarseness tends to zero (and the standard Bayesian case is recovered). The following quote from Bagehot (1971) provides the perfect summary of my model's microstructure:

> It is well known that market makers of all kinds make surprisingly little use of fundamental information. Instead they observe the relative pressure of buy and sell orders and attempt to find a price that equilibrates these pressures. The resulting market price at any point in time is not merely a consensus of the transactors in the marketplace, it is also a consensus of their mistakes. Under the heading of mistakes we may include errors in computation, errors of judgment, factual oversights and errors in the logic of analysis.

Even with substantial trade volume generated by coarse expectations, prices may still be non-random due to laws of large numbers, which may be invoked when there are many traders.[1] However, the market maker considered here does not scale prices using a factor $\frac{1}{n}$ as in standard models (where n denotes the number of traders), and this causes prices to converge not to a steady state, but in distribution. I show that this alternative market maker can be established from a reasonable set of assumptions. The combination of coarse expectations and this non-Walrasian market maker results in a volatile price discovery process.

Trading anomalies around stock market milestones are studied empirically by Donaldson and Kim (1993). The authors find that the Dow Jones Industrial Average tends not to cross increments of one hundred; the index faces "resistance" levels from above and "support" levels from below. Having broken through an increment of one hundred, or what they call a price barrier, the Dow Jones Industrial Average then moves more than otherwise warranted. Price barriers demarcate market regimes as defined in Day and Huang (1990), for example bull and bear markets. Returning to my simple example with either a low or high expected price, the cutoff between these two regimes acts precisely as a price barrier. As prices approach the cutoff from below, traders believe the asset is overpriced in a bear market; the resulting selling puts downward pressure on prices. As prices approach the cutoff from above, traders believe the asset is underpriced in a bull market; the resulting buying puts upward pressure on prices. In this way, price barriers arise endogenously in my setting.

In the main extension to my model, I match empirically relevant trade volume and price return moments. While the lack of trading volume generated in standard rational expectations models is well-known (Campbell (2017)), my model generates power law trade volume. Notably, I do not assume power law shocks; the result holds for a broad family of exogenous random variables. Instead, my (necessary and) sufficient assumptions relate to trader preferences: when they trade aggressively based on perceived arbitrage opportunities, trade volume converges *in distribution* to a power law. Figure 2 compares trade volume predicted by my model to the S&P 500 daily volume for a three year period ending on August 31, 2021.[2]

It is also well-known that empirical price returns do not conform to a normal distribution (Hoechstoetter et al. (2005)); instead, returns have heavy tails. Using the same

[1]For an example of this, see Aiyagari (1994).

[2]The power law coefficient in this calibration is $\alpha = 1.9$.

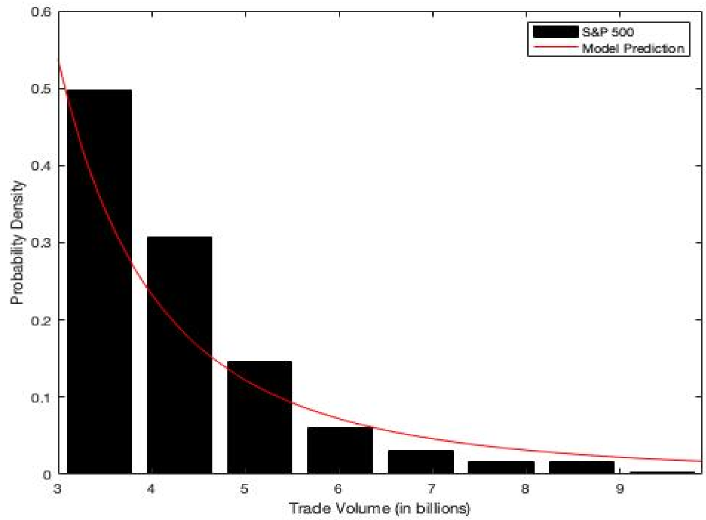

Figure 2: Trade Volume Histogram

calibration as in Figure 2, in Figure 3 I plot stable-α returns predicted by my model versus the S&P 500 daily returns for the same three year period. Like with large trade volumes, aggressive trades based on perceived arbitrage opportunities lead to large price movements (i.e. heavy tails). And like with trade volume, I do not assume heavy-tailed prices; I show that prices converge in distribution to heavy-tailed prices for a broad family of exogenous random variables.

I establish an equilibrium in the following way. Traders myopically believe prices are volatile; at first this is a strange belief because there are no exogenous shocks. However, based on their coarse beliefs, I show that traders trade. Via my market maker, trade is translated into price movements and I show that prices converge in distribution precisely to the original trader belief. So traders believe prices are volatile, trade based on this belief, and generate this volatility in reality.

The only exogenous random variable in my model is a heterogeneity parameter drawn at time zero. Because this parameter is realized at time zero, I must consider whether traders can learn each other's demands (and hence learn the future price path) from the information contained in prices. Prices are not invertible for the following reason. Demand functions considered here are not monotone: for small price changes, asset demand functions are downward sloping for the usual reasons. However, for large price changes,

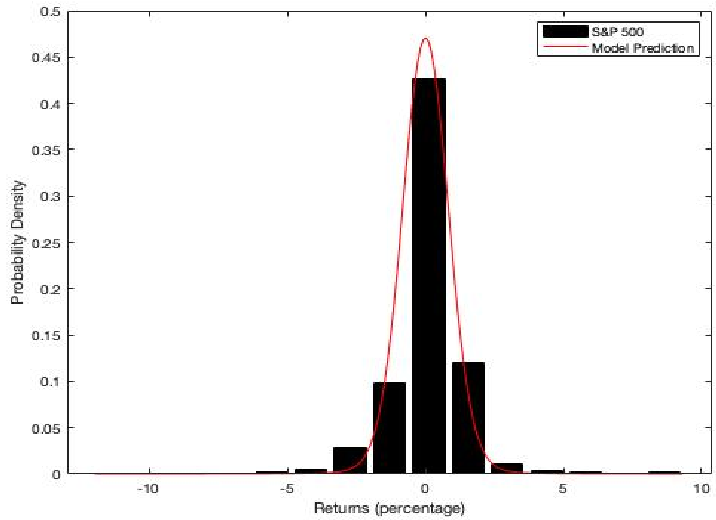

Figure 3: Price Returns Histogram

traders use higher prices today as indications of higher prices tomorrow and discretely shift their expectations upward. Demand functions that are not mathematically one-to-one lead to a pricing function that is not one-to-one; this leads to non-revealing prices in the sense of Radner (1979). That is, there are multiple demands that could have generated the same price. I call the solution concept a perfectly non-revealing equilibrium: it is the disagreement that generates volatile prices, and it is the volatility that impedes traders from learning their differences. Even though all random variables have been realized at time zero, they are not *revealed* through information contained in prices and hence traders are unable to predict future prices. This approach presents a compelling alternative to i.i.d. shocks over time assumed in many models today.

Literature

Leading explanations for fundamental asset price volatility *amplify* exogenous risk. For example, the habit formation model (Campbell and Cochrane (1999)) uses utility curvature to amplify shocks, and the long-run risks model (Bansal and Yaron (2004)) uses persistent shocks to alter long-run expectations. The disaster risks literature (Barro (2006)) assumes a jump process on top of the standard Brownian motion. In contrast, here I will not assume any exogenous driving process. Competing explanations for non-fundamental volatility identify asset *mispricing*. For example, overconfident traders over-react to news (Daniel et al. (1998)); price movements are amplified with herd behavior and information externalities (Park and Sabourian (2011)); and bubbles can persist when there are coordination issues (Abreu and Brunnermeier (2003)). In contrast to mispricing at a fixed moment in time, here I will consider volatility defined as the second moment of a stochastic process.

My model addresses concerns about the lack of trading volume generated in standard rational expectations models compared to what is empirically observed (Campbell (2017)). Cochrane (2007) went so far as to suggest that the next revolution in asset pricing will consist of models that can explain empirically observed levels and patterns of trading volume. While this debate goes back to Milgrom and Stokey (1982) and Varian (1989), a more recent paper by Alvarez and Atkeson (2021) hits empirically relevant volume moments using shocks to risk aversion. Here, heterogeneous expectation errors are what generate substantial trade volume.

Coarse Bayesian updating, which nests standard Bayesian updating as a special case, is axiomatized by Jakobsen (2021) and is related to categorical thinking from Mullainathan (2002) and hypothesis testing from Ortoleva (2012).[3] The original intent of this literature was to address a growing set of experimental findings refuting the standard Bayesian model. The set of posteriors over which an agent must optimally select (optimal in the sense of minimal deviation from Bayes' Rule) induces categories, or competing theories of the world; in my setting these will exactly correspond to market regimes. Mullainathan (2002) points out that this type of behavior naturally creates under- and over-reactions to information. That is, some data may not be "drastic enough" to induce a change in the posterior category (under-reactions); however, other data that induces a change in category may move the posterior more than in the standard Bayesian case (over-reactions). I will show that coarse Bayesian updating is capable of generating momentum and reversals over the asset price time series as well.

The idea of learning from information revealed in prices was formally introduced by Radner (1979). While fully and partially revealing equilibria have been studied by many subsequent authors, the current paper investigates perfectly non-revealing equilibria. For an example of a partially revealing equilibrium, see Allen (1985) who assumes that price

[3]Behaviors similar in spirit to coarse Bayesian updating have been applied to financial market settings by Eyster and Piccione (2013), Gul et al. (2017), and Steiner and Stewart (2015). These authors, however, do not study non-fundamental volatility.

observations are noisy. Polemarchakis and Siconolfi (1993) study a fully non-revealing equilibrium, which obtains from indeterminacy of equilibrium prices. Here, prices are perfectly non-revealing without any reliance on additional noise or multiplicity.

An endogenous Brownian motion price process has, in itself, drawn considerable attention. One literature has derived a Brownian motion price process from rational traders who are subject to an exogenous Brownian motion. Pakkanen (2010) derives Brownian prices from noise in the demand function, and Raimondo (2005) derives geometric Brownian prices by assuming endowments and asset returns follow a geometric Brownian motion. Deriving a Brownian motion by assuming yet another Brownian motion is unsatisfactory to some degree, and is not the goal of the present exposition. A second literature derives the Brownian motion by deviating from the utility maximization paradigm. Horst and Rothe (2008) solve for a Brownian motion by assuming a Poisson process for order arrivals and agent switching, and Cox et al. (1979) find a Brownian motion as a limit of simple coin flips. In my model, traders are fully rational, aside from their coarse expectations.

The remainder of the paper is organized as follows. In Section 2, I provide a simplified version of the model defining coarse expectations, deriving non-monotone demands, and plotting price paths that resemble a random walk for some parameter values and price barriers for others. In Section 3, I lay out the assumptions required for my equilibrium. In Section 4, I show that prices are non-revealing and analytically prove that prices converge in distribution to a Brownian motion, establishing the desired equilibrium. In Section 5, I extend the result to heavy-tailed distributions in equilibrium and discuss policy. Section 6 concludes.

2 Simplified Model

Random Walk

Time is discrete and infinite $t = 0, 1, 2, ...$ There are n traders who observe today's price p_t of a single durable asset that pays no dividends, then form expectations about tomorrow's price p_{t+1}. Because there is no exogenous information arrival in this model, this could be viewed as a form of technical trading. Preferences are mean-variance over wealth w_{t+1}:

$$E[u(w_{t+1})|p_t] = E\left[w_{t+1} - \frac{\rho}{2} w_{t+1}^2 \middle| p_t \right]$$

subject to wealth being derived from capital gains:

$$w_{t+1} = (p_{t+1} - p_t)x_t$$

where holdings of the asset are given by x_t and risk aversion is given by ρ. The constraint is kept intentionally simple so as to only capture the speculative motive for trade.[4] Note that I have implicitly assumed that traders have access to an unlimited line of credit at zero interest rates. Traders guess (and must later verify in equilibrium) that prices are a unit-variance martingale $E[p_{t+1}|p_t] = p_t$, so there is no expected gain or loss. For now, this is simply taken as a form of myopia. Under this guess, the demand function takes the following form:

$$x(p_t) = \frac{E[p_{t+1}|p_t] - p_t}{\rho} \tag{1}$$

and there would be no trade in my model due to the martingale guess, $E[p_{t+1}|p_t] = p_t$. Instead, I deviate from this standard setup and assume that traders form expectations using coarse Bayesian updating. Consistent with the axiomatization of coarse Bayesian updating in an abstract setting (see Jakobsen (2021)), I assume posteriors are chosen from an exogenous, fixed set $\mathcal{P}$. The interpretation is that traders face unmodeled cognitive costs of attention beyond this coarse set of posteriors $\mathcal{P}$.

In my application to financial markets, I assume that this set $\mathcal{P}$ consists of all distributions with a mean in the set $\{\ell\varepsilon\}_{\ell\in\mathbb{Z}}$, where $\mathbb{Z}$ denotes the set of all integers. That is, expectations are restricted to an equally-spaced grid with spacing ε. The equal spacing assumption is made in the spirit of empirical applications, such as whole numbers (Shiller (2015)) and price barriers (Donaldson and Kim (1993)). If $\varepsilon = 1$, traders think in terms of integers, and if $\varepsilon = 100$, traders think in terms of hundreds. As $\varepsilon \to 0$, the trader becomes Bayesian and considers the entire number line. The coarse expectation is then

[4]While this trader may seem myopic in the sense that she does not look beyond period $(t+1)$, this is without loss of generality in such a simple setup. This is a problem of speculation, not savings.

defined:

$$
\begin{aligned}
\tilde{E}[p_{t+1}|p_t] =& \text{argmin}_p |p - E[p_{t+1}|p_t]| \;\text{ s.t. } p \in \{\ell\varepsilon\}_{\ell \in \mathbb{Z}} \\
=& \text{argmin}_p |p - p_t| \;\text{ s.t. } p \in \{\ell\varepsilon\}_{\ell \in \mathbb{Z}}
\end{aligned}
\tag{2}
$$

The first line of the formula above says that traders pick the closest posterior mean to the true expected value of prices, $E[p_{t+1}|p_t]$. The second line of the formula follows from the martingale guess. Altogether, (2) reads: traders believe prices are *as close as possible* to a martingale. The interpretation is that traders use simple "rules of thumb." For example, even if the price of milk is a martingale, shoppers may quote $2.00 instead of the actual price paid today of $1.89. Note that (2) is a not a restriction on the *support* of the posterior, only on its *mean*. Hence traders are not surprised when they see a price p_{t+1} outside their grid; this is simply viewed as a realization not equal to the mean. In Figure 4, I have combined equations (1) and (2) and plotted the demand function with $\varepsilon = \rho = 1$.

