Price Distortions under Coarse Reasoning with Frequent Trade*

Jakub Steiner†
CERGE-EI and University of Edinburgh

Colin Stewart‡
University of Toronto

June 29, 2015

Abstract

We study the effect of frequent trading opportunities and categorization on pricing of a risky asset. Frequent opportunities to trade can lead to large distortions in prices if some agents forecast future prices using a simplified model of the world that fails to distinguish between some states. In the limit as the period length vanishes, these distortions take a particular form: the price must be the same in any two states that a positive mass of agents categorize together. Price distortions therefore tend to be large when different agents categorize states in different ways, even if each individual's categorization is not very coarse.

JEL: D53, D84

Keywords: Categorization, bounded rationality, prices

*This paper was formerly circulated under the title "Price Distortions in High-Frequency Markets." We thank the associate editor and anonymous referees, Milo Bianchi, Christian Hellwig, Johannes Hörner, Philippe Jehiel, Willemien Kets, Filip Matějka, John Moore, Stephen Morris, Andreas Park, Rani Spiegler, Muhamet Yildiz, and seminar and conference participants at CERGE-EI, CETC-2013, Columbia, Edinburgh, ESEM-2013, Georgetown, Harvard/MIT, Hebrew U, Humboldt/TU Berlin, INEXC New York 2014, Kellogg, Koç, Oxford, Paris Dauphine, Rice, SAET-2013, São Paulo School of Economics, Tel Aviv, TIGER Forum 2013, Toronto, Toulouse, UCL, and UT-Austin for helpful comments. Ludmila Matýsková provided excellent research assistance. Steiner was supported by Purkyně fellowship of the Czech Academy of Sciences and by GAČR grant 13-34759S. Stewart is grateful to SSHRC for financial support of this research.

†email: jakub.steiner@cerge-ei.cz

‡email: colinbstewart@gmail.com

1 Introduction

Forecasting prices in financial markets is notoriously difficult. Prices depend on so many factors that it is seemingly impossible to identify all of them or to perfectly assess the influence of each one. Agents therefore must use simplified theories of the world to predict prices—theories that are likely to feature some disagreement about what the relevant factors are. How does the use of simplified theories affect the ability of the market to efficiently aggregate information? We present a model to show that, if agents' theories are sufficiently precise, prices tend to be close to rational expectations prices when opportunities to trade are infrequent. However, if the time between trades is small, heterogeneity in agents' theories leads to large distortions even when every individual agent uses a precise (but imperfect) theory.

We study pricing of a single risky asset that is traded at discrete times. The asset pays a flow dividend that depends on the current state, which is publicly observed and evolves according to a Markov process. In choosing prices, agents consider both the current dividend and the resale price in the next period. A key assumption of our model is that, when forming price forecasts, some agents employ a simplified model of the world in which they fail to distinguish among some states. These agents group states into categories and form forecasts in each state that are correct on average for the category containing that state.

We show that whenever two states are categorized together by a positive mass of agents, the price in those two states becomes identical in the limit as the time between trading periods vanishes. This result implies that prices are identical whenever two states are connected by a chain of states along which adjacent states are categorized together (possibly by different agents); prices may be identical even across states with different fundamentals that *no* agent groups together. Thus distortions tend to be large when categorization is heterogeneous. Moreover, if agents' demands take a particular simple form, limiting prices admit a characterization as rational expectations prices associated with a coarsened process—one in which each state corresponds to a set of states in the true process, and dividends and transition probabilities are convex combinations of those in the true process.

Convergence of prices across large sets of states generates a particular pattern of price behavior over time exhibiting sudden large adjustments. Much of the time, prices do not respond to new information, but occasionally there is an overreaction to small changes in fundamentals. These relatively large price jumps occur when the state transitions between two sets of states with differing prices. The net effect on price volatility is ambiguous.

Coarse prices arise from a combination of two effects. First, despite using different theories, agents' expectations of the asset value become identical in each state as the period length vanishes. Second, each agent's expectation becomes constant on each of her own categories. Together, these two effects imply that all agents' expectations are constant on sets of states that are larger than each individual's categories. More precisely, expectations are constant on each element of the finest common coarsening of all agents' categorizations.

Both effects arise when the period length becomes short. In the limit, the per-period dividend becomes negligible and the perceived value of the asset to each agent is based entirely on her forecast of the resale price. This gives rise to the second effect since price forecasts are constant on individual categories.

The first effect—the coordination of individual expectations in each state—is driven by a strong speculative motive that arises when there are frequent trading opportunities. The resale price in the next period is a function of the market forecasts of the price in the following period, which in turn depends on forecasts of the price in the next period, etc. Thus, in a sense, forecasting the resale price can be thought of in terms of forecasting others' forecasts. When agents place a great deal of weight on the resale price, any distinction that an individual makes between two states has little effect on her forecast unless others make the same distinction. If the time between trades is very short, this effect is so strong that a distinction is useful for one's forecast only if *all* other agents make the same distinction. A group of agents failing to distinguish between two states tends to dampen any price difference between those states. This in turn affects the resale price for other agents, further dampening the difference, and thus multiplying the effect. When the period length is short, making the speculative motive strong, this multiplication is powerful enough to drive the prices together.

The coarse pricing result can alternatively be understood in terms of market demand. Learning using categories leads to outcomes as if agents form rational expectations based on incorrect beliefs about the process governing the evolution of states. These incorrect beliefs take a particular form: in addition to transitions that occur under the true process, agents behave as if they have assigned some probability to the state changing to another state in the same category. These beliefs are nonvanishing as the period length vanishes. At the same time, since the current dividend is proportional to the length of time the asset is held, shortening the period length strengthens the speculative motive in the sense that more weight is placed on changes in price relative to dividends. To mitigate this effect and allow markets to clear, prices must become closer together within each category.

Our main result holds under general conditions on the relationship between market prices and individual forecasts. We require only that the market price lie between the highest and lowest expected gain from owning one unit of the asset across all agents (so that, for example, it cannot be that all agents receive an expected gain from buying the asset and an expected loss from selling), and, in addition, that it be bounded away from the lowest expected gain by some fixed convex combination of the two extremes. This assumption is satisfied in the framework of Harrison and Kreps (1978), where all agents are risk neutral, the asset is in fixed supply, and short-selling is limited.

The coordination of individual expectations is most transparent in the special case in which the market price is equal to the average of all agents' expected values of owning the asset, which we study in Section 5. In that case, the model becomes a dynamic variation of a Morris and Shin (2002) beauty contest. This additional structure allows us to explicitly express the steady-state prices as a sum of higher-order expectations of future dividends. When the period length is short, prices are driven by very high orders of expectations, which depend only on those features of individuals' theories that are common to all.[1] Based on the connection to beauty contests, we identify and solve a version of the model that is tractable even when period lengths are non-vanishing and heterogeneous across agents. While prices tend to be less coarse away from the limit than in the limit of vanishing period length, the coarsening effect remains visible. Moreover, if some agents form correct expectations, increasing the frequency of trade only for those agents can amplify distortions in prices due to coarseness of other agents' expectations.

Our main result is stark and should not be taken too literally; the main goal of the paper is to elucidate a mechanism by which trading opportunities that are frequent relative to the arrival

[1]The role of high orders of expectations in our coarse pricing result is akin to their observation that private information has little influence on high orders of expectations (see also Samet 1998).

of information may amplify distortions resulting from imperfect rationality. One could consider an alternative setting in which categorization is replaced with heterogeneous beliefs about transition rates. Although it is difficult to formalize, we believe the same effect is present in that setting. Roughly speaking, given any two states, the effect would lead to prices that are closer together with more frequent trade when the heterogeneity in beliefs about transition rates between those states is relatively high. Section 6 describes some other variations on the model that may go against the constant price result (although again the effect remains).

2 Related Literature

Our focus on categorization places this paper within the burgeoning literature on analogical and similarity-based reasoning (Gilboa and Schmeidler 1995, Jehiel 2005, Jehiel and Samet 2007, Mullainathan, Schwartzstein, and Shleifer 2008, Al-Najjar and Pai 2014). In particular, the starting point of our analysis is a characterization of steady-state outcomes that is closely related to the characterization in Steiner and Stewart (2008); although the setting is different, both papers characterize behavior in terms of equilibrium play in a model with distorted beliefs.

Several other papers have studied the use of coarse theories in asset pricing. Each of those papers considers a fixed trading frequency and therefore has a quite different focus from that of our paper. Most closely related is Eyster and Piccione (2013), who model coarse theories as categorizations of the state space in a very general setting, with steady-state price forecasts formed in the same way as in this paper. They focus on how the composition of theories in the market affects prices and agents' individual performance with a fixed trading frequency. The style investing of Barberis and Shleifer (2003) is related to the categorization in our model. They consider a case with a large fraction of investors who divide assets into a common set of styles over a fixed time horizon, while we focus on the effect of the trading frequency when agents use a variety of categorizations. Another key difference is that Barberis and Shleifer assume that demands are based on relative past performance, while in our model they are based on absolute prices. Similar comments apply to Bianchi and Jehiel (2010), who show that bubbles and crashes can arise when some agents form expectations about price movements that are incorrect but consistent with the average across multiple periods.

A number of earlier papers have highlighted the role of strategic complementarities in amplifying the effect of irrational agents (e.g. Haltiwanger and Waldman 1985, Haltiwanger and Waldman 1989, Fehr and Tyran 2005). Our main result is driven in part by this effect, which is compounded in our model because shorter period lengths strengthen strategic complementarities.

Our results can be understood in terms of higher order expectations about future prices. Among others, Allen, Morris, and Shin (2006) and Bacchetta and Van Wincoop (2008) have highlighted the role of higher order expectations in financial markets with asymmetric information. Since our model is one of complete information, the main thrust of those papers is somewhat orthogonal to the present one.

De Long, Shleifer, Summers, and Waldmann (1990b) show that irrational traders can induce rational agents to behave in a way that destabilizes prices: if irrational traders chase trends, rational traders' demands increase ahead of an upturn in anticipation of greater demand from irrational traders. Our model does not have this feature. Rational traders act as a stabilizing influence, but do not fully stabilize prices. Similarly, in De Long, Shleifer, Summers, and Waldmann (1990a) the same authors study how noise trader risk can distort prices, and show that risk aversion, by limiting the size of positions, can cause rational traders to receive lower expected returns than do noise traders. These papers are related in spirit to our point that irrationality can drive prices away from rational expectations, but the mechanisms are very different.

3 Model and main result

We consider a single asset whose dividend depends on a state $\omega(t)$ drawn from a finite set Ω. The state evolves according to an ergodic continuous-time stationary Markov process with transition rates $q(\omega, \omega')$ from ω to $\omega' \neq \omega$. Trading occurs at discrete times $t = 0, \Delta, 2\Delta, \ldots$. We refer to time $k\Delta$ as period k, and write ω_k for the state $\omega(k\Delta)$ in period k. Sampling the continuous-time process q at times $k\Delta$ gives rise to a discrete-time Markov process (sometimes called the discrete skeleton of q at scale Δ) with transition probabilities $q_\Delta(\omega, \omega')$ from ω to ω'. The state affects the flow dividend $d(\omega_k)$ of the asset, which is paid at a constant rate from time $k\Delta$ to $(k+1)\Delta$.[2]

[2]The assumption that the dividend remains constant between periods instead of changing according to the continuous-time Markov process simplifies the notation but makes no difference for our results. If dividends were instead paid at fixed times (independent of the period length), we conjecture that the results would be similar except that prices would vary depending on the time until the next dividend payment, and would also vary across states for

A continuum of agents of measure one trades the asset in each trading period k. We focus on steady-state prices $P_\Delta : \Omega \longrightarrow \mathbb{R}$ that depend only on the current state. The market price $P_\Delta(\omega)$ is determined by the current dividend and by agents' forecasts of prices in the following period. Agents form these forecasts as follows. Each agent i categorizes states according to a partition of Ω that is fixed across all periods. Letting $\Pi_1, \ldots, \Pi_N$ denote those partitions belonging to a positive measure of agents, we write π_n for the measure of agents using partition Π_n, and refer to the set of agents using this partition as group n. The group of which agent i is a member is denoted $n(i)$. For each state ω, $\Pi(\omega)$ denotes the element of the partition Π containing ω.

Each agent i observes the current state ω_k and forms expectations that are measurable with respect to her categorization $\Pi_{n(i)}$ and are correct on average within each category; that is, given prices $P_\Delta(\omega)$ and any category $C \in \Pi_{n(i)}$, the forecasts E^i satisfy

$$\sum_{\omega \in C} \phi(\omega) E^i \left[P_\Delta \left(\omega_{k+1} \right) \mid \omega_k = \omega \right] = \sum_{\omega \in C} \phi(\omega) E_{q_\Delta(\omega, \omega')} \left[P_\Delta \left(\omega' \right) \right],$$

where ϕ denotes the stationary distribution of states with respect to the true process q.[3] It follows that E^i is identical to the expectation with respect to the *modified process* $m_\Delta^{n(i)}$ given by

$$m_\Delta^n(\omega, \omega') = \sum_{\omega'' \in \Pi_n(\omega)} \phi\left(\omega'' \left| \Pi_n(\omega) \right. \right) q_\Delta(\omega'', \omega'). \tag{1}$$

Note that these expectations differ from those of a fully rational agent with information partition $\Pi_{n(i)}$ (i.e. an agent who understands how price paths differ across states, but cannot observe which state occurs within each category). Such an agent could make additional inferences about future prices based on current and past prices, and on her own history of observations. In contrast, our approach assumes away any such inferences to capture the idea that the agent does not understand that differences within each category could be relevant for forecasting prices.

Following Harrison and Kreps (1978), suppose that all agents are risk neutral and short selling is limited. Given prices P_Δ, the demand of each agent in group n is based on a reservation price $P_\Delta^n(\omega)$ proportional to her net expected profit from holding one unit of the asset for one period;

a vanishingly short time leading up to each dividend payment.