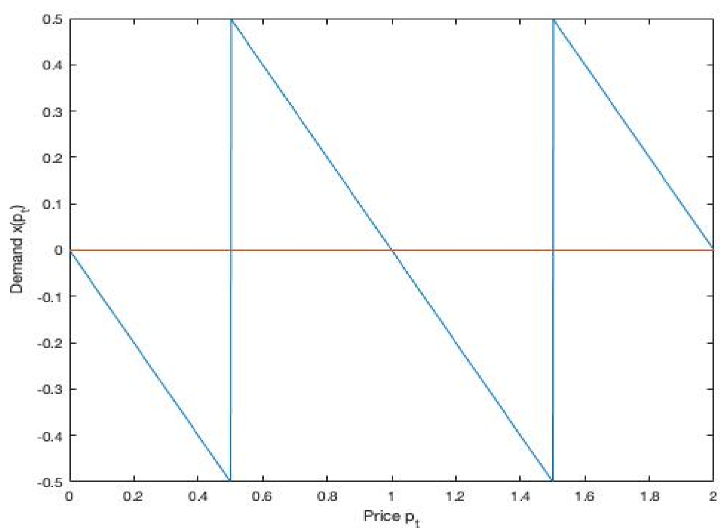

Figure 4: Demand Function

Demand curves slope downwards for the usual reasons. However, at certain prices, the trader uses higher prices today as a signal for higher prices tomorrow and shifts her demand upwards. The fact that she uses prices today to predict prices tomorrow follows from the martingale guess. And demand shifts discretely because expectations (the demand intercept) shift discretely. Each downward-sloping interval should be interpreted as a regime, for example a bull and bear market. Figure 4 captures the notion of a trigger strategy. For example, when prices fall below the threshold $p_t = 0.5$, investors liquidate their positions, commonly known as the stop-loss order (Lei and Li (2009)).

As noted by Mullainathan (2002), coarse Bayesians can both under-and over-react to news. Here, the news is simply the current price realization. Say $p_{t-1} = 1$ in Figure 4, so the prior is $E[p_{t+1}|p_{t-1}] = 1$. For modest price changes, for example $p_t = 1.4$, the posterior remains unchanged, $E[p_{t+1}|p_t] = 1$. This represents an under-reaction with respect to a standard Bayesian. For larger price changes, for example $\tilde{p}_t = 1.6$, the posterior moves more than that of a Bayesian, $E[p_{t+1}|\tilde{p}_t] = 2$. This represents an over-reaction. There are some trades that lead to price reversals (say prices rise to $p_t = 1.4$, which induces selling), and others that induce price momentum (say prices rise to $\tilde{p}_t = 1.6$, which induces yet more buying).

My economy is populated by n coarse Bayesians indexed by j, who are heterogeneous based on one parameter. I capture this heterogeneity in a particularly tractable way: with a scalar shift parameter r^j which shifts their demands horizontally, and hence shifts regimes horizontally. Traders may disagree on where the bear market ends and the bull market begins. For this reason, I call this heterogeneity parameter the *reference point.* More formally, I assume that traders have heterogeneous posteriors $\mathcal{P}^j$, where $\mathcal{P}^j$ consists of all distributions with a mean in the set $\{r^j + \ell\varepsilon\}_{\ell\in\mathbb{Z}}$. The coarse expectation equation (2) is then updated to:

$$\tilde{E}_j[p_{t+1}|p_t] = \text{argmin}_p|p - p_t| \;\; \text{s.t.} \;\; p \in \{r^j + \ell\varepsilon\}_{\ell\in\mathbb{Z}} \tag{3}$$

I assume that r^j is drawn independently across traders at time $t = 0$ from a standard normal distribution; so traders tend to agree, but not entirely, on the location of regimes.[5] Notice that r^j is a heterogeneity parameter drawn at time zero, not a per-period shock. Also note that equations (1) and (3) together define a demand function based on coarse expectations, which I denote $x(p_t, r^j)$.

Finally, instead of implicitly assuming market clearing prices each period, I explicitly model a market maker who attempts to clear markets. She observes prices today p_t and aggregate demand $z_n(p_t, r) = \sum_j x(p_t, r^j)$ and chooses p_{t+1} to maximize the following

[5]Heterogeneous spacing ε^j is covered both later in this example and in Section 5.

objective:[6]

$$H\left(p_{t+1}-p_t, z_n(p_t,r)\right) = c(p_{t+1}-p_t)z_n(p_t,r) - \frac{1}{2}(p_{t+1}-p_t)^2$$

In the first term, the market maker increases (decreases) price when there is positive (negative) excess demand. In the second term, the market maker adjusts prices as little as possible. The parameter $c > 0$ captures the relative weight the market maker places on correctly signing (versus minimizing) price adjustments. Notice that when aggregate demand is zero $z_n(p_t, r) = 0$, the market maker settles prices at the steady state.[7] The first-order condition yields a simple linear rule:[8]

$$p_{t+1} = p_t + c\sum_j x(p_t, r^j) \tag{4}$$

This market maker explicitly carries out the price discovery process, and the ensuing dynamics are reminiscent of the classical tâtonnement (trial-and-error) process. Per period, markets may not clear and the market maker would need to absorb these extra positions; in this sense, this is a model of disequilibrium. Traders trade, the market maker adjusts price accordingly, and the process repeats. With a starting price p_0, (1), (3), and (4) describe a dynamical system. Figure 5 shows the first 500 iterations of the nonlinear map (1), (3), and (4), where parameters have been chosen as $p_0 = 1000$, $n = 1000$, and $c = 0.1$. The three price paths correspond to different draws of r^j, holding all other parameters fixed.

The price paths follow something that resembles a random walk to the naked eye, from nothing more than heterogeneous traders with coarse expectations. This volatile price discovery process is not driven by any exogenous shocks, but instead by trade volume that stems from trader disagreement on whether the asset is under- or over-priced. Some traders buy, thinking that the asset is under-priced in a bull market; others sell, thinking that the asset is over-priced in a bear market.

Price Barriers

[6] r denotes the vector $[r^1, ..., r^n]$.

[7] One benefit of this market maker over one who minimizes the distance between supply and demand is a lower informational requirement. This market maker only needs to observe the *value* of aggregate demand, not its functional form.

[8] The model of Kyle (1985) employs a similar linear market maker.

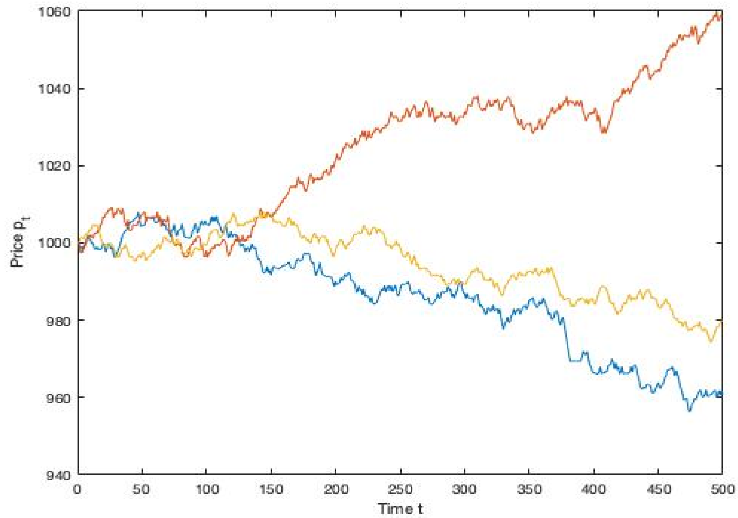

Figure 5: Price Paths

With small $\varepsilon = 1$ as in Figure 5, traders are close to being Bayesian. With large $\varepsilon = 100$ as in Figure 6, traders are far from Bayesian. Expectations can take on only a restrictive set of values, and as in the price barriers literature, this leads to prices that tend to stay above "support" levels and below "resistance" levels at increments of 100. The intervals between price barriers, for example $p_t \in [1000, 1100]$ in Figure 6, correspond one-to-one with market regimes, which were the continuous portions of demand in Figure 4. As prices near the upper (lower) limit of a regime, selling (buying) pressure tends to keep prices within that regime. In Figure 6, the distribution over r^j has been adjusted to $N(50, 625)$, and parameters to $\varepsilon = 100$, $p_0 = 1000$, $n = 15$, $c = 1.5$, and $\rho = 5$. If I were to send the variance of the distribution over r^j to infinity, the price path begins to resemble a random walk. It is the *agreement* between traders, captured by the relatively small standard deviation of this distribution (standard deviation is 25), that makes price barriers so pronounced in Figure 6.

Albeit only informally so far, I have nested two extreme cases within the Motivating Example. As the grid size ε gets smaller, prices begin to resemble a random walk. When the grid ε is calibrated to a large number, price dynamics match those described in the price barriers literature. Reality likely consists of a combination of these two cases. In fact, to match the S&P data in Figure 1, the model was calibrated with 85% of the pop-

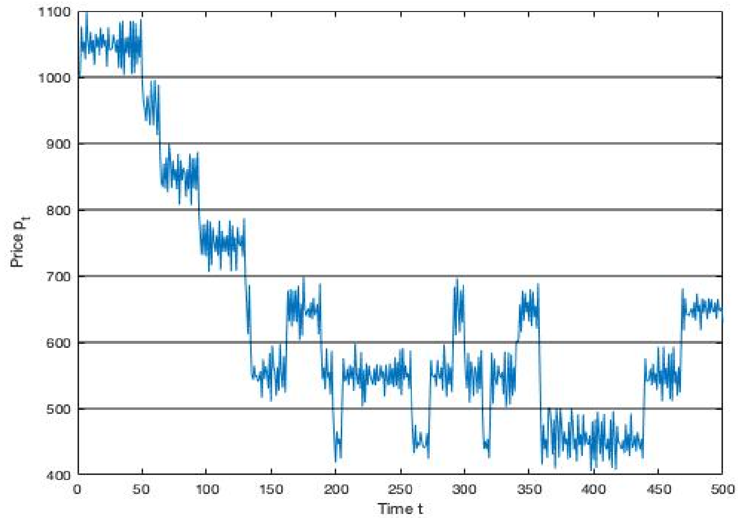

Figure 6: Price Barriers

ulation thinking in terms of integers ($\varepsilon = 1$), and the remaining 15% thinking in terms of hundreds ($\varepsilon = 100$). Parameters were kept the same as each Motivating Example, except $p_0 = 3500.31$, $n = 1000$, and $c = 0.1$.[9]

In the remainder of the paper, I will justify assumptions made in this example − an important one is that traders believe prices are a martingale − and discuss in what sense prices converge to a random walk.

3 Model Assumptions

In this section, I state my assumptions; some generalize the previous setup, while others restrict it. As a starting point, I consider second-order Taylor series expansions of both a generalized market maker objective $H_n(p_{t+1} - p_t, z_n(p_t, r))$ and trader utility function $\tilde{E}_j[u(w_{t+1})|p_t]$. For now, I continue to treat the martingale guess as a form of myopia; I will return to this issue in Section 4, where I establish my equilibrium.

Market Maker

[9] To select the proportion 85/15, reference points were drawn 2000 times for each proportion, and the closest price path (in terms of ℓ^1) to the S&P was selected. Hence the claim is that there exists a set of reference points consistent with the S&P, not that my model predicts future prices out-of-sample.

The Taylor series expansion yields a linear market maker rule:

$$p_{t+1} = a_n z_n(p_t, r) + b_n p_t + c_n$$

where, as before, $z_n(p_t, r)$ denotes aggregate demand. This starting point is for nothing more than tractability. There are likely many interesting non-linear market maker rules; however, the analysis here will focus on the linear case. I allow constants a_n, b_n, and c_n to depend on the number of traders n. My first two assumptions place restrictions on these three constants.

A1: a_n, b_n, and c_n are selected so $p_{t+1} \to p_t$ in probability $\iff \lim_{n\to\infty} z_n(p_t, r) = 0$.
A2: For fixed p_t, p_{t+1} converges in distribution as $n \to \infty$.

A1 is a necessary and sufficient condition for price convergence, in line with the market microstructure literature (Farmer and Joshi (2002), Kyle (1985)). If A1 did not hold, the market maker might have to absorb nonzero excess demand positions into perpetuity. Notice that the sufficient direction ($\Leftarrow$) of A1 places the following restrictions on the market maker rule as $n \to \infty$:

$$p_{t+1} = p_t + a_n z_n(p_t, r) \tag{5}$$

That is, b_n tends to one and c_n tends to zero. The weaker of the two assumptions is A2. Prices must converge in the weakest possible sense; stated another way, prices do not diverge. While weak, A2 does impose a scaling property on the market maker's constant a_n: as there are more traders, prices respond less sensitively to each. One scaling that may immediately come to mind is a Walrasian market maker $a_n = \frac{1}{n}$.[10] Because excess demands are centered around zero, I could then invoke the weak law of large numbers and claim that $p_{t+1} \to p_t$ in probability. But the proposed rule of flat prices, irrespective of the the realization of excess demand, is nonsensical and violates the necessary direction ($\Rightarrow$) of A1. In my first lemma, I show that, without loss of generality, I can consider the following market maker rule:

$$p_{t+1} = p_t + \frac{c z_n(p_t, r)}{\sqrt{n}} \tag{6}$$

[10] Many notable models implicitly assume the Walrasian market maker when they normalize the mass of agents to one. For example, see Aiyagari (1994).

where c is a constant that does not depend on n. This rule will allow me to invoke central limit theorems, instead of laws of large numbers as in many standard models (see Aiyagari (1994)). Prices will then converge in *distribution*, not in probability. Intuitively, this non-Walrasian market maker (6) allows prices to move more freely than a Walrasian one. If, instead of A1, I had considered A1′, where $p_{t+1} \to p_t$ in probability $\iff \lim_{n\to\infty} \frac{1}{n} z_n(p_t, r) = 0$, then the Walrasian market maker would have been a viable candidate. I argue that price convergence when there is zero aggregate demand (A1) is more reasonable than when there is zero demand per capita (A1′).