[3]Thus ϕ solves the global balance equations $\sum_{\omega' \neq \omega} \phi(\omega) q(\omega, \omega') = \sum_{\omega' \neq \omega} \phi(\omega') q(\omega', \omega)$.

that is, in each state ω, her *reservation price* is

$$P^n_\Delta(\omega) = \int_{k\Delta}^{(k+1)\Delta} d(\omega)e^{-(t-k\Delta)}dt + e^{-\Delta}E^\omega_n\left[P_\Delta(\omega')\right] = (1-e^{-\Delta})d(\omega) + e^{-\Delta}E^\omega_n\left[P_\Delta(\omega')\right], \quad (2)$$

where $E^\omega_n\left[P_\Delta(\omega')\right]$ denotes the expected resale price $E_{m^n_\Delta(\omega,\cdot)}\left[P_\Delta(\cdot)\right]$ in the next period given that the current state is ω. The agents have a common discount factor normalized to $1/e$. Each agent has zero demand above, arbitrary demand at, and infinite demand below her reservation price. Accordingly, the market-clearing price is

$$P_\Delta(\omega) = \max_n P^n_\Delta(\omega). \quad (3)$$

We say that $\left(P_\Delta(\omega)\right)_\omega$ are *Harrison-Kreps equilibrium prices* if (3) holds for reservation prices P^n_Δ satisfying (2) for each n.

Proposition 1. *Harrison-Kreps equilibrium prices exist and are unique.*

Proof. Let $T : \mathbb{R}^\Omega \longrightarrow \mathbb{R}^\Omega$ be the mapping defined by

$$T(P)(\omega) = \max_n \left\{\left(1-e^{-\Delta}\right)d(\omega) + e^{-\Delta}E^\omega_n\left[P(\omega')\right]\right\};$$

thus T maps each price profile to the profile corresponding to the largest reservation price in each state. The mapping T is increasing in P, and satisfies, for any P and positive constant a,

$$T(P+\overline{a})(\omega) \le T(P)(\omega) + e^{-\Delta}a$$

for every ω, where $\overline{a} \in \mathbb{R}^\Omega$ is the profile with every component equal to a. Therefore, T satisfies Blackwell's sufficient conditions for a contraction, and has a unique fixed point by the Contraction Mapping Theorem.[4] □

Recall that the meet of a collection of partitions is their finest common coarsening and let Π denote the meet of $\Pi_1, \ldots, \Pi_N$. We refer to the elements of Π as *aggregate categories*. Two states ω and ω' lie in the same aggregate category if and only if there exists a sequence $\omega_1, \ldots, \omega_r$ of states

[4]We are grateful to an anonymous referee for suggesting this proof.

such that $\omega = \omega_1$, $\omega' = \omega_r$, and for each $\ell = 1, \dots, r-1$, $\omega_{\ell+1} \in \Pi_n(\omega_\ell)$ for some $n \in \{1, \dots, N\}$. In particular, given two states in different aggregate categories, every agent distinguishes between those two states, but the converse is not true in general: two states in the same aggregate category may be distinguished by all agents.

Focusing on vanishing Δ leads to a striking result: equilibrium prices generally fail to distinguish among states that may differ substantially in terms of fundamentals.

Proposition 2. *As Δ vanishes, prices become constant on each aggregate category; that is, for the Harrison-Kreps equilibrium prices P_Δ,*

$$\lim_{\Delta \to 0} \left(P_\Delta(\omega) - P_\Delta(\omega') \right) = 0$$

whenever $\Pi(\omega) = \Pi(\omega')$.

Proposition 2 follows directly from Proposition 3, which is proved in the following section. The result indicates that, for short period lengths, whenever a positive mass of agents fail to distinguish between two states the market price will be the same in those states. However, that is not all: prices may often be the same in two states even if *no* agent categorizes them together. This is the case if there is an overlapping chain of categories connecting those states. Indeed, if agents do not use the same categories, aggregate categories can be large—potentially leading to large distortions in prices—even if all individual categories are small. Put differently, market prices represent a coarser view of the world than that held by individual market participants.

While the coarsening of prices may suggest that volatility is reduced when Δ is small, the effect is ambiguous in general. Relative to prices when all agents forecast correctly (and prices vary little with Δ), it is easy to construct examples in which average volatility—measured in terms of the standard deviation of single-period logarithmic returns—is higher when some agents use coarse forecasts. Volatility is clearly reduced within each aggregate category. However, when the state moves from one aggregate category to another, the jump in prices may be larger when Δ is smaller, leading to an increase in volatility.

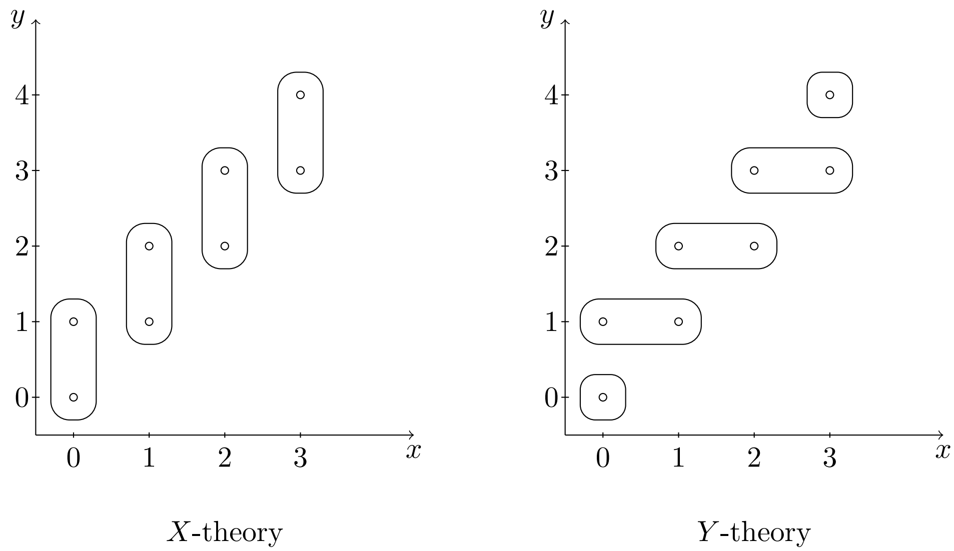

Figure 1: State space and categorizations for Section 3.1 when $S = 3$ for $z = 0$ or 1.

3.1 Example

We now illustrate Proposition 2 in a relatively simple special case in which categorizations are based on dimensions of the state space.

Consider

$$\Omega = \big\{(x, y, z) : x \in \{0, \ldots, S\}, y \in \{x, x+1\}, \text{ and } z \in \{0, 1\}\big\}.$$

Figure 1 depicts a cross-section of the state space for a given value of z. The state follows a continuous-time Markov process with transition rates $q(\omega, \omega') = 1/(4S + 4)$ from state ω to state $\omega' \neq \omega$. This process can be thought of as drawing a state at times that are distributed according to a Poisson process with arrival rate 1, with the new state drawn uniformly from the entire state space. The flow dividend in state $\omega = (x, y, z)$ is

$$d(x, y, z) = (x + y)z/(2S + 1).$$

Thus flow dividends range from 0 to 1 when $z = 1$ and are equal to 0 whenever $z = 0$.