Lemma 1 (Market Maker): For any linear market maker rule that satisfies A1-A2, there is an alternative rule described by (6) that achieves the same resulting price distribution as $n \to \infty$.

Proof: See Appendix A.

Unlike a problem where the normalization constant is given and one might make use of standard sufficient conditions for invoking the central limit theorem, here convergence was assumed and the normalization constant was derived using lesser known necessary conditions for the central limit theorem. Recall the interpretation of (6) given in Section 2: it is the first-order condition of a market maker objective. This objective had two components: on one hand the market maker had an incentive to correctly sign price changes, and on the other she had an incentive to minimize price movements. Equation (6) indeed correctly signs price changes: it maps excess demand into higher prices and excess supply into lower prices. Furthermore, (6) scales the sensitivity of the response based on the number of traders; price movements are dampened as the number of traders increases.

Traders

A second-order Taylor series expansion of trader utility yields the demand function (1) from Section 2; hence I do not generalize in this direction. My final two assumptions are imposed on risk aversion ρ and the distribution of reference points r^j.[11]

A3: Traders are risk-neutral in the sense that $\rho \sim \varepsilon$.

[11] $f(\varepsilon)$ is on the order of ε, $f(\varepsilon) \sim \varepsilon$, if $\lim_{\varepsilon\to 0} \frac{f(\varepsilon)}{\varepsilon} = 1$.

A4: r^j are i.i.d. across traders with a tight and absolutely continuous distribution.

As previously discussed, small ε represents the standard Bayesian benchmark. In A3, I assume that $\rho \sim \varepsilon$ and therefore $\lim_{\varepsilon \to 0} \rho = 0$; traders are risk-neutral. The intent of sending ε to zero is to recover the case of the Bayesian trader. But because traders believe prices are a martingale, risk-averse traders with correct expectations will choose not to trade. Risk-neutrality is an unintended but understandable condition to avoid this no-trade result.

A4 is a strict generalization of the Motivating Example; the standard normal imposed in the example was entirely unnecessary. Importantly, these reference points are i.i.d. across traders, but not over time. These traders are born different and stay different. A distribution F is tight if for all $\delta > 0$ there exists M such that $F(M) - F(-M) \geq 1 - \delta$. In words, probability mass cannot escape to infinity. And finally, an absolutely continuous distribution has a density.[12] Tightness and absolute continuity are not without loss of generality, but the family of distributions considered under A4 is very broad.

To summarize, these four assumptions A1-A4 include two on the market maker and two on traders. Because A2 and A4 are extremely weak, A1 and A3 should be thought of as the assumptions with the most "bite." A3 ensures nonzero trade when $\varepsilon \to 0$, and A1 ensures prices are a nondegenerate random variable when $n \to \infty$.

My upcoming results will apply more generally than to linear demands (which follow from second-order Taylor expansions of utility). These Taylor expansions were introduced only for clarity of exposition. In Appendix B, I introduce three additional assumptions, which generalize linear demands at the expense of being more technical, for which my results hold. Put differently, the second-order Taylor series expansion of utility is one special case of the additional assumptions A5-A7. Upcoming proofs are written for the more general case.

[12] This rules out degenerate r^j.

4 Equilibrium

In this section, I establish an equilibrium in the following sense. Traders believe prices are drawn from some distribution (because there are no shocks, this is a strange belief!). Based on their coarse beliefs, I have shown they trade; and based on the market maker rule, prices move. I need to show that the distribution to which prices converge is precisely the one from traders' initial beliefs. In short, traders trade based on volatile price beliefs, and it is their trades that generate this volatility in reality. The main result in this section will be that prices converge in distribution to a discrete time Brownian motion.

Incomplete Information

Before establishing my equilibrium, I must formally address trader beliefs and learning. The traders in this setting are different, and I assume that each knows only her own r^j; hence this is a setting with incomplete information. Asset returns, at least in the first period, are random variables because they are a function of r^j. Each subsequent period, I allow traders to observe that period's price p_t (and only this information), and allow them to condition their demands on this realization. The question then becomes: What do traders learn from the information revealed in prices? Put another way: Are traders able to learn from disequilibrium? As intuition suggests, with fixed n and ε, traders would successfully and fully learn each other's information in finite time. Instead of modelling this learning process, I consider conditions under which learning fails entirely (not as an assumption, but as a result).

While r^j differ from information as defined in standard models, there is nonetheless an incentive for traders to learn the random vector $r = [r^1, ..., r^n]$. Say one sophisticated trader learns the entire vector r, while all others learn nothing. The sophisticated trader can infer the aggregate demand function, and hence the future price path, leading to arbitrage opportunities even with coarse expectations.

The next lemma provides conditions under which prices will never reveal trader information. Effectively, I am slowly lifting the trader's myopia, but considering conditions under which their beliefs coincide with the myopic case. When I make claims about non-revealing prices, the result is not a product of coarse beliefs; they hold for any Bayesian.[13]

Lemma 2 (Non-Revealing Prices): Given p_t, a linear market maker, and assuming A1-A5, price increments $(p_{t+1} - p_t)$ and references points r^j are independent random variables as $n \to \infty$ or as $\varepsilon \to 0$.[14]

Proof: See Appendix C.

I allow traders to condition trades based on the realization p_t, and I want to know if they can infer any information about r^j. If they are independent random variables as stated in Lemma 2, the answer must be that traders infer nothing. Notice that the lemma is a very strong statement: traders learn *nothing* about reference points, so that learning fails even with an arbitrary number of price observations. If r^j were to be interpreted as a trader's private information, the lemma says that prices never reveal any information. And this happens for two reasons.

The large number of traders intuitively *washes out* any individual's price contribution. Each individual's price contribution is normalized by the market maker constant $\frac{1}{\sqrt{n}}$, and so is zero in the limit. It is intuitive that, when there are an infinite number of traders, each must have zero price impact. When each trader does not even impact the price, inverting that price to learn an individual r^j becomes impossible.

Recall that $\varepsilon \to 0$ represents a Bayesian trader. In this case, prices become *non-revealing* in the sense of Radner (1979). Imagine an extreme case with only one trader; the periodic demands proposed here are not one-to-one. Figure 7 captures the lack of invertibility: prices tomorrow could be used to back out aggregate demand, given by x. Today's price is given by p. All four points, r^j_1, r^j_2, r^j_3, and r^j_4, are viable candidates for the reference point which generated this demand. As $\varepsilon \to 0$ and the period of demand functions tend to zero, non-invertibility problems are further exacerbated.[15]

[13] Non-revealing prices, if constrained to coarse beliefs, are a weaker result. Because I take limits $\varepsilon \to 0$ and recover the Bayesian special case, the stronger result is needed.

[14] Two random variables with densities f_{xn} and f_{yn} are independent in the limit if $\lim_{n\to\infty} f_{xn,yn}(x,y) = \lim_{n\to\infty} f_{xn}(x) f_{yn}(y)$.

[15] The reader might wonder why this multiplicity in r^j even matters; they all generate the same demand. First, this lemma is a stepping stone to the main result, which is that traders must not learn others' *demands*. Second, it is not immediately clear that $\{r^j + \ell\varepsilon\}_{\ell\in\mathbb{Z}}$ is a sufficient statistic for the demand function. Consider a case where there are many (but finite) traders, so that a single price observation is

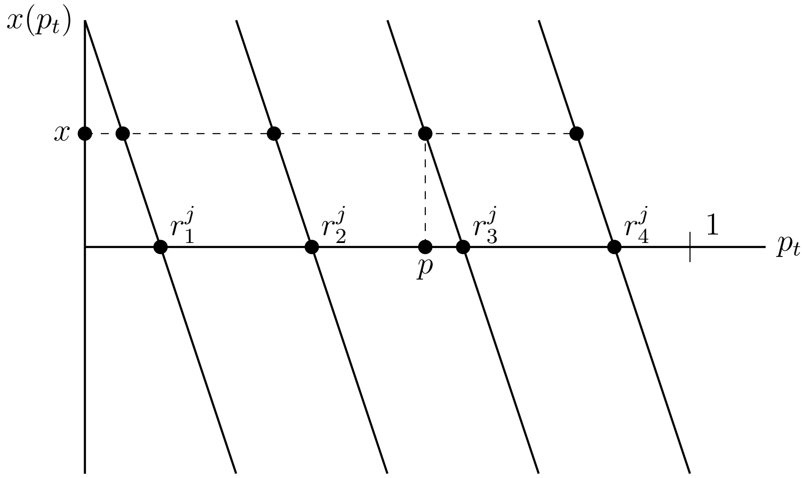

Figure 7: Partially Revealing Prices

Why do I send $n \to \infty$ and $\varepsilon \to 0$? One potential answer is that we truly believe there are many risk-neutral Bayesian traders, and another is that we want to understand in what sense the non-asymptotic version of the model approximates an equilibrium. There is a third and more candid reason, which is for tractability. Computing coarse expectations when traders learn some nonzero information from prices requires modelling higher-order beliefs (how traders believe other traders learn); common knowledge of non-revealing prices alleviates such issues.

The notion of asymptotically independent random variables is non-standard. The definition reads: in the limit, the joint density can be written as the product of two densities. But outside the limiting point, the interpretation is less clear. For any fixed n and ε, prices and reference points are not independent; independence is a Boolean (either true or false). Let me first reiterate the lemma: prices converge in distribution to a random variable that is independent from reference points. But an alternative and more intuitive interpretation will be discussed later in this section: for any precision statistical test for independence, there exists an n and ε such that the test will return a positive result.

Brownian Prices

insufficient for learning all information. Traders would need to combine the data observation with the prior, the distribution over r^j, to form their posterior. This prior has location (it is not periodic), and hence the posterior is over r^j (not the entire set $\{r^j + \ell\varepsilon\}_{\ell\in\mathbb{Z}}$).

Next I work towards establishing a perfectly non-revealing equilibrium: aggregated trades induce a Brownian motion price process, yet the price process impedes traders from learning their differences. The previous lemma directly addressed the second part of the statement, and my main proposition, discussed next, addresses convergence to a Brownian motion.

Proposition 1 (Brownian Prices): Assuming A1-A7 and a linear market maker, prices converge in distribution to a discrete time Brownian motion as $\varepsilon \to 0$ and $n \to \infty$.

Proof: See Appendix D.

As the number of traders tends to infinity, the aggregate demand function begins to resemble noise centered around zero over the price domain. It is a random object, but it is static in the sense that it is drawn at time $t = 0$. With the appropriate scaling guaranteed by Lemma 1, the central limit theorem can be applied to conclude that aggregate demand at any price is normally distributed.

However, of the requirements of a Brownian motion, it is not the normality that should come as a surprise. More surprising is the i.i.d. nature of price increments over time, which is equivalent, via the linear market maker, to the i.i.d. nature of excess demand over the price domain. To see how this works, consider periodic demands with an increased frequency induced by smaller ε. Take any initial distribution over reference points like the normal denoted by the red curve in Figure 8. Split the domain into increasingly smaller, equi-spaced intervals as shown by the black demand lines. The demand function takes probability mass from each interval and folds it across the range of the function, effectively stretching out probability mass as interval sizes tend to zero. At extremely fine partitions, only the shape of demand curves − not the original distribution − matters for determining the excess demand distribution. This observation, combined with the fact that demand is periodic, implies that excess demand becomes identically distributed at any price as ε tends to zero. This could be seen analytically in the proof of Lemma 2, when the characteristic function for demand lost all dependence on the initial distribution of heterogeneity, $h(r)$, as $\varepsilon \to 0$.

Interestingly, this "folding and stretching" also creates independently distributed demand across the price domain. This could be seen analytically in the proof of Proposition 1, when the covariance between demand evaluated at two different prices tended to zero as $\varepsilon \to 0$. Although proofs apply more generally, in the linear demand example depicted in Figure 8, excess demand approaches a uniform distribution independently of any price p_t and any initial distribution $h(r)$ as $\varepsilon \to 0$.

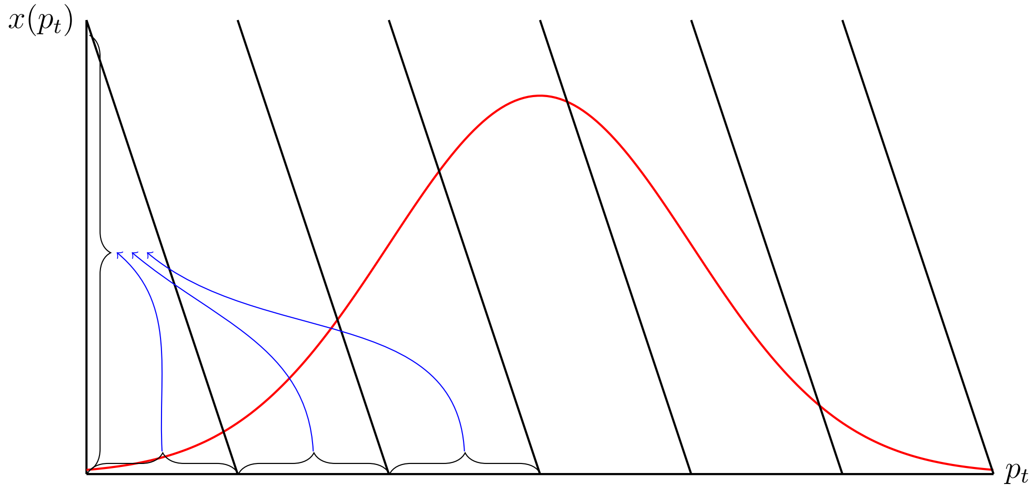

Figure 8: Folding and Stretching

Equilibrium

Finally, I address the assumption of myopia on the part of traders. Although previously traders simply guessed that prices are a unit-variance martingale, the discrete time Brownian motion satisfies these two properties and hence this guess is confirmed. I can now define a δ-equilibrium, which is a set of demand functions $x(p_t, r^j)$, a perceived stochastic process for prices characterized by measure μ_b, and an actual stochastic process for prices with measure μ satisfying:

(a) Given μ_b, $x(p_t, r^j)$ solves each trader's optimization problem.