As a benchmark, suppose that all agents' forecasts are based on the true process q_Δ, which we capture by assuming that that all agents use the finest partition of the state space. Then, the

Harrison-Kreps equilibrium prices satisfy

$$P_\Delta(\omega) = (1 - e^{-\Delta})d(\omega) + e^{-\Delta} E_{q_\Delta(\omega,\omega')} \left[P_\Delta(\omega')\right]. \tag{4}$$

The solution to this system of linear equations is given by

$$P_\Delta(\omega) = \frac{4d(\omega) + e^{-\Delta}}{4(1 + e^{-\Delta})}.$$

In this case, as usual, the price in each state is equal to the sum of the expected discounted future dividends.[5]

Consider two variations in which agents categorize some states together. First, suppose that each agent does not understand that the variable y is relevant for forecasting future prices; forecasts depend only on the current values of x and z. This coarse theory, which we call the X-theory, is represented by a partition Π_X of Ω into doubleton categories $\{(x, x, z), (x, x + 1, z)\}$ for $x = 0, \ldots, S$ and $z = 0, 1$, as depicted in Figure 1. Agents' beliefs are coarse but unbiased: their forecasts are measurable with respect to Π_X, but are correct on average within each category. In other words, agents form expectations *as if* they believe transition probabilities between states are given by

$$q^X_\Delta\big((x, y, z), \omega'\big) = \frac{q_\Delta\big((x, x, z), \omega'\big) + q_\Delta\big((x, x + 1, z), \omega'\big)}{2}.$$

The two states within the category are given equal weight because the stationary distribution of the true process q_Δ is uniform.

As in the benchmark case, the prices $P^X_\Delta(\omega)$ given these coarse forecasts satisfy (4) except with the true process q_Δ replaced by q^X_Δ. When S is large, making states within each category similar in terms of fundamentals, the use of a coarse theory has only a small effect on prices: one can show that, $P^X_\Delta(\omega) = P_\Delta(\omega)$ in each state ω with $z = 0$, and

$$\left|P^X_\Delta(\omega) - P_\Delta(\omega)\right| = \frac{e^{-2\Delta}}{2(1 + e^{-\Delta})(2S + 1)} < \frac{1}{4(2S + 1)}$$

in each state ω with $z = 1$. When all agents use the same theory, the magnitude of price distortions

[5] Prices depend on Δ only because of the simplifying assumption that the flow dividend does not change between trading periods.

corresponds to the precision of the theory.

Now suppose that agents are heterogenous in the theories on which they base their forecasts. One group uses the X-theory while the remaining agents use the Y-theory corresponding to the partition Π_Y of Ω into categories $\{(x, y', z') \in \Omega \mid (y', z') = (y, z)\}$ for $y = 0, \ldots, S+1$ and $z = 0, 1$, as depicted in Figure 1.[6] For $W \in \{X, Y\}$, agents using the W-theory form expectations as if they believe the process has transition probabilities $q^W_\Delta(\omega, \omega')$ equal to the average of $q_\Delta(\tilde{\omega}, \omega')$ over all $\tilde{\omega}$ lying in the same element of Π_W as ω. Heterogeneity of theories effectively translates into heterogenous prior beliefs about the underlying process. As in Harrison and Kreps (1978), differences in beliefs motivate agents to trade even though, objectively, there can be no gains from trade.

One can show that, for any Δ, when $z = 1$, the steady-state prices $P^{X,Y}_\Delta(\omega)$ when both theories are present in the market are increasing with the dividend $d(\omega)$. It follows that agents using the X-theory have a higher willingness to pay in states $(x, y, 1)$ along the diagonal (i.e. those with $y = x$), while those using the Y-theory have a higher willingness to pay in the off-diagonal states (those with $y = x + 1$). Hence the prices $P^{X,Y}_\Delta(\omega)$ satisfy (4) except with the true process q_Δ replaced by q^X_Δ if ω lies on the diagonal, and replaced by q^Y_Δ if ω is off the diagonal. When $z = 0$, the two groups have the same willingness to pay.

Figure 2 depicts, for each state ω, the price $P^{X,Y}_\Delta(\omega)$ as a function of Δ when $S = 3$. When the period length is not too short, prices are similar to the benchmark prices $P_\Delta(\omega)$ because both theories are fairly precise in the way they partition the state space. However, as Δ vanishes, prices collapse across states, generating large distortions relative to fundamental values.

Why, when Δ is small, do prices fail to respond to changes in the state even though all agents use theories that are not very coarse? When the trading period is short, reservation prices place little weight on the current dividend relative to the expected resale price. Since any given agent's expectation of the resale price is constant within any of that agent's categories, it must be the case that, in the limit, prices are constant within those categories. To see this, consider a state $\omega = (x, x, 1)$ and let $\omega' = (x, x+1, 1)$. The price in state ω is equal to the reservation price of agents using the X-theory. In the limit, with vanishing weight placed on the dividend, this reservation

[6] Prices in this case would be exactly the same if there was a third group forming rational expectations, since agents in that group would always have a lower willingness to pay than members of one of the two coarse groups.

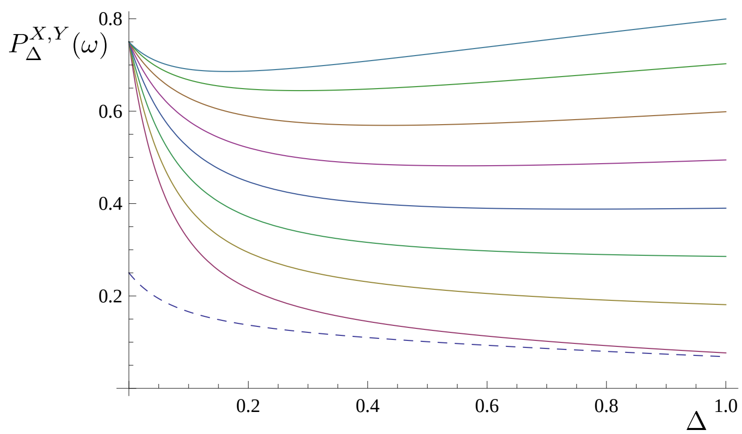

Figure 2: Prices as a function of the period length for $S = 3$ when both the X-theory and the Y-theory are present in the market. Each solid curve depicts the price in a particular state with $z = 1$, while the dashed curve depicts the price in all states with $z = 0$.

price is equal to the average of the prices in states ω and ω'. This implies that the prices in the two states must be equal. Since the same argument applies to every agent and category, prices in the limit must be measurable with respect to *both* individual categorizations, and therefore with respect to the *meet*—the finest common coarsening—of the two individual categorizations. In this case, the finest common coarsening distinguishes between states based only on the value of z. Accordingly, prices in the limit do not respond to changes in x and y, and adjust only when z changes.