(b) The market maker rule implies that $d(\mu_b, \mu) \leq \delta$.[16]

[16] While it is not the only choice, one metric that may be used here is $\limsup_{T\to\infty} D_T$, where D_T is defined in Appendix E.

I call the special case when $\delta = 0$ a *perfectly non-revealing equilibrium.* And an immediate corollary of Lemma 2 and Proposition 1 is that a discrete time Brownian motion is a perfectly non-revealing equilibrium as $n \to \infty$ and $\varepsilon \to 0$. When $\delta = 0$, what makes this otherwise-standard equilibrium definition perfectly non-revealing is that the random variable r^j is drawn only once. I allow traders to condition their trades on past price realizations, yet they learn nothing about the process ($\mu_{bt} = \mu_b$ for all t) beyond its realization.

In a standard model of speculative trade, Brownian prices are a result of exogenous information arrival. Here, Brownian prices are a result of price discovery, with considerable trade volume generated by coarse beliefs. However, these two narratives are not at odds with one another. One could construct a model with both informed traders who receive a signal (a martingale) and coarse Bayesians who do not receive a signal. In simpler terms, this market would consist of both fundamental and technical traders. While the equilibrium price process may differ from a Brownian motion, coarse Bayesians act exactly as described above and generate the noise typically assumed in standard models. That is, one use case for the results derived here is a microfoundation for noise trading.

Now consider δ-equilibria with $\varepsilon > 0$. Because traders make ε-sized errors when forming beliefs, and because these errors are heterogeneous, it is impossible for all traders to agree with one another, let alone on equilibrium prices. To allow for these errors, I relax the equilibrium concept up to some predetermined error band $\delta > 0$. Intuitively, δ should be thought of as a number proportional to (or equal to) ε. This gives the following interpretation: although traders have heterogeneous coarse expectations, they must be *as correct as they can be* about equilibrium prices, given their constraint. I maintain the limit $n \to \infty$ so I can continue to invoke Lemma 2, prices are non-revealing, and traders never learn their differences.

With nonzero ε, the resulting distribution of price increments is neither independent nor identically distributed. It is normally distributed by the central limit theorem. It follows from arguments in Proposition 1 that the resulting distribution of price increments approaches $N(g_1(\varepsilon), g_2(\varepsilon))$ where $\lim_{\varepsilon \to 0} g_1(\varepsilon) = 0$ and $\lim_{\varepsilon \to 0} g_2(\varepsilon) = (t_2 - t_1)$. And this is intuitive. Imagine a price path like that of Figure 6 with large ε. Prices fail to be a martingale; prices near the edge of a regime tend to revert toward the center of the regime. The variance also becomes non-stationary, because traders tend to trade larger volumes near the edges of each regime. This highlights an important difference between the δ-equilibrium with $\delta > 0$ and $\varepsilon > 0$ (a non-martingale which is *not* robust to rational arbitrage) and the perfectly non-revealing equilibrium (a martingale which is robust to rational arbitrage).

There are, in fact, two reasons why the equilibrium must be defined using a relaxed distance between beliefs and reality. The first reason, mentioned above, relates to the mean of the belief. Traders are constrained to have heterogeneous posterior means so they can never agree with each other, let alone on the correct equilibrium process. The second reason relates to the variance of the belief: the unit-variance guess is incorrect when $\varepsilon > 0$. While a natural next step might be to have traders form beliefs over the complicated process shown in Figure 6 (a non-martingale with non-stationary variance), I refrain from this exercise for the following reason. It is a highly complex exercise from the perspective of traders, which goes against the very spirit of my narrative: traders use approximations and rules of thumb. Instead, when $\varepsilon > 0$, I let traders form an incorrect (but simple) guess, which I have shown to be a very good guess for small but nonzero values of ε.

Numerical Interpretation

I have shown that prices converge in distribution to a discrete time Brownian motion as $n \to \infty$ and $\varepsilon \to 0$; this is the large market, risk-neutral, Bayesian limit of my model. Here I offer an interpretation of this statement off of the limiting point, that is, for finite n and nonzero ε. For any precision statistical test, there exist N and E such that, for $n \geq N$ and $\varepsilon \leq E$, the test cannot distinguish between the system (1), (3), and (6) and a Brownian motion. At these non-asymptotic values, prices *look* like a Brownian motion. For details of this Kolmogorov-Smirnov test, see Appendix E.

The highlights of the exercise are as follows. The test, as well as the naked eye, can easily distinguish between the dynamical system and a Brownian motion when $\varepsilon = 100$ (see Figure 6). The test cannot distinguish between the dynamical system and a Brownian motion when $\varepsilon = 1$ (see Figure 5) when less than 75 iterations of the system are considered. As the number of iterations under consideration is increased, ε must tend to zero and n must tend to infinity rather quickly for the test to be unable to distinguish between the system and a Brownian motion.

The Econometrician

The discussion can quickly become philosophical when considering an omniscient econometrician who knows the realization r. To her, the price path is a pre-determined trajectory indistinguishable from a draw of a Brownian motion. It is only from the perspective of the traders who populate this economy that the price path is stochastic. To them, there is an aggregate demand function that determines prices, and that function is i.i.d. over its domain. It is drawn at time $t = 0$. Traders observe individual realizations of this aggregate demand function through the information revealed in prices, but precisely due to its i.i.d. nature, these realizations reveal nothing about the global shape of the aggregated demand function. Traders remain in a state of ignorance well past time $t = 0$, and hence are unable to predict the future price path.

The key to my result is that I do not allow traders to observe (and condition trades on) others' reference points r^j. Each trader's information is restricted to only observing prices p_t, and I show that these prices are non-revealing. An appropriate analogy might be a die that is rolled, but covered by the dealer's hand. If there is no hope of peaking around the dealer's hand, the die is treated in a similar fashion to one that has yet to be rolled. All random variables are *realized* at time zero, but they are not *revealed* to traders through the information contained in prices.

In Appendix F, I consider a scenario where the realization r is known to the econometrician, and both $n < \infty$ and $\varepsilon > 0$ are fixed. I show that, to the econometrician, the dynamical system is chaotic but perfectly deterministic.

5 Extensions and Policy

There are several immediate extensions to the model. To achieve a geometric Brownian motion from Black and Scholes (1973), I redefine prices in logs. Otherwise, the derivation from the previous sections is used; it is the percentage change in price that is now normally distributed.

The discrete time Brownian motion can be rescaled to one that assigns probability one to paths that are uniformly continuous; see Durrett (2017) Theorem 7.1.2 for a formal treatment. The valuable learning from the continuous time limit comes from the rescaling. The market maker constant must be redefined with an additional factor of $\frac{1}{\sqrt{T}}$, where T denotes the frequency of trade in the time interval $[0, 1]$. Intuitively, there is an invariance to whether there are more traders n or a higher frequency of trade T.

Heavy-Tailed Distributions

In my main extension, I extend the Brownian motion price process to a broader family of heavy-tailed distributions. The question of which family of distributions best describes stock market returns has yet to be settled empirically; however, many studies find support for the stable-α family first proposed by Benoit Mandelbrot and Eugene Fama in the 1960's (Hoechstoetter et al. (2005), Lux (1996), Rachev et al. (2005)). There is also a theoretical reason for focusing on the stable-α family mentioned in the proof of Lemma 1. The Durrett necessary and sufficiency condition states that a summed and normalized i.i.d. sequence can only converge to a stable-α law. Therefore, the only theoretical candidate for price increments is the stable-α law, making this extension, in some sense, maximal.

I choose to maintain assumptions A1-A3 for the same reasons outlined previously. Relaxing A4 is the key to breaking the Durrett necessary and sufficient condition; non-i.i.d. reference points can lead to an even broader family of price distributions. While this is an interesting line of reasoning for future work, in the current paper I leave A4 as is. A5 remains intact: differentiable utility allows me to take first-order conditions. The assumptions I will need to revisit, A6 and A7, relate to the magnitude and scaling of demand. In words, these two assumptions state that demand must be square integrable (not too large) and centered (zero mean); see Appendix B for their formal descriptions. In fact, the Mikosch necessary and sufficient condition from the proof of Lemma 1 gives the exact condition required to achieve stable-α price increments.

In words, this condition states that demand volume must follow a power law. Although it is derived as a necessary and sufficient theoretical condition, there is also extensive evidence that trading volume does, in fact, obey a power law (Gabaix et al. (2003), Plerou et al. (2004), Balakrishnan et al. (2008)). Power law tails imply that demands are unbounded; there are prices where traders desire arbitrarily large positions based on perceived arbitrage opportunities. Figure 9 shows a plausible demand function satisfying these conditions. Like in the Motivating Example, market regimes arise from coarse expectations. The only difference is that traders now trade more aggressively near the edge of these regimes.

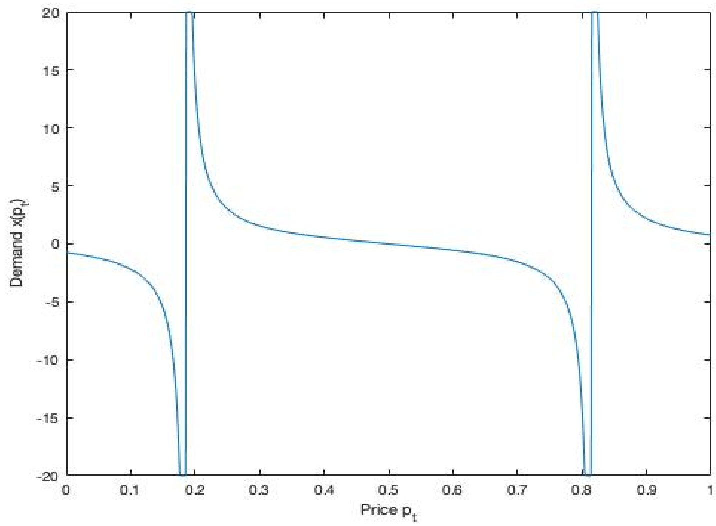

Figure 9: Unbounded Demand

Figure 9 is in the spirit of Black (1986), who claimed the farther the price of a stock gets from its value, the more aggressive traders become. Day and Huang (1990) also discuss similarly shaped demands, but they are not periodic. I now formally replace assumptions A6 and A7 with:

A6′: Power law demand volume, $\int_{-0.5}^{0.5} 1_{\{|(L')^{-1}(p)|>y\}} dp \sim ky^{-\alpha}$ for $1 < \alpha < 2$ and $k > 0$.
A7′: Demand is symmetric, $L(-x) = L(x)\ \forall x$.

Both assumptions are imposed on the generalized loss function $L(x)$ in trader utility functions, described by:

$$u(r^j, p_t, x_t) = (\tilde{E}_j[p_{t+1}|p_t] - p_t)x_t - \rho L(x_t) \tag{7}$$

where $L(x)$ was previously quadratic in the case of the Brownian motion. One interpretation of (7) is a higher-order Taylor series approximation, where ρ now represents aversion to higher moments. The first-order condition is:

$$L'(x_t) = \frac{\tilde{E}_j[p_{t+1}|p_t] - p_t}{\varepsilon}$$

which holds for small ε, assuming A3. Then demand functions are explicitly defined by taking the inverse, $x(p) = (L')^{-1}(p)$. A6′ says that the mass of prices at which demand volume $|x(p)|$ is greater than y is asymptotically $y^{-\alpha}$, up to the factor of a constant.[17] While this is only an assumption on the *shape* of demand functions, arguments provided in Appendix G show that this implies that the demand volume (the random variable) follows a power law. A6′ is consistent with power law volumes observed empirically, and is (almost) a necessary condition theoretically for heavy-tailed price distributions. Any further weakening would need to stay within the Mikosch necessary and sufficient condition.[18] Intuitively, this new loss function penalizes the trader at a much *lower* rate than the quadratic loss, leading to much *higher* trade volumes. These higher trade volumes, in turn, lead to higher probability of tail events (heavy-tailed prices).

Most empirical studies indicate that the relevant range for stable-α price increments is $1 < \alpha < 2$. An additional reason for omitting the case $\alpha = 2$ is that it corresponds to the Brownian motion already explored in the previous section. When $\alpha \leq 1$, the mean of price increments is undefined, making coarse expectations not well-defined. A6′ replaces A6; square integrable demands were what previously guaranteed finite second moments, which allowed me to apply the central limit theorem.

A7′ is a slight strengthening of the previous A7, and is for tractability only. Intuitively, traders receive the same disutility from selling as buying. While it may be intuitive that the loss function being even implies demand symmetry, it follows formally from the

[17]The range of integration, $[-0.5, 0.5]$, is the range of the right-hand side of the first-order condition, and hence the domain of the function $x(p)$.