4 Generalized aggregation rules

To check the robustness of Proposition 2, we prove a more general result that allows for the market price to be determined in ways that differ from the model of Harrison and Kreps. All elements of the model are the same except for the way in which individual reservation prices are aggregated to determine market prices. An *aggregation rule* ψ is a collection of continuous mappings $\psi_\Delta : (\mathbb{R}^\Omega)^N \longrightarrow \mathbb{R}^\Omega$, one for each $\Delta > 0$, from profiles $(P^n)_n$ of reservation prices to market prices P. For example, (3) defines an aggregation rule. Given Δ, we say that prices P_Δ are *equilibrium prices for aggregation rule* ψ if $\psi_\Delta(P_\Delta^1, \ldots, P_\Delta^N) = P_\Delta$ when, for each n, the reservation prices P_Δ^n satisfy (2) given P_Δ. Since ψ_Δ is continuous, so is its composition with the continuous mapping described by (2) from prices P_Δ to the profile of reservation prices. Brouwer's Fixed-Point Theorem therefore implies that equilibrium prices exist within the compact set bounded by the minimal and maximal dividends.

We say that an aggregation rule ψ is *regular* if there exists $\mu \in (0,1]$ such that for all ω and Δ, the price $P_\Delta(\omega) = \psi_\Delta(P^1_\Delta, \ldots, P^N_\Delta)(\omega)$ satisfies

$$(1-\mu)\min_n P^n_\Delta(\omega) + \mu \max_n P^n_\Delta(\omega) \leq P_\Delta(\omega) \leq \max_n P^n_\Delta(\omega). \tag{5}$$

Roughly speaking, regularity requires that (i) the price is never so high as to make every agent receive an expected loss from buying the asset, and (ii) if agents differ in their expectations of the asset's value, the price is higher than the lowest of the expected values (by at least some fixed amount relative to the difference in expectations). The Harrison and Kreps (1978) aggregation rule in (3) is regular with $\mu = 1$. In particular, as in Section 3.1, this abstract formulation implicitly allows for investors who have long horizons. Another regular aggregation rule, which we study in Section 5, is the one that maps reservation prices to their average, weighted according to the size of each group. This aggregation rule arises in standard rational expectations models with CARA utilities and normal distributions (see, e.g., Allen, Morris, and Shin 2006).

Lemma 1. *Suppose that, for some regular aggregation rule, P_Δ are equilibrium prices for each $\Delta > 0$. Then, as $\Delta \to 0$, all agents' reservation prices become identical in each state; that is,*

$$\lim_{\Delta \to 0} (P^n_\Delta(\omega) - P^m_\Delta(\omega)) = 0$$

for all ω, m, and n.

The idea of the proof can be illustrated by considering a simple beauty contest in which each agent chooses a real number with the goal of matching the average of all agents' choices. In equilibrium, agents must perfectly coordinate on the same number, for otherwise, the agent with the minimal action would rather increase her action toward the average action. Similarly, in our setting, if agents differ in their reservation prices, then there exists a state-agent pair such that the agent's reservation price is minimal. The proof shows that the agent would increase her reservation price in that state, establishing a contradiction.

Proof of Lemma 1. Fix a regular aggregation rule and suppose for contradiction that the statement does not hold. Then we can find a sequence Δ_ℓ converging to 0 such that the limits $P^n(\omega) =$

$\lim_\ell P^n_{\Delta_\ell}(\omega)$, and $P(\omega) = \lim_\ell P_{\Delta_\ell}(\omega)$ exist for all ω and n, and $P^n(\omega) \neq P^m(\omega)$ for some n, m, and ω.

Let R be the set of (ω, n) for which there exists m such that $P^n(\omega) \neq P^m(\omega)$. Consider a pair

$$(\omega^*, n^*) \in \underset{(\omega,n)\in R}{\arg\min} P^n(\omega). \tag{6}$$

By (5), we have

$$P(\omega^*) \geq (1-\mu)P^{n^*}(\omega^*) + \mu \max_n P^n(\omega^*) > P^{n^*}(\omega^*).$$

The strict inequality follows from the fact that $(\omega^*, n^*) \in R$.

Notice that $P^{n^*}(\omega) = P^{n^*}(\omega^*)$ for all $\omega \in \Pi_{n^*}(\omega^*)$ since, for a given group, the reservation prices differ between two states in the same category only through the difference in dividends, which is of order Δ.

In addition, we have

$$P^n(\omega) \geq P^{n^*}(\omega^*) \tag{7}$$

for all $\omega \in \Pi_{n^*}(\omega^*)$ and all n because either $(\omega, n) \notin R$, in which case $P^n(\omega) = P^{n^*}(\omega) = P^{n^*}(\omega^*)$, or $(\omega, n) \in R$, in which case $P^n(\omega) \geq P^{n^*}(\omega^*)$ by (6).

Inequalities (5) and (7) together imply that the market price $P(\omega)$ is at least $P^{n^*}(\omega^*)$ in all states $\omega \in \Pi_{n^*}(\omega^*)$. We have shown that $P(\omega^*) > P^{n^*}(\omega^*)$. Finally, in state ω^*, the probability $m^{n^*}_\Delta(\omega^*, \omega^*)$ that agents in group n^* assign to the state in the next period being ω^* converges to a positive limit $\phi(\omega^* \mid \Pi_{n^*}(\omega^*))$. Therefore, the limit reservation price of group n^* at ω^* must exceed $P^{n^*}(\omega^*)$, which establishes the desired contradiction. $\square$

The following result directly extends Proposition 2 to all regular aggregation rules.

Proposition 3. *For any regular aggregation rule, equilibrium prices P_Δ satisfy*

$$\lim_{\Delta\to 0} \left(P_\Delta(\omega) - P_\Delta(\omega')\right) = 0$$

whenever $\Pi(\omega) = \Pi(\omega')$.

Proof. From (2), limit reservation prices are constant on individual categories, that is, $P^n(\omega) = P^n(\omega')$ whenever $\Pi_n(\omega) = \Pi_n(\omega')$. Lemma 1 establishes that limit reservation prices are also constant across groups in each state. Hence the limit reservation prices are measurable with respect to the aggregate categorization Π. Since $P_\Delta(\omega) \in [\min_n P^n_\Delta(\omega), \max_n P^n_\Delta(\omega)]$, the limit market price is itself measurable with respect to Π. $\square$

5 Beauty contests

This section clarifies the connection between our model and that of Morris and Shin (2002), and builds on it to better understand the properties of equilibrium prices in and away from the limit. We show that, for a particular aggregation rule, our model can be reinterpreted as a beauty contest game played by agents with coarse expectations. Following the insights of Morris and Shin, we can interpret the equilibrium price in terms of higher-order expectations of dividends, and characterize prices in specific settings.