[18]Such weakening involves slowly varying functions and $o(1)$ terms as described in Appendix A.

following line of logic:

$$\begin{aligned} & \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{p>L'(y)\}}dp = \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{p<-L'(y)\}}dp \ \ \forall y \\ \iff & \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{p>L'(y)\}}dp = \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{p<L'(-y)\}}dp \ \ \forall y \\ \iff & \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{x(p)>y\}}dp = \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{x(p)<-y\}}dp \ \ \forall y \end{aligned}$$

where the first equality follows from the symmetry of a Uniform$[-\frac{1}{2}, \frac{1}{2}]$ random variable, the second equality follows from A7$'$ and the fact that $L(x)$ being even implies $L'(x)$ is odd, and the last equality follows from the definition of demand. While some studies find an asymmetry between positive and negative asset returns, these results are still highly debated (see Jondeau and Rockinger (2003)), and so I choose to proceed with symmetric demands.

My final proposition shows that these updated assumptions are sufficient to conclude that prices converge in distribution to a Lévy process with stable-α price increments. Much like the Brownian motion, the Lévy process has stationary and independent increments. The price increments are now expanded from the normal to the heavy-tailed stable-α family.

Proposition 2 (Heavy-Tailed Prices): Assuming a linear market maker, A1-A5, A6$'$, and A7$'$, prices converge in distribution to a discrete time Lévy process with stable-α increments as $\varepsilon \to 0$ and $n \to \infty$.

Proof: See Appendix G.

In a setting where price increments are proportional to excess demand, a necessary condition for heavy-tailed prices is unbounded demand: traders trade aggressively based on perceived arbitrage opportunities. In Appendix G, I show that the market maker must adjust her scaling from $\frac{1}{\sqrt{n}}$ to $\frac{1}{n^{1/\alpha}}$; that is, she must restrict price movements more so than in the Brownian case. The net effect is heavy-tailed prices.

Finally I am ready to characterize equilibrium as in the previous section. A corollary of

Lemma 2 and Proposition 2 is that a discrete time Lévy process with stable-α increments is a perfectly non-revealing equilibrium as $n \to \infty$ and $\varepsilon \to 0$. Note that Figures 2 and 3 from the Introduction, which plot power law volume and stable-α prices, respectively, satisfy the conditions of this last proposition.

Proposition 2 holds true assuming that trader preferences are given by (7). However, an implicit assumption made in condition (7) requires a careful reinterpretation. In the previous case of the Brownian motion, the variance of price increments, $\text{Var}[p_{t+1} - p_t | p_t] = 1$, multiplied the quadratic loss term. For any $\alpha < 2$, all price moments higher than the first are undefined. While this is not problematic at an intuitive level because of neutrality to risk (and higher moments), in Appendix H, I outline a formal mathematical alternative. This last observation highlights a broader takeaway: risk neutrality is a necessary condition for nonzero trade in any model with heavy-tailed stable-α prices.

Heterogeneous ε

When ε is large, I have argued that the definition of equilibrium must be relaxed to accommodate disagreement in posteriors. Instead of considering heterogeneous ε in such a relaxed setting, I consider a stronger result. I ask whether heterogeneous ε can be supported in a perfectly non-revealing equilibrium. Since, by definition, grid size will need to tend to zero, I consider a setting with heterogeneous grids and let the coarsest partition (and hence all partitions) tend to zero.

Details of this exercise are relegated to Appendix I. The intuition is the following: I define heterogeneous i.i.d. *groups* of individuals, and apply previous arguments group-by-group. Then I make use of a general property of stable-α, and hence also normally distributed, random variables: linear combinations of independent stable-α random variables also have stable-α distributions.

Policy

Instead of evaluating welfare from the perspective of traders with coarse expectations, I take it as given that volatility has negative overall welfare implications. This should not be overly controversial, and would be justified in the presence of additional risk-averse traders who choose to avoid volatile markets. In the case of a Brownian motion, a proportional capital gains tax is effective at reducing volatility as measured by the variance of price increments. The new trader utility function would be defined:

$$u(r^j, p_t, x_t) = (1 - \tau)(\tilde{E}_j[p_{t+1}|p_t] - p_t)x_t - \frac{\varepsilon}{2}x_t^2$$

where τ is the capital gains tax. The tax multiplies the previous demand functions, and ultimately multiplies the market maker rule:

$$p_{t+1} = p_t + (1 - \tau)\frac{cz_n(p_t, r)}{\sqrt{n}}$$

The resulting normal distribution for price increments will have a dampened variance as desired.

The capital gains tax is, however, ineffective in the case of Lévy prices because variance is undefined. With unbounded demand, policies targeted at restricting arbitrarily large trade volumes are more effective. A social planner can impose trading constraints for individual traders such that demand must stay within some range $|x(p_t)| \leq M$. M can be made large so that constraints bind rarely. Truncated demand functions are now square integrable; in fact, any bounded function is square integrable.[19] Then Proposition 1 applies, and prices converge in distribution to a Brownian motion. So trading constraints on both long and short positions can be used to alter the equilibrium price distribution from one of heavy tails to a standard Brownian motion with finite variance. A combination of the two policies guarantees that the variance of price increments will be dampened as desired.

6 Conclusion

There are three main takeaways from the paper. First, a connection between coarse expectations and price barriers was established. When traders' expectations were restricted to take on only a discrete set of possible values, demand discontinuities naturally arose. A discontinuity was interpreted as a cutoff between regimes such as a bull and bear market, and provided a foundation for trigger strategies like the stop-loss order. Under- and over-reactions, as well as momentum and reversals, could be explained using coarse expectations. Price barriers could be supported in equilibrium when beliefs and reality

[19]This last statement is true on any finite measure space.

could differ by no more than δ.

Second, a perfectly non-revealing equilibrium was established in the risk-neutral Bayesian limit. Equilibrium consisted of two components. First, there was a belief of a particular stochastic process. Instead of explicitly modelling trader learning and price invertibility, I considered conditions under which such learning fails entirely. Under the large market or risk-neutral Bayesian limit, traders learned nothing from price and entered into the second period with the same information set as in the first. The other half of equilibrium was the trader's action. Given a belief of a particular stochastic process, aggregated trades were shown to confirm that price process in reality: coarse beliefs led to disagreement and sizeable trade volume, and the ensuing price discovery process converged to a Brownian motion. Altogether, in equilibrium, traders induced a Brownian motion price process because of their differences, yet it was this very process that impeded traders from learning their differences.

Surprisingly, neither of the two nonstandard assumptions required for the perfectly non-revealing equilibrium consisted of coarse Bayesian updating. As coarseness tended to zero, the standard Bayesian case was recovered. However, trader risk-neutrality amplified even the smallest expectation errors to ensure nonzero trade. The second nonstandard assumption related to the market maker; prices were scaled in such a way that they converged to an equilibrium distribution, as opposed to a steady state equilibrium.

Third, I showed that the same mechanism discussed for a Brownian motion was capable of generating heavy-tailed price distributions in equilibrium. The necessary condition for such heavy tails was unbounded demand, or prices at which traders traded arbitrarily large amounts based on their coarse expectations. In terms of policy, I find that trading constraints on excessive trade volume can eliminate heavy tails in equilibrium. In the case of the simpler Brownian motion, a capital gains tax is effective at reducing the variance of price increments.

Throughout the paper, I have discussed several avenues for future work, such as considering a nonlinear market maker, adding informed traders, and relaxing the i.i.d. requirement across traders. Two broader questions remain; the first is: To which other markets might this mechanism apply? Stock markets were a natural starting point, due to not only the "rule of thumb" interpretation for technical traders but also the tick size interpretation for algorithmic trading. Perhaps the broadest open question relates to the fact that coarse expectations were not a necessary component of the equilibrium established here; these simply guaranteed nonzero trade. Any heterogeneity that creates a “noisy” aggregate demand function would suffice: the irregularities in excess demand prevent traders from learning their differences, and their differences are what create an irregularly shaped aggregate demand. Why are prices volatile? Many macroeconomic models assume shocks, many finance models assume noise traders, and here I have shown that volatility can arise from heterogeneous mistakes. If the question of the origin of financial market fluctuations is to be taken seriously, we must investigate our differences that make us trade.

References

D. Abreu and M. Brunnermeier. Bubbles and crashes. *Econometrica*, 71(1):173–204, 2003.

R. Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684, 1994.

B. Allen. The existence of fully rational expectations approximate equilibria with noisy price observations. *Journal of Economic Theory*, 37(2):213–253, 1985.

F. Alvarez and A. Atkeson. Random risk aversion and liquidity: a model of asset pricing and trade volumes. *Working Paper*, 2021.

W. Bagehot. The only game in town. *Financial Analysts Journal*, 27(2):12–14, 1971.

P. Balakrishnan, J. Miller, and G. Shankar. Power law and evolutionary trends in stock markets. *Economics Letters*, 98(2):194–200, 2008.

R. Bansal and A. Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004.

R. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866, 2006.

F. Black. Noise. *Journal of Finance*, 41(3):529–543, 1986.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

J. Campbell. *Financial Decisions and Markets: a Course in Asset Pricing*. Princeton University Press, 2017.

J. Campbell and J. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.

J. Cochrane. Efficient markets today. *Talk at Conference of Chicago Economics*, Nov, 2007.

J. Cox, S. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, 1979.

K. Daniel, D. Hirshleifer, and A. Subrahmanyam. Investor psychology and security market under- and overreactions. *Journal of Finance*, 53(6):1839–1885, 1998.

R. Day and W. Huang. Bulls, bears and market sheep. *Journal of Economic Behavior and Organization*, 14(3):299–329, 1990.

G. Donaldson and H. Kim. Price barriers in the dow jones industrial average. *Journal of Financial and Quantitative Analysis*, 28(3):313–330, 1993.

R. Durrett. *Probability: Theory and Examples*. Cambridge University Press, 2017.

E. Eyster and M. Piccione. An approach to asset pricing under incomplete and diverse perceptions. *Econometrica*, 81(4):1483–1506, 2013.

D. Farmer and S. Joshi. The price dynamics of common trading strategies. *Journal of Economic Behavior and Organization*, 49(2):149–171, 2002.

X. Gabaix, P. Gopikrishnan, V. Plerou, and E. Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423:267–270, 2003.

F. Gul, W. Pesendorfer, and T. Strzalecki. Coarse competitive equilibrium and extreme prices. *American Economic Review*, 107(1):109–137, 2017.

M. Hoechstoetter, S. Rachev, and F. Fabozzi. Distributional analysis of the stocks comprising the dax 30. *Probability and Mathematical Statistics*, 25(2):363–383, 2005.

U. Horst and C. Rothe. Queuing, social interactions, and the microstructure of financial markets. *Macroeconomic Dynamics*, 12(2):211–233, 2008.

A. Jakobsen. Coarse bayesian updating. *Working Paper*, 2021.

E. Jondeau and M. Rockinger. Testing for differences in the tails of stock-market returns. *Journal of Empirical Finance*, 10(5):559–581, 2003.

A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

A. Lei and H. Li. The value of stop loss strategies. *Financial Services Review*, 18(1): 23–51, 2009.

T. Lux. The stable paretian hypothesis and the frequency of large returns: An examination of major german stocks. *Applied Financial Economics*, 6(6):463–475, 1996.

T. Mikosch. *Regular Variation, Subexponentiality and their Applications in Probability Theory.* EURANDOM Report 99-013, 1999.

P. Milgrom and N. Stokey. Information, trade, and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

S. Mullainathan. Thinking through categories. *Working Paper*, 2002.

P. Ortoleva. Modeling the change of paradigm: Non-bayesian reactions to unexpected news. *American Economic Review*, 102(6):2410–2436, 2012.

M. Pakkanen. Microfoundations for diffusion price processes. *Mathematics and Financial Economics*, 3(2):89–114, 2010.

A. Park and H. Sabourian. Herding and contrarian behavior in financial markets. *Econometrica*, 79(4):973–1026, 2011.

V. Plerou, P. Gopikrishnan, X. Gabaix, and E. Stanley. On the origin of power-law fluctuations in stock prices. *Quantitative Finance*, 4(1):7–11, 2004.

H. Polemarchakis and P. Siconolfi. Asset markets and the information revealed by prices. *Economic Theory*, 3:645–661, 1993.

S. Rachev, S. Stoyanov, A. Biglova, and F. Fabozzi. An empirical examination of daily stock return distributions for u.s. stocks. *Data Analysis and Decision Support*, pages 269–281, 2005.

R. Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, 47(3):655–678, 1979.

R. Raimondo. Market clearing, utility functions, and security prices. *Economic Theory*, 25:265–285, 2005.

R. Shiller. *Irrational Exuberance.* Princeton University Press, 2015.

J. Steiner and C. Stewart. Price distortions under coarse reasoning with frequent trade. *Journal of Economic Theory*, 159:574–595, 2015.

H. Varian. Differences of opinion in financial markets. *Financial Risk: Theory, Evidence and Implications*, pages 3–37, 1989.

Appendices

Appendix A

To prove Lemma 1, first I begin with the linear rule (5), which follows from A1. Durrett (2017) Theorem 3.8.8 states that Y is the limit of a summed and normalized i.i.d. sequence X_i if and only if Y has a stable law; I will refer to the result as the Durrett necessary and sufficient condition. Therefore, if $a_n z_n(p_t, r)$ converges to a nondegenerate distribution (it converges either to a degenerate or nondegenerate distribution by A2), it must be in the stable-α family. The parameter α nests several important distributions: $\alpha = 2$ is the Gaussian, and $\alpha = 1$ is the Cauchy.

Mikosch (1999) Theorem 1.4.7 states that a summed and normalized sequence of i.i.d. random variables with distribution F converges to a stable-α law for $\alpha < 2$ if and only if:

$$F(-x) = \frac{q + o(1)}{x^\alpha} S(x), \quad 1 - F(x) = \frac{p + o(1)}{x^\alpha} S(x), \quad x \to \infty$$

for slowly varying $S(x)$ and nonnegative p, q such that $p + q > 0$; I will refer to the result as the Mikosch necessary and sufficient conditions.[20] Since demand functions in (1) are bounded, the case $\alpha < 2$ is ruled out. The only remaining possibility is the normal $\alpha = 2$, for which we can apply the standard Lindeberg-Lévy theorem. Now say $a_n z_n(p_t, r)$ converges to a degenerate distribution, in other words, a constant. If a random variable converges in distribution to a constant, it converges in probability to that constant. By the sufficient direction ($\Leftarrow$) of A1, the constant must equal zero. But then $p_{t+1} \to p_t$ in probability, which contradicts the necessary direction ($\Rightarrow$) of A1, and the proof is complete. $\square$

Appendix B

In a more general setting, I assume that traders get utility and disutility from trade according to:

$$u(r^j, p_t, x_t) = (\tilde{E}_j[p_{t+1}|p_t] - p_t)x_t - \rho L(x_t)$$

[20] A slowly varying function $S(x)$ satisfies $\lim_{x\to\infty} \frac{S(tx)}{S(x)} = 1$ for all $t > 0$. Examples include functions with a nonzero limit and $ln(x)$.

where $L(x_t)$ is the disutility function, which was previously quadratic. These preferences can be interpreted as a higher-order Taylor series approximation, where ρ now represents aversion to higher moments. With this generalized setup, I introduce three assumptions on traders that supplant the previous second-order Taylor series approximation:

A5: $L(x)$ is strictly convex and differentiable.
A6: Demand functions are square integrable, $\int_{-0.5}^{0.5} \left((L')^{-1}(p)\right)^2 dp < \infty$.
A7: Demand functions have zero mean, $\int_{-0.5}^{0.5} (L')^{-1}(p)dp = 0$.

A5, A6, and A7 are assumptions on the generalized loss function $L(x_t)$. As $\varepsilon \to 0$, I simply write $\rho = \varepsilon$ by A3. Then demand is now implicitly defined by the following first-order condition, when $L(x_t)$ is convex and differentiable:

$$L'(x_t) = \frac{\tilde{E}_j[p_{t+1}|p_t] - p_t}{\varepsilon} = \text{mod}\left(\frac{r^j - p_t}{\varepsilon} - \frac{1}{2}, 1\right) - \frac{1}{2} \tag{8}$$

where the second equality follows from explicit calculation of the updating rule. Hence, I denote the demand function as $x(p) = (L')^{-1}(p)$. The class of demand functions considered under (8) are periodic with period ε. As in the Motivating Example, the periodic nature of demand stems from traders discretely shifting price expectations at regular price intervals. A6 says that demands must be square integrable, $\int_{-0.5}^{0.5} x(p)^2 dp < \infty$, which translates to demands that are not too large.[21] A7 says that traders are just as likely to buy as they are to sell; this will ensure that the resulting prices are a martingale.

In the special case of quadratic utility, A5 holds and $(L')^{-1}(p)$ is the identity function. Then it is an algebraic exercise to verify that both A6 and A7 are satisfied.

Appendix C

To prove Lemma 2, I begin with asymptotics in n. The price increment is given by (6):

$$p_{t+1} - p_t = \frac{c z_n(p_t, r)}{\sqrt{n}}$$

[21] The range of integration, $[-0.5, 0.5]$, is the range of the right-hand side of (8), and hence the domain of the function $x(p)$.

Two random variables are independent if and only if their characteristic functions separate: $\varphi(s_1, s_2) = \varphi(s_1)\varphi(s_2)$. Here I denote $\varphi(s_1)$ as the characteristic function of normalized excess demand when $n \to \infty$, which exists by A2, and $\varphi(s_2)$ as the characteristic function of a reference point r^j. For simplicity of notation, I solve for individual one; the same argument applies to all individuals. Also, for brevity, I omit the constant c, which does not affect independence. $h(r)$ denotes the density of the reference point, which exists by A4. Then:

$$
\begin{aligned}
&\varphi_n(s_1, s_2) \\
=&E\left[\exp\left(is_1\frac{1}{\sqrt{n}}\sum_{j=1}^{n} x(p_t, r^j) + is_2 r^1\right)\right] \\
=&\int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} \exp\left(is_1\frac{1}{\sqrt{n}}\sum_{j=1}^{n} x(p_t, r^j) + is_2 r^1\right)\prod_{j=1}^{n} h(r^j)dr^j \\
=&\int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} \exp\left(is_1\frac{1}{\sqrt{n}}\sum_{j=2}^{n} x(p_t, r^j)\right)\exp\left(is_1\frac{1}{\sqrt{n}}x(p_t, r^1) + is_2 r^1\right)\prod_{j=1}^{n} h(r^j)dr^j \\
=&\int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} \exp\left(is_1\frac{1}{\sqrt{n}}\sum_{j=2}^{n} x(p_t, r^j)\right)\prod_{j=2}^{n} h(r^j)dr^j \int_{-\infty}^{\infty} \exp\left(is_1\frac{1}{\sqrt{n}}x(p_t, r^1) + is_2 r^1\right)h(r^1)dr^1 \\
\to&\varphi(s_1)\int_{-\infty}^{\infty} \exp(is_2 r^1)h(r^1)dr^1
\end{aligned}
$$

where the fourth equality follows from the independence guaranteed by A4, and the asymptotic result follows from dominated convergence. Notice that the same argument above can now be applied to the two-period-ahead price increment:

$$
p_{t+2} - p_t = \frac{c}{\sqrt{n}}\sum_{j=1}^{n} x(p_t, r^j) + \frac{c}{\sqrt{n}}\sum_{j=1}^{n} x(p_{t+1}, r^j)
$$

Applying the same line of reasoning above to the two-period-ahead increment, the fourth equality now holds because p_{t+1} and r^j are asymptotically independent. This logic then recursively applies to any price increment $p_{t+s} - p_t$.[22] For the asymptotic result in ε, I consider an extreme case with $n = 1$. I again manipulate the joint characteristic function,

[22]This general result is stronger than what is needed for the statement of the lemma, but will be referenced in later propositions.

now making use of the demand functional form (8):

$$
\begin{aligned}
\varphi_\varepsilon(s_1,s_2) &= \int_{-\infty}^{\infty} \exp\left(is_1 x \left[\text{mod}\left(\frac{r-p_t}{\varepsilon} - \frac{1}{2}, 1 \right) - \frac{1}{2} \right] + is_2 r \right) h(r)dr \\
&= \sum_{\ell=-\infty}^{\infty} \int_{p_t+\ell\varepsilon}^{p_t+(\ell+1)\varepsilon} \exp\left(is_1 x \left[\text{mod}\left(\frac{r-p_t}{\varepsilon} - \frac{1}{2}, 1 \right) - \frac{1}{2} \right] + is_2 r \right) h(r)dr \\
&= \sum_{\ell=-\infty}^{\infty} \int_0^1 \exp(is_1 x \left[\text{mod}\left(u - 0.5, 1 \right) - 0.5 \right]) \exp(is_2 [p_t + \varepsilon(u+\ell)]) h(p_t + \varepsilon(u+\ell))\varepsilon du \\
&= \int_0^1 \exp(is_1 x \left[\text{mod}\left(u - 0.5, 1 \right) - 0.5 \right]) \sum_{\ell=-\infty}^{\infty} \exp(is_2 [p_t + \varepsilon(u+\ell)]) h(p_t + \varepsilon(u+\ell))\varepsilon du
\end{aligned}
$$

where the third equality follows from a change of variable $u = \frac{r-p_t}{\varepsilon} - \ell$, and the fourth equality from dominated convergence. The sum in the expression above is a Riemann sum, $\sum_{\ell=-\infty}^{\infty} \exp(is_2 x_\ell^*) h(x_\ell^*)\varepsilon$, with:

$$
\begin{aligned}
x_\ell^* &= p_t + \varepsilon(u+\ell) \\
x_{\ell-1} &= p_t + \varepsilon\ell \\
x_\ell &= p_t + \varepsilon(\ell+1)
\end{aligned}
$$

so that the function is always evaluated at a point within the Riemann partition $x_{\ell-1} \leq x_\ell^* \leq x_\ell$, and the length of the partition tends to zero $x_\ell - x_{\ell-1} = \varepsilon$. I can then rewrite the characteristic function:

$$
\begin{aligned}
\varphi_\varepsilon(s_1,s_2) &= \int_0^1 \exp(is_1 x \left[\text{mod}\left(u - 0.5, 1 \right) - 0.5 \right]) \sum_{\ell=-\infty}^{\infty} \exp(is_2 x_\ell^*) h(x_\ell^*)\varepsilon du \\
&\to \int_{-\frac{1}{2}}^{\frac{1}{2}} \exp(is_1 x(u)) du \int_{-\infty}^{\infty} \exp(is_2 x^*) h(x^*) dx^*
\end{aligned}
$$

where again I have used dominated convergence and explicitly calculated the range of the modulo term. Riemann sums converge to Riemann integrals on compact intervals only. If the distribution is tight by A4, then the measure along the entire real line can be approximated by an appropriately large compact interval. □

Appendix D

To prove Proposition 1, first I recursively solve (6) from Lemma 1:

$$p_{t_2} - p_{t_1} = \sum_{t=t_1}^{t_2-1} \frac{cz_n(p_t, r)}{\sqrt{n}}$$

The goal is to show that the term above is normal with appropriate mean and variance. In order to do so, I choose to manipulate its characteristic function $\varphi_{n,\varepsilon}(s)$. Let $\mathbb{T}$ denote the set $\{t_1, ..., t_2\}$, and let $\mathbb{S}$ denote the set $\{q, t \mid t_1 \leq t < q \leq t_2 - 1\}$. For now I guess that a joint density $f_{\mathbb{T}}\left(\times_{q,t,\in\mathbb{S}}\, p_{qt}\right)$ for price increments $p_{qt} \equiv (p_q - p_t)$ exists; this will later be verified. In general, the density may depend on $\mathbb{T}$, and it is defined over the product of all price increments. The absolute continuity assumption A4 guarantees the existence of a density $h(r)$, which I use next in the manipulation of the characteristic function:

$$\begin{aligned}
\varphi_{n,\varepsilon}(s) =& E\left[\exp\left(\sum_{t=t_1}^{t_2-1} \frac{iscz_n(p_t, r)}{\sqrt{n}}\right)\right] \\
\approx& \int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} \exp\left(\sum_{t=t_1}^{t_2-1}\sum_{j=1}^{n} \frac{iscx(p_t, r^j)}{\sqrt{n}}\right) \prod_{j=1}^{n} h(r^j) dr^j f_{\mathbb{T}}(\cdot) \prod_{q,t\in\mathbb{S}} dp_{qt} \\
=& \int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} \exp\left(\sum_{t=t_1}^{t_2-1} \frac{iscx(p_t, r^1)}{\sqrt{n}}\right) h(r^1) dr^1\right)^n f_{\mathbb{T}}(\cdot) \prod_{q,t\in\mathbb{S}} dp_{qt}
\end{aligned}$$

In the approximate equality, I have applied the argument provided in Lemma 2: the independence of reference points and general price increments $p_{t+s} - p_t$. This will hold with equality when limits are taken. In the last equality, I have used the i.i.d assumption A4 for random variable r^j. Expectations are taken over both reference points and future prices; this is a necessary condition when considering prices multiple periods ahead.[23]

[23] To see why, consider prices two periods ahead. These depend on excess demand today and tomorrow, which are a function of prices tomorrow.

Then, using a Taylor series expansion, the characteristic function equals:

$$\left(1+\frac{is}{\sqrt{n}}E_{\mathbb{T}}\left[\sum_{t=t_1}^{t_2-1}cx(p_t,r^1)\right]-\frac{s^2}{2n}E_{\mathbb{T}}\left[\left(\sum_{t=t_1}^{t_2-1}cx(p_t,r^1)\right)^2\right]+E_{\mathbb{T}}[\xi(n,s)]\right)^n$$

$$=\left(1-\underbrace{\frac{s^2}{2n}E_{\mathbb{T}}\left[c^2\sum_{t=t_1}^{t_2-1}x(p_t,r^1)^2\right]}_{variance}-\underbrace{\frac{s^2}{n}E_{\mathbb{T}}\left[c^2\sum_{t,q\in\mathbb{S}}x(p_q,r^1)x(p_t,r^1)\right]}_{covariance}+\underbrace{E_{\mathbb{T}}[\xi(n,s)]}_{error}\right)^n$$

where the second equality follows from the fact that demand is centered, A7. I call the second term in the large parenthesis the variance term, and the third term in the large parenthesis the covariance term. The final term is the Taylor series error. As before, expectations $E_{\mathbb{T}}[\cdot]$ are taken with respect to both $h(\cdot)$ and $f_{\mathbb{T}}(\cdot)$. I first manipulate the variance term, for now ignoring the outer integral over prices:

$$\begin{aligned}
c^2\int_{-\infty}^{\infty}x(p_t,r^1)^2h(r^1)dr^1 =&c^2\int_{-\infty}^{\infty}x\left(\mathrm{mod}\left[\frac{r^1-p_t}{\varepsilon}-\frac{1}{2},1\right]-\frac{1}{2}\right)^2h(r^1)dr^1\\
=&c^2\sum_{\ell=-\infty}^{\infty}\int_{p_t+\ell\varepsilon}^{p_t+(\ell+1)\varepsilon}x\left(\mathrm{mod}\left[\frac{r^1-p_t}{\varepsilon}-\frac{1}{2},1\right]-\frac{1}{2}\right)^2h(r^1)dr^1\\
=&c^2\sum_{\ell=-\infty}^{\infty}\int_0^1 x\left(\mathrm{mod}\left[u-0.5,1\right]-0.5\right)^2h(p_t+\varepsilon(u+\ell))\varepsilon du\\
=&c^2\int_0^1 x\left(\mathrm{mod}\left[u-0.5,1\right]-0.5\right)^2\sum_{\ell=-\infty}^{\infty}h(p_t+\varepsilon(u+\ell))\varepsilon du\\
\to&c^2\int_{-\frac{1}{2}}^{\frac{1}{2}}x(u)^2du=1
\end{aligned}$$

with an appropriate choice of the normalization constant c. Notice that the last expression is well defined by A6.[24] The third equality above follows from a change of variables $u=\frac{r^1-p_t}{\varepsilon}-\ell$, and the fourth equality follows from dominated convergence, which applies due to A6. Using the same Riemann sum technique as Lemma 2, the sum inside the integral approaches the value one as $\varepsilon\to 0$. The entire variance term approaches $\frac{s^2}{2n}(t_2-t_1)$ even after the integral over prices is taken; the expectation of a constant is that constant. Next I manipulate the covariance for a fixed (q,t); for brevity in notation

[24]If demands have zero L^2 norm, traders never trade, and the price process approaches a Brownian motion with variance zero.