To simplify notation, we often omit the subscript Δ in this section. We focus on the *linear aggregation rule* defined by

$$P(\omega) = \sum_n \pi_n P^n(\omega) = (1 - e^{-\Delta})d(\omega) + e^{-\Delta} \sum_n \pi_n E^\omega_n \left[P(\omega')\right]. \tag{8}$$

This aggregation rule gives the market-clearing prices if the asset has zero net supply, and agents from group n with reservation price $P^n(\omega)$ demand a quantity proportional to $P^n(\omega) - P(\omega)$ (where the reservation prices are formed according to (2)). Although we do not provide a micro-foundation, such demands arise from utility optimization in a class of models with CARA utility and normally distributed shocks.

This case is essentially equivalent to a dynamic beauty contest game in which each agent seeks to match her action with the current state and the average action in the following period. More precisely, the following abstract game has a unique equilibrium that coincides with equilibrium prices under the linear aggregation rule. Each agent i from the same set of agents as in the main model chooses actions $P^i_k \in \mathbb{R}$ in periods $k = 0, 1, 2, \ldots$ to maximize the flow payoffs

$$-(1 - e^{-\Delta}) \left(P^i_k - d(\omega_k)\right)^2 - e^{-\Delta} \left(P^i_k - P_{k+1}\right)^2,$$

where $P_k = \int P_k^i di$ is the average action. The agents have coarse beliefs as in the main model; that is, agents from group n hold beliefs about the state in the next period described by (1). Note that, unlike the reservation price in the main model, in this reinterpretation P_k^i is an action chosen by agent i.

It is straightforward to verify that the reservation prices $P^n(\omega)$ together with equilibrium prices $P(\omega)$ for the linear aggregation rule correspond precisely to stationary equilibria of this dynamic beauty contest (that is, equilibria in which each agent's action in each period k depends only on the current state ω_k). To see this, notice that the best response of an agent from group n in the dynamic beauty contest is $(1 - e^{-\Delta})d(\omega_k) + e^{-\Delta} E_n^{\omega_k}[P_{k+1}]$, which is equivalent to the defining equation (2) for the reservation price.

As in Morris and Shin (2002), one can characterize the equilibrium outcome as a sum of higher-order expectations. Iterating (8) leads to

$$P(\omega) = \left(1 - e^{-\Delta}\right) d(\omega) + \left(1 - e^{-\Delta}\right) \sum_{k=1}^{\infty} e^{-k\Delta} \left(\overline{E}^{\omega}\right)^k [d(\cdot)], \tag{9}$$

where $\overline{E}^{\omega} = \sum_n \pi_n E_n^{\omega}$ is the population average expectation and $\left(\overline{E}^{\omega}\right)^k$ is its k-fold iteration. As the period length Δ vanishes, increasing weight in (9) is placed on higher-order expectations (i.e. on higher values of k).

Whereas first-order expectations are based on agents' individual beliefs about the dividend in the next period, high-order expectations are based on a common understanding of the underlying process shared by the whole population. Just as high-order expectations in Morris and Shin (2002) converge to the expectation conditional only on public information, high-order expectations in our model converge to the expectation conditional only on aggregate categories. Since low-order expectations receive little weight when Δ is small, it follows that reservation prices converge across groups.

The dependence of prices on high-order expectations may appear to be somewhat inconsistent with the bounded rationality of agents in our model insofar as computing those expectations requires a sophisticated understanding of other agents' categorizations. In the supplementary appendix, we show that such sophistication is not necessary: the same expectations arise when agents forecast using a simple naïve rule. More precisely, equilibrium prices for the linear aggregation rule arise almost surely as the long-run outcome of a process in which agents forecast future prices using past data from all states that they categorize together with the current state. In period k, each agent i forms a forecast $E^i[P_{k+1}]$ of the price in period $k+1$ according to

$$E^i[P_{k+1}] = \frac{\sum_{s<k-1:\omega_s \in \Pi_{n(i)}(\omega_k)} p_{s+1}}{\sum_{s<k-1:\omega_s \in \Pi_{n(i)}(\omega_k)} 1}$$

whenever the denominator is nonzero (otherwise the forecast is some arbitrary fixed number), where p_s denotes the market price in period s. Thus the price forecast $E^i[P_{k+1}]$ is formed by averaging all prices that occurred in periods immediately following those in which the state was in the same category as the current one (according to $\Pi_{n(i)}$). In addition, the supplement allows for the presence of rational agents who know all parameters of the model, including other agents' forecasting rules. As in the steady-state analysis, the long-run behavior of these rational agents is identical to that of agents who forecast using the finest partition of the state space. This convergence result indicates that agents do not need to forecast others' forecasts for prices to satisfy (9). We view the expansion in terms of higher-order expectations simply as an analytical tool for analyzing the equilibrium.

5.1 Prices in the limit

Since the linear aggregation rule is regular, Proposition 3 implies that prices become constant within each aggregate category as the period length vanishes. This section sheds additional light on the exact form that prices take by showing that they approach rational expectations prices with respect to a process that is coarser than the true process, with each aggregate category playing the role of an individual state.

Define the *coarse dividend function* $\overline{d} : \Pi \longrightarrow \mathbb{R}$ by

$$\overline{d}(C) = \sum_{\omega \in C} \phi(\omega|C)d(\omega)$$

for each $C \in \Pi$. That is, the coarse dividend is obtained by averaging dividends on each aggregate category with weights determined by the stationary distribution ϕ. Similarly, define the *coarse* *process* $\overline{q}$ to be the continuous-time Markov process on the state space Π with transition rates

$$\overline{q}(C, C') = \sum_{\omega \in C} \phi(\omega|C) \sum_{\omega' \in C'} q(\omega, \omega')$$

for all distinct states $C, C' \in \Pi$. That is, the coarse process $\overline{q}$ is obtained by averaging the true process q across each aggregate category with weights determined by the stationary distribution ϕ.

Define the *rational expectations prices* with respect to a continuous-time Markov process $\tilde{q}$ on Π and a dividend function $\tilde{d} : \Pi \longrightarrow \mathbb{R}$ to be the unique solution P to the system of equations

$$P(C) = \frac{\tilde{d}(C) + \sum_{C' \neq C} \tilde{q}(C, C') P(C')}{1 + \sum_{C' \neq C} \tilde{q}(C, C')}.$$

Proposition 4. *For each* $\Delta > 0$*, let* P_Δ *be an equilibrium price vector for the linear aggregation rule. Then* $\lim_{\Delta \to 0} P_\Delta(\omega)$ *is equal to the rational expectations price with respect to* $\overline{q}$ *and* $\overline{d}$ *in state* $\Pi(\omega)$*.*

The proof is in the appendix.

The rational expectations prices that arise in the market are derived from a dividend process that is typically much coarser than the true underlying process, and also much coarser than the theories held by individual agents. In the limit, prices are as if all agents' relatively precise individual theories were replaced by a coarser "market theory."

5.2 Prices away from the limit

This section studies a particularly tractable variation of our model in which categorization is replaced with weighting according to similarity. When states and similarity weights follow a normal distribution, we can apply the techniques of Morris and Shin (2002) to explicitly characterize prices for non-vanishing Δ, and for cases in which agents differ in their trading frequencies.

The state $\omega_t \in \mathbb{R}$ follows a continuous-time Markov process with transition times that have a state-independent arrival rate of 1. At each transition time, the next state is drawn from the standard normal distribution, independent of the current state. In the state ω, the asset pays flow dividend $d(\omega) = \omega$.

As before, a continuum of agents is divided into N groups of sizes π_n, with $\sum_n \pi_n = 1$. Agents from group n are endowed with a similarity function $\sigma_n : \mathbb{R}^2 \longrightarrow \mathbb{R}_+$. We interpret $\sigma_n(\omega, \omega')$ as capturing group n's perception of the degree of similarity between states ω and ω'. The categorization model can be thought of as a case in which σ_n takes on values of 0 and 1. Along the same lines as (1), we assume that, in any state ω_t at time t, each agent forecasts prices as if she forms expectations about the state $\omega_{t+\Delta}$ according to a convex combination of the objective expectations across those states $\tilde{\omega}$ that the agent perceives as similar to ω_t. More precisely, we define the expectation of members of group n to be

$$E_n[\omega_{t+\Delta} \mid \omega_t] = \frac{\int_{-\infty}^{\infty} \phi(\tilde{\omega})\sigma_n(\omega_t, \tilde{\omega}) E\left[\omega_{t+\Delta} \mid \tilde{\omega}\right] d\tilde{\omega}}{\int_{-\infty}^{\infty} \phi(\tilde{\omega})\sigma_n(\omega_t, \tilde{\omega}) d\tilde{\omega}},$$

where ϕ is the standard normal density and $E[\omega_{t+\Delta} \mid \tilde{\omega}]$ denotes the expectation of $\omega_{t+\Delta}$ according to the true process conditional on $\omega_t = \tilde{\omega}$. According to this formulation, states $\tilde{\omega}$ are given relatively large weight in expectation formation if they are similar to the current state, and if they are objectively more likely to occur. As in the case of categorization, we show in the supplementary appendix that expectations of this form arise as the long-run outcome of a simple similarity-based learning process.

We allow for groups to differ in their trading frequency. Members of group n trade at intervals Δ_n, and form reservation prices

$$P^n(\omega_t) = \int_t^{t+\Delta_n} e^{-(t'-t)} E[d(\omega_{t'}) \mid \omega_t] dt' + e^{-\Delta_n} E_n\left[P\left(\omega_{t+\Delta_n}\right) \mid \omega_t\right].$$

Note that, in this formulation, similarity affects only the expectations of future prices; expected dividends within each trading period are based on the objective process. This is a conservative assumption in that it tends to weaken price distortions relative to the alternative in which dividend forecasts are also coarse. We restrict attention to the linear aggregation rule $P(\omega) = \sum_n \pi_n P^n(\omega)$ and to prices of the form $P(\omega) = r\omega$.[7]

Equilibrium prices can be computed explicitly when each group n either uses a similarity func-

[7]The constant weights π_n are based on an implicit assumption that members of each group do not all trade simultaneously; rather, at each trading time, a fixed fraction of each group participates, and π_n represents that fraction's share of the market (as opposed to the size of the group as a whole).

tion of the form

$$\sigma_n(\omega, \omega') = \exp\left[-\frac{(\omega - \omega')^2}{2\rho_n^2}\right],$$

where $\rho_n > 0$, or forms rational expectations, in which case we let $\rho_n = 0$. The parameter ρ_n can be interpreted as the degree of coarseness of the expectations of group n. We assume that $\rho_n > 0$ for some n.

Proposition 5. *Prices $P(\omega) = r\omega$ with*

$$r = \frac{1}{2}\frac{1 - \sum_n \pi_n e^{-2\Delta_n}}{1 - \sum_n \pi_n \frac{e^{-2\Delta_n}}{1+\rho_n^2}} \tag{10}$$

form an equilibrium for the linear aggregation rule. Prices become less responsive to the state if any group forms coarser expectations or trades more frequently; that is, for each n, r is decreasing in ρ_n and increasing in Δ_n. Moreover, r tends to 0 in the limit as Δ_n vanishes for every n.

This proposition highlights how the degree of coarseness in pricing is directly linked to the coarseness and trading frequency of each group; as either of these increases for any group, prices respond less to fundamentals. As the time between trades vanishes for each group (possibly at different rates), we obtain coarse pricing, just as in Propositions 2 and 3; in this case, prices are completely unresponsive to fundamentals. At the other extreme, as each Δ_n grows large, prices approach rational expectations prices (which is not surprising in light of the assumption that dividend forecasts are correct between trading periods).

Reducing the period length in our model is analogous to increasing the relative weight on matching the average action in Morris and Shin (2002). In their model, as the relative weight on matching the state vanishes, agents ignore their private information and choose an action equal to the public signal. In contrast, in the corresponding limit in our model, agents also ignore public information and choose a constant action independent of the state.

The speed with which prices converge as the horizons Δ_n vanish depends on the parameters. For example, suppose there are two groups of equal size with identical time Δ between trades. Figure 3 plots r as a function of Δ for three different profiles (ρ_1, ρ_2) of coarseness parameters. As Proposition 5 implies, convergence is relatively slow when the similarity functions are relatively narrow. The numerical results indicate that price distortions can be non-negligible even for parameters that are

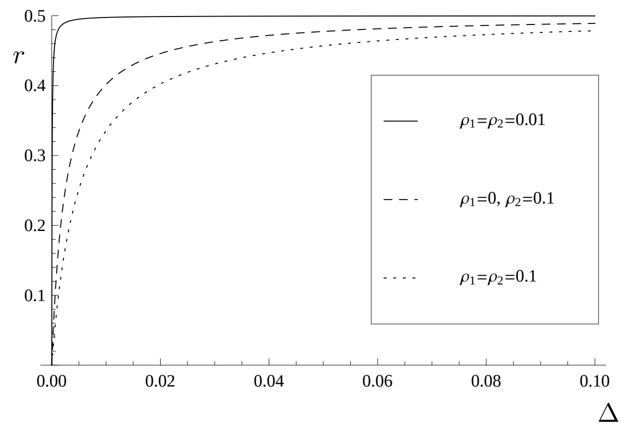

Figure 3: The equilibrium price coefficient r as a function of Δ for two equal-sized groups with $\Delta_1 = \Delta_2 = \Delta$. Lower values of r correspond to coarser prices, and higher values of ρ_1 and ρ_2 correspond to coarser expectations.

moderate relative to the speed of evolution and standard deviation of the state, such as $\rho_1 = 0$, $\rho_2 = 0.1$, and $\Delta = 0.02$ (recall that the standard deviation of the stationary distribution and the average time between transitions are both equal to 1).

One might think that more sophisticated traders—those with smaller values of ρ_n—tend also to be the ones with more frequent opportunities to trade. If some traders using coarse similarity have a fixed trading frequency, what is the effect on prices of increasing the frequency of trade of other, rational traders? Proposition 5 implies that doing so reduces the value of r, amplifying distortions in prices. Figure 4 illustrates this effect by depicting the price coefficient r as a function of the time between trading periods for a group forming rational expectations that is equal in size to a second group using a nontrivial similarity function (i.e., $\pi_1 = \pi_2 = 1/2$). In this case, distortions are smaller than if Δ_n vanishes for both groups, but the effect is present nonetheless. Prices remain close to rational expectations prices if the coarse group trades infrequently. However, when that is not the case, increasing the frequency of trade of the group that forms rational expectations can generate sizable distortions.