I suppress the modulo term inside of demand, as well as integrals over $\mathbb{S} \setminus (q,t)$:

$$
\begin{aligned}
E_{\mathbb{T}}[x(p_q,r^1)x(p_t,r^1)] &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x\left(\frac{r^1-p_q}{\varepsilon}\right)x\left(\frac{r^1-p_t}{\varepsilon}\right)h(r^1)dr^1 f_{\mathbb{T}}(p_{qt},\cdot)dp_{qt} \\
&= \int_{-\infty}^{\infty}\sum_{\ell=-\infty}^{\infty}\int_{p_q+\ell\varepsilon}^{p_q+(\ell+1)\varepsilon} x\left(\frac{r^1-p_q}{\varepsilon}\right)x\left(\frac{r^1-p_t}{\varepsilon}\right)h(r^1)dr^1 f_{\mathbb{T}}(p_{qt},\cdot)dp_{qt} \\
&= \int_{-\infty}^{\infty}\sum_{\ell=-\infty}^{\infty}\int_{0}^{1} x\left(u\right)x\left(u+\frac{p_q-p_t}{\varepsilon}\right)h(\varepsilon(u+\ell)+p_q)\varepsilon du f_{\mathbb{T}}(p_{qt},\cdot)dp_{qt} \\
&= \int_{-\infty}^{\infty}\int_{0}^{1} x\left(u\right)x\left(u+\frac{p_q-p_t}{\varepsilon}\right)\sum_{\ell=-\infty}^{\infty}h(\varepsilon(u+\ell)+p_q)\varepsilon du f_{\mathbb{T}}(p_{qt},\cdot)dp_{qt}
\end{aligned}
$$

where, like before, the third equality follows from a change of variable $u = \frac{r^1-p_q}{\varepsilon} - \ell$, and the fourth equality follows from dominated convergence, which I can invoke due to A6 and the Cauchy-Schwarz inequality:

$$
\int_0^1 \left| x\left(u\right)x\left(u+\frac{p_q-p_t}{\varepsilon}\right)\right| du \le \left(\int_0^1 \left| x\left(u\right)\right|^2 du \cdot \int_0^1 \left| x\left(u+\frac{p_q-p_t}{\varepsilon}\right)\right|^2 du\right)^{1/2}
$$

By Fubini's Theorem I can change the order of integration, and the covariance expression becomes:

$$
\begin{aligned}
&= \int_0^1\int_{-\infty}^{\infty} x\left(u\right)x\left(u+\frac{p_{qt}}{\varepsilon}\right)\sum_{\ell=-\infty}^{\infty}h(\cdot)\varepsilon f_{\mathbb{T}}(p_{qt},\cdot)dp_{qt}du \\
&= \int_0^1\sum_{k=-\infty}^{\infty}\int_{\varepsilon(k-u)}^{\varepsilon(k+1-u)} x\left(u\right)x\left(u+\frac{p_{qt}}{\varepsilon}\right)\sum_{\ell=-\infty}^{\infty}h(\cdot)\varepsilon f_{\mathbb{T}}(p_{qt},\cdot)dp_{qt}du \\
&= \int_0^1\sum_{k=-\infty}^{\infty}\int_0^1 x\left(u\right)x\left(v\right)\sum_{\ell=-\infty}^{\infty}h(\varepsilon[v+\ell]+p_t)\varepsilon f_{\mathbb{T}}(\varepsilon[v-u+k],\cdot)\varepsilon dvdu \\
&= \int_0^1\int_0^1 x\left(u\right)x\left(v\right)\sum_{\ell=-\infty}^{\infty}h(\varepsilon[v+\ell]+p_t)\varepsilon\sum_{k=-\infty}^{\infty}f_{\mathbb{T}}(\varepsilon[v-u+k],\cdot)\varepsilon dvdu \\
&\to \int_0^1\int_0^1 x\left(u\right)x\left(v\right)dvdu\int_{-\infty}^{\infty}f_{\mathbb{T}}(x^*,\cdot)dx^* \\
&= \int_0^1 x\left(u\right)du\int_0^1 x\left(v\right)dv\int_{-\infty}^{\infty}f_{\mathbb{T}}(x^*,\cdot)dx^* = 0
\end{aligned}
$$

where the third equality again follows from a change of variable $v = u + \frac{p_{qt}}{\varepsilon} - k$, and the fourth equality from dominated convergence. The asymptotic result as $\varepsilon \to 0$ is an application of the Riemann sum technique, and the final equality follows from centered demand, A7. $\int_{-\infty}^{\infty} f_{\mathbb{T}}(x^*, \cdot)dx^*$ denotes the marginal density obtained after integrating out p_{qt}. While the sum over $h(\cdot)$ is dealt with similarly to previous proofs, the Riemann sum over $f_{\mathbb{T}}(\cdot)$ is handled with caution because the random variable $(v-u)$ has a range of two. Formally, I split the sum:

$$\sum_{k=-\infty}^{\infty} f_{\mathbb{T}}(\varepsilon[v-u+k], \cdot)\varepsilon = \sum_{k=-\infty}^{\infty} f_{\mathbb{T}}(\varepsilon[v-u+2k], \cdot)\varepsilon + \sum_{m=-\infty}^{\infty} f_{\mathbb{T}}(\varepsilon[v-u+2m+1], \cdot)\varepsilon$$

Then with a step size of 2ε and the Riemann partitions:

$$\begin{aligned} x_k^* &= \varepsilon(v-u+2k) & x_m^* &= \varepsilon(v-u+2m+1) \\ x_{k-1} &= (2k-1)\varepsilon & x_{m-1} &= 2m\varepsilon \\ x_k &= (2k+1)\varepsilon & x_m &= (2m+2)\varepsilon \end{aligned}$$

I conclude that each sum approaches $\frac{1}{2}$ times the marginal density, or the total sum approaches the marginal density. Next I use an explicit bound on the Taylor series error term, $E_{\mathbb{T}}[\xi(n,s)]$, from Durrett (2017) Lemma 3.3.19:[25]

$$|\xi(n,s)| \leq \min \left\{ \frac{s^3}{6n^{3/2}} \left| \sum_{t=t_1}^{t_2-1} cx(p_t, r^1) \right|^3, \ \frac{s^2}{n} \left| \sum_{t=t_1}^{t_2-1} cx(p_t, r^1) \right|^2 \right\}$$

The bound is particularly useful for showing that $nE_{\mathbb{T}}[\xi(n,s)] \to 0$ as $n \to \infty$. To see why, notice that the second argument of the $\min\{\cdot\}$ function combined with A6 implies that $nE_{\mathbb{T}}[\xi(n,s)]$ is bounded; I can now apply dominated convergence. Then the first argument of the $\min\{\cdot\}$ function implies $nE_{\mathbb{T}}[\xi(n,s)] \to 0$. I can now take the limit as $n \to \infty$ of the entire characteristic function:

$$\lim_{n\to\infty} \left(1 - \frac{s^2}{2n}(t_2-t_1) + \frac{nE_{\mathbb{T}}[\xi(n,s)]}{n}\right)^n = \exp\left(-\frac{s^2}{2}(t_2-t_1)\right)$$

which follows from the definition of the exponential function, e^x. This is the characteristic function of the $N(0, t_2-t_1)$ distribution, as desired.

To prove the independence of price increments, I show the independence of arbitrary

[25]More generally $|e^{ix} - \sum_{m=0}^{n} \frac{(ix)^m}{m!}| \leq \min\{\frac{|x|^{n+1}}{(n+1)!}, \frac{2|x|^n}{n!}\}$.

$\frac{cz_n(p_t,r)}{\sqrt{n}}$ and $\frac{cz_n(p_q,r)}{\sqrt{n}}$; these steps are similar to those detailed above.[26] Two random variables are independent if and only if their characteristic functions satisfy:

$$\varphi(s_1, s_2) = \varphi(s_1)\varphi(s_2)$$

where $\varphi(s_1)$ is the characteristic function of normalized demand evaluated at p_t, and $\varphi(s_2)$ is the characteristic function of normalized demand evaluated at p_q. I begin with the joint characteristic function:

$$\begin{aligned}\varphi_{n,\varepsilon}(s_1, s_2) &= E\left[\exp\left(\frac{cis_1}{\sqrt{n}}z_n(p_t, r) + \frac{cis_2}{\sqrt{n}}z_n(p_q, r)\right)\right] \\ &\approx \int_{\infty}^{\infty}\left(\int_{-\infty}^{\infty}\exp\left(\frac{cis_1}{\sqrt{n}}x(p_t, r^j) + \frac{cis_2}{\sqrt{n}}x(p_q, r^j)\right)h(r^j)dr^j\right)^n f_{\mathbb{T}}(p_{qt})dp_{qt}\end{aligned}$$

In the approximate equality, I have used the i.i.d. assumption A4 on the random variable r^j, as well as the independence of general price increments $p_{t+s} - p_t$ and reference points r^j from Lemma 2. A Taylor series expansion with zero mean yields:

$$\begin{aligned}&\left(1 - \frac{c^2}{n}E_{\mathbb{T}}\left[\frac{1}{2}s_1^2x(p_t, r^j)^2 + s_1s_2x(p_t, r^j)x(p_q, r^j) + \frac{1}{2}s_2^2x(p_q, r^j)^2\right] + E_{\mathbb{T}}[\xi(n, s_1, s_2)]\right)^n \\ \to &\left(1 - \frac{s_1^2}{2n} - \frac{s_2^2}{2n} + E_{\mathbb{T}}[\xi(n, s_1, s_2)]\right)^n\end{aligned}$$

where the asymptotic result is for small ε, I have normalized variance to one using the constant c, and the covariance term is zero by previous arguments. I can rearrange terms:

$$\varphi_n(s_1, s_2) = \left(1 - \frac{s_1^2}{2n} + E_{\mathbb{T}}\left[\zeta(n, s_1, s_2)\right]\right)^n\left(1 - \frac{s_2^2}{2n}\right)^n$$

The term in the second set of large parentheses converges to the correct limit, but I need to ensure that the term in the first set of large parentheses converges to $\exp(-s_1^2/2)$. I explicitly write out the new error term:

$$E_{\mathbb{T}}\left[\zeta(n, s_1, s_2)\right] = \frac{1}{1 - s_2^2/(2n)}\left(E_{\mathbb{T}}\left[\xi(n, s_1, s_2)\right] - \frac{s_1^2s_2^2}{4n^2}\right)$$

Following previous arguments, $nE_{\mathbb{T}}[\zeta(n, s_1, s_2)]$ tends to zero by the Durrett bound and the dominated convergence theorem; therefore, the desired result holds for large n and

[26] While it may seem that zero covariance is sufficient to conclude independence of two normal random variables, this in fact requires the random variables be multivariate normal. This has yet to be shown.

the proof is complete:

$$\varphi_n(s_1, s_2) \to e^{-s_1^2/2} \cdot e^{-s_2^2/2}$$

□

Appendix E

To numerically test the difference between the dynamical system and a Brownian motion, I choose to run the Kolmogorov-Smirnov test, which works in the following way. First define the empirical distribution for i.i.d. draws $(X_t)_{t=1}^T$ as:

$$F_T(x) = \frac{1}{T}\sum_{t=1}^{T} 1_{[-\infty,x]}(X_t)$$

The test then calculates the distance:

$$D_T = \sup_x |F_T(x) - F(x)|$$

where $F(x)$ is the true distribution function. Here, $F(x)$ is the standard normal distribution, and X_t are the price increments generated by the dynamical system, $X_t = p_{t+1} - p_t$. Notice I am jointly testing for independence and normality of price increments, which together satisfy the definition of the discrete time Brownian motion. I test the null hypothesis that the data come from a standard normal, and use the cutoff $p = 0.05$ to reject the null hypothesis. It turns out that, when the dynamical system has been run for few periods and T is small, the test consistently fails to reject the null. For higher values of T, the test often rejects the null. In Table 1, I show the values of n and ε for which the test fails to reject the null more often than not at the specified value of T.

	Traders n	Coarseness ε
$T = 5$	15	100
$T = 75$	1000	1
$T = 150$	10^6	10^{-6}

Table 1: Kolmogorov-Smirnov Test

The Motivating Example (with $\varepsilon = 100$) can be easily distinguished, even with the naked eye, from a Brownian motion. The Motivating Example (with $\varepsilon = 1$) cannot be distinguished from a Brownian motion when $T = 75$ iterations are considered. Beyond these examples, the table gives a sense of the rate of convergence of the dynamical system to a Brownian motion.

Appendix F

Consider an economy with $n < \infty$, $\varepsilon > 0$, and the realization r is known to an omniscient econometrician. A deterministic dynamical system is typically classified as chaotic if it has sensitivity to initial conditions. The Lyapunov exponent is the mathematical object that characterizes the rate of separation of infinitesimally close trajectories. For a discrete dynamical system, it is defined:

$$\lambda(X_0) = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \ln|f'(X_t)|$$

where the dynamical system starts at X_0 and is described by $X_{t+1} = f(X_t)$. In the current model, $X_t = p_{t+1} - p_t$ are price increments. A positive Lyapunov exponent is typically taken as an indication of chaos, and numerical tests confirm that the exponent is well above zero for the current model calibrated exactly as in the Motivating Example (with $\varepsilon = 1$). Instead of showing the details of such numerical tests, I visually show what sensitivity means for the dynamical system in Figure 10.

The three price paths were calculated using the same draw of reference points r^j. In other words, the dynamical system was held constant across the three price paths. The only difference between the paths was that one begins at $p_0 = 999.99$, another at $p_0 = 1000$, and the third at $p_0 = 1000.01$. These small differences in starting point are visually indistinguishable at first, but lead to entirely different price paths, hence the terminology "sensitivity to initial conditions." Although the paths in Figure 10 may resemble stochastic processes, non-asymptotic versions of the model with fixed r are chaotic but perfectly deterministic.

Appendix G

The proof of Proposition 2 will proceed in three steps. First I show that demands evaluated at two different prices are independent. Then I pin down the market maker

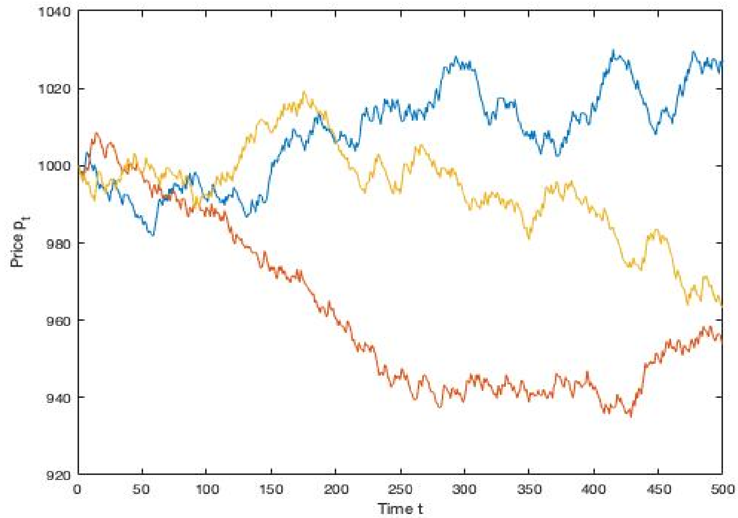

Figure 10: Sensitivity to Initial Conditions

functional form and directly invoke the generalized central limit theorem to show that aggregated demand converges to a stable-α distribution. The desired result then follows almost immediately. For brevity of notation, I suppress the modulo term inside of demand. $h(r^j)$ denotes the density of reference points, and $f_{\mathbb{T}}(p_{qt})$ denotes the density of price increments $p_{qt} \equiv (p_q - p_t)$, which for now I guess exists.[27] By the asymptotic independence of price increments and reference points from Lemma 2, which only required A1-A5, I can write the joint characteristic function of demand evaluated at two different prices p_q and p_t as:

$$\begin{aligned}
\varphi_{n,\varepsilon}(s_1, s_2) &\approx \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \exp\left[is_1 x\left(\frac{r^j - p_q}{\varepsilon}\right) + is_2 x\left(\frac{r^j - p_t}{\varepsilon}\right)\right] h(r^j) dr^j f_{\mathbb{T}}(p_{qt}) dp_{qt} \\
&= \int_{-\infty}^{\infty} \sum_{\ell=-\infty}^{\infty} \int_{p_q+\ell\varepsilon}^{p_q+(\ell+1)\varepsilon} \exp\left[is_1 x\left(\frac{r^j - p_q}{\varepsilon}\right) + is_2 x\left(\frac{r^j - p_t}{\varepsilon}\right)\right] h(r^j) dr^j f_{\mathbb{T}}(p_{qt}) dp_{qt} \\
&= \int_{-\infty}^{\infty} \sum_{\ell=-\infty}^{\infty} \int_{0}^{1} \exp\left[is_1 x(u) + is_2 x\left(u + \frac{p_q - p_t}{\varepsilon}\right)\right] h(\varepsilon[u+\ell] + p_q)\varepsilon du f_{\mathbb{T}}(p_{qt}) dp_{qt} \\
&= \int_{-\infty}^{\infty} \int_{0}^{1} \exp\left[is_1 x(u) + is_2 x\left(u + \frac{p_q - p_t}{\varepsilon}\right)\right] \sum_{\ell=-\infty}^{\infty} h(\varepsilon[u+\ell] + p_q)\varepsilon du f_{\mathbb{T}}(p_{qt}) dp_{qt}
\end{aligned}$$

where I have used a change of variables, $u = \frac{r^j - p_q}{\varepsilon} - \ell$, in the third equality, and dominated

[27]Stable-α distributed price increments do not always have a closed form density. However, stable distributions are absolutely continuous, so a density does exist.

convergence in the fourth equality. Next I apply Fubini's Theorem to flip the order of integration:

$$\int_0^1 \int_{-\infty}^{\infty} \exp\left[is_1 x(u) + is_2 x\left(u + \frac{p_{qt}}{\varepsilon}\right)\right] \sum_{\ell=-\infty}^{\infty} h(\cdot)\varepsilon f_{\mathbb{T}}(p_{qt}) dp_{qt} du$$
$$= \int_0^1 \sum_{k=-\infty}^{\infty} \int_{-\varepsilon u + k\varepsilon}^{-\varepsilon u+(k+1)\varepsilon} \exp\left[is_1 x(u) + is_2 x\left(u + \frac{p_{qt}}{\varepsilon}\right)\right] \sum_{\ell=-\infty}^{\infty} h(\cdot)\varepsilon f_{\mathbb{T}}(p_{qt}) dp_{qt} du$$
$$= \int_0^1 \sum_{k=-\infty}^{\infty} \int_0^1 \exp\left[is_1 x(u) + is_2 x(v)\right] \sum_{\ell=-\infty}^{\infty} h(\varepsilon[v+\ell] + p_t)\varepsilon f_{\mathbb{T}}(\varepsilon[v-u+k])\varepsilon dv du$$
$$= \int_0^1 \int_0^1 \exp[is_1 x(u)] \exp[is_2 x(v)] \sum_{\ell=-\infty}^{\infty} h(x_\ell^*)\varepsilon \sum_{k=\infty}^{\infty} f_{\mathbb{T}}(x_k^*)\varepsilon dv du$$

where, as before, the second equality follows from a change of variables $v = u + \frac{p_{qt}}{\varepsilon} - k$, and the third equality follows from dominated convergence. By the Riemann sum technique, the characteristic function converges to:

$$\varphi_{n,\varepsilon}(s_1, s_2) \to \int_0^1 \exp[is_1 x(u)] \int_0^1 \exp[is_2 x(v)] dv du$$

as $\varepsilon \to 0$. This concludes the proof of independence, which is if and only if $\varphi(s_1, s_2) = \varphi(s_1)\cdot\varphi(s_2)$. In the next step, I consider the market maker constant a_n. Normalized excess demand $a_n z_n(p_t, r)$ cannot converge to a constant, for this would violate A1. Then, by A2, it converges to a nondegenerate distribution. By the Durrett necessary and sufficient condition from Lemma 1, it must then converge to a stable distribution. Furthermore, in the small ε limit:

$$P\left(\left|x\left(\text{mod}\left[\frac{r^j - p_t}{\varepsilon} - \frac{1}{2}, 1\right] - \frac{1}{2}\right)\right| > y\right) = \int_{-\infty}^{\infty} 1_{\{|x\left(\text{mod}\left[\frac{r^j - p_t}{\varepsilon} - \frac{1}{2}, 1\right] - \frac{1}{2}\right)| > y\}} h(r^j) dr^j$$
$$= \sum_{\ell=-\infty}^{\infty} \int_{\varepsilon\ell + p_t}^{\varepsilon(\ell+1)+p_t} 1_{\{|x\left(\text{mod}\left[\frac{r^j - p_t}{\varepsilon} - \frac{1}{2}, 1\right] - \frac{1}{2}\right)| > y\}} h(r^j) dr^j$$
$$= \sum_{\ell=-\infty}^{\infty} \int_0^1 1_{\{|x(\text{mod}(u-0.5,1)-0.5)|>y\}} h(\varepsilon[u+\ell] + p_t)\varepsilon du$$
$$= \int_0^1 1_{\{|x(\text{mod}(u-0.5,1)-0.5)|>y\}} \sum_{\ell=-\infty}^{\infty} h(\varepsilon[u+\ell] + p_t)\varepsilon du$$
$$\to \int_{-\frac{1}{2}}^{\frac{1}{2}} 1_{\{|x(u)|>y\}} du$$

where the third equality follows from a change of variables $u = \frac{r^j - p_t}{\varepsilon} - \ell$, and the fourth equality follows from dominated convergence. By A6′, the last term asymptotically approaches $ky^{-\alpha}$ as $y \to \infty$. By the Mikosch necessary and sufficient conditions, aggregate demand can only converge to a stable-α distribution, where α is given in A6′. Without loss of generality, I set the market maker constant to $a_n = cn^{-1/\alpha}$ which, along with A4, A6′ and A7′, are sufficient to invoke the generalized central limit theorem from Durrett (2017) Theorem 3.8.2, the result of which guarantees convergence of the normalized sum of i.i.d. random variables to the stable-α distribution. The limiting characteristic function of normalized aggregate demand evaluated at any price is given by:[28]

$$\varphi_n(s) \to \exp(-|s|^\alpha)$$

Putting these arguments together, the price increment over multiple time periods $p_{t_2} - p_{t_1}$ has the following characteristic function:

$$\begin{aligned}\varphi_{n,\varepsilon}(s) &= E_{\mathbb{T}}\left[\exp\left(is\sum_{t=t_1}^{t_2-1} a_n \sum_{j=1}^{n} x(p_t, r^j)\right)\right] \\ &\to \exp(-|s|^\alpha)^{t_2-t_1} \\ &= \exp(-(t_2-t_1)|s|^\alpha)\end{aligned}$$

where the asymptotic result invokes both independence of demands at different prices and the generalized limit theorem. □

Appendix H

In this alternative interpretation of the model, I assume traders receive utility and disutility from trade according to:

$$u(r^j, p_t, x_t) = (\tilde{E}_j[p_{t+1}|p_t] - p_t)x_t - L(x_t, \varepsilon)$$

The interpretation is as follows: traders receive disutility from trading with coarse expectations parametrized by ε, and this is captured by the loss term on the right-hand side of the equation above. Traders realize they may be erring by using "rules of thumb."

[28] My assumptions allow simplification of the general form $\varphi(s) = \exp(is\mu - b|s|^\alpha[1 + ik\text{sgn}(s)w_\alpha(s)])$.

Demand is then implicitly defined by:

$$L'(x_t, \varepsilon) = \tilde{E}_j[p_{t+1}|p_t] - p_t$$

where the derivative is taken with respect to the first argument. Now assumption A3, risk-neutrality, can be rewritten as $L'(x_t, \varepsilon) \sim L'(x_t)\varepsilon$. The intent of the assumption is exactly as before: traders must trade a nonzero and finite amount. Since the utility term, $(\tilde{E}_j[p_{t+1}|p_t] - p_t)x_t$, is $\mathcal{O}(\varepsilon)$, the disutility term must also be on the same order if this requirement is to be met. For small ε, I implicitly define demand as before:

$$L'(x_t) = \frac{\tilde{E}_j[p_{t+1}|p_t] - p_t}{\varepsilon} \tag{9}$$

Albeit mathematically equivalent, these two assumptions, $\rho \sim \varepsilon$ and $L'(x_t, \varepsilon) \sim L'(x_t)\varepsilon$, carry different interpretations. The latter says the following: traders with small ε, and hence more accurate expectations, also trade more sensitively to price. This can be seen in condition (9): a small ε in the denominator amplifies price changes. Although expectations are more likely to be correct, when they are incorrect, these traders trade aggressively to exploit the perceived mispricing. This is my preferred interpretation over mean-variance traders, who were introduced only for familiarity. In summary, trader preferences cannot depend on any moments higher than the first, which are undefined when $\alpha < 2$ and prices have heavy tails.

Appendix I

Consider an economy with heterogeneous ε. I evaluate the more tractable Brownian case assuming A1-A7; however, analogous arguments apply in the heavy-tailed case. Say $\varepsilon^i = \varepsilon\beta^i$ for some heterogeneous β^i, and furthermore, say that there are only a finite number I of groups. That is, $n^i(n)$ traders use partitions with coarseness $\varepsilon\beta^i$. I let $\gamma^i = \lim_{n\to\infty} n^i(n)/n > 0$, which denotes the proportion of traders of type i. I suppress the dependence on n for brevity. Normalized aggregate demand is given by:

$$\begin{aligned}\frac{cz_n(p_t, r)}{\sqrt{n}} =& \frac{c}{\sqrt{n}} \sum_{j=1}^{n^1} x^1(p_t, r^j) + ... + \frac{c}{\sqrt{n}} \sum_{j=n-n^I+1}^{n} x^I(p_t, r^j) \\ \approx& c\sqrt{\frac{\gamma^1}{n^1}} \sum_{j=1}^{n^1} x^1(p_t, r^j) + ... + c\sqrt{\frac{\gamma^I}{n^I}} \sum_{j=n-n^I+1}^{n} x^I(p_t, r^j)\end{aligned}$$

where the approximate equality follows from the definition of γ^i. Prices are still non-revealing due to Lemma 2, and by the arguments provided in Appendix G, aggregate demands evaluated at different prices are independent when each group is analyzed separately. Now fix a price p_t. Excess demands for each group are independent from one another by construction, but now they are not guaranteed to be identically distributed; in general, they will not be. Applying the Lindeberg-Lévy theorem to each group, each is normally distributed $N(0, \sigma^2(\beta^i, \gamma^i))$, where the variance depends on (β^i, γ^i). But the sum of independent normally distributed random variables is also normally distributed $N\left(0, \sum_i \sigma^2(\beta^i, \gamma^i)\right)$. The market maker constant c can then be used to normalize this summed variance as desired, and the price process converges in distribution to a discrete time Brownian motion. This additional source of heterogeneity contributes nothing to the non-revealing nature of prices; this last exercise was simply a check for robustness.