The main focus of our results is on the limit in which period lengths vanish while the degree of coarseness is fixed. If instead we fix period lengths and allow the similarity parameters ρ_n to vanish, then, by (10), prices converge to those under rational expectations. One might expect that, as trade becomes more frequent, traders' refine their understanding of the market and reduce the

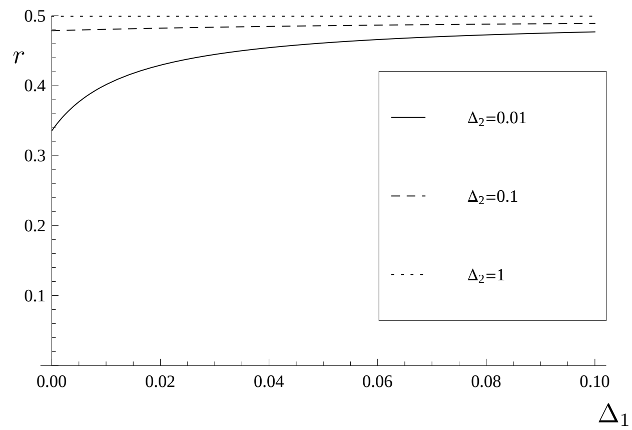

Figure 4: The equilibrium price coefficient r for two groups of equal size as a function of the period length Δ_1 for group 1. Group 1 forms rational expectations while group 2 forms coarse expectations with parameter $\rho_2 = 0.1$. The period length Δ_2 for group 2 is fixed. Lower values of r correspond to coarser prices.

coarseness of their theories. In that case, whether prices become coarse depends on the relative sizes of the period lengths and the similarity parameters. To see this, note that when every ρ_n and Δ_n is close to 0, (10) can be approximated by

$$r \approx \frac{\sum_n \pi_n \Delta_n}{2\sum_n \pi_n \Delta_n + \sum_n \pi_n \rho_n^2}.$$

Thus if decreasing period lengths are accompanied by an increase in traders' sophistication, prices may be less coarse than our main result suggests.

6 Concluding remarks

In order to highlight the effect of trade being frequent relative to information flows, we have focused on a simple tractable model in which the resulting prices take a stark form. A number of natural modifications of the model may moderate the effect while retaining significant price distortions with frequent trade. In this section, we speculate about the consequences of various extensions and modifications to the main model.

Traders in our model live forever and have no limits on losses. Since traders using coarse models tend to lose money against traders using refinements of those models, forcing traders to exit once reaching a given loss threshold could drive all agents out of the market except those who use the finest possible categorization (if any such agents exist), thereby eliminating price distortions. However, since our results hold independent of the fractions of agents using various partitions, we conjecture that they also hold as long as there is continual entry of a nonvanishing mass of new traders using coarse categories.

Risk aversion of traders can limit the size of the positions they take. Since agents who form rational expectations perceive the risk of state transitions over a short horizon to be very low, risk aversion tends to limit their positions less than those of agents who use coarse categorization. Thus we expect that increasing risk aversion should reduce price distortions.

Agents in our model form forecasts that are measurable with respect to their own categorization and correct on average within each category. More generally, our main result holds as long as agents' forecasts are bounded by prices at states within the current category and place at least some weight on each other state in the category. For example, one could consider a model in which agents form forecasts by averaging past prices at states in the same category as the current one, with the weight applied to each past state diminishing in the length of time since it occurred. While such a model is not amenable to steady-state analysis, we conjecture that our main result would extend to long-run prices in the sense that, if the period length is short, prices at any given point in time will eventually be approximately constant within each aggregate category (although the price associated with a given aggregate category may vary over time depending on the history). On the other hand, the result would not hold if forecasts were based on a fixed number of the most recent observations in the current category, since the weight assigned to different states would then vanish along with the period length.

For the sake of parsimony, we have assumed that agents employ categories that are fixed across time. Alternatively, one might expect agents to adjust their categories as they learn the correct model. If learning leads to successive refinements in categorization toward the finest categorization, then our results may not hold in the long-run. As with limits on losses, however, we expect that continual entry of agents using coarse categories would suffice to generate persistent price distortions.

Appendices

A Proof of Proposition 4

Let $m_\Delta(\omega,\omega') = \sum_n \pi_n m^n_\Delta(\omega,\omega')$ denote the population-average belief about transition probabilities. Given an aggregate category C, for each Δ let $\tilde{m}_\Delta$ denote the transition probabilities of the restriction of m_Δ to C, that is, the probabilities defined on $C \times C$ obtained by conditioning $m_\Delta(\omega,\omega')$ on $\omega' \in C$. Let ϕ^Δ_C denote the stationary distribution of $\tilde{m}_\Delta$.

Lemma 2. *There exists $K(\eta,\Delta)$ such that, for each $\eta > 0$,*

1. *for each $\Delta > 0$ and $\omega \in C$,*

$$\frac{1}{K(\eta,\Delta)} \sum_{k=0}^{K(\eta,\Delta)-1} \|\tilde{m}^k_\Delta(\omega,\cdot) - \phi^C_\Delta\| < \eta,$$

   where $\|\cdot\|$ is the 1-norm and $\tilde{m}^k_\Delta$ are the transition probabilities for k steps of $\tilde{m}_\Delta$; and

2. *$K(\eta,\Delta)\Delta \to 0$ as $\Delta \to 0$.*

Proof. We claim that, for each $\eta > 0$ and $\omega \in C$, there exists K_0 such that

$$\|\tilde{m}^k_\Delta(\omega,\cdot) - \phi^C_\Delta\| < \eta/2$$

for every $k \geq K_0$, and $K_0\Delta \to 0$ as $\Delta \to 0$. Since $\|\tilde{m}^k_\Delta(\omega,\cdot) - \phi^C_\Delta\| \leq 2$ for every k, taking $K(\eta,\Delta) = 4K_0/\eta$ proves the result.

We will show that there exists ε_Δ such that (i)

$$\|\tilde{m}^k_\Delta(\omega,\cdot) - \phi^C_\Delta\| \leq 2(1-\varepsilon_\Delta)^{k-1} \tag{11}$$

for every k and ω, and (ii) $\lim_{\Delta\to 0} \varepsilon_\Delta/\Delta = \infty$. Then, letting

$$K_0(\eta,\Delta) = 2 + \frac{\log(\eta/4)}{\log(1-\varepsilon_\Delta)},$$

straightforward algebraic manipulation shows that $2(1-\varepsilon_\Delta)^{k-1} < \eta/2$ for every $k \geq K_0$, as needed.

Moreover, $K_0(\eta, \Delta)\Delta \to 0$ as $\Delta \to 0$ since $\lim_{\Delta\to 0} \Delta/\log(1-\varepsilon_\Delta) \to 0$ by (ii).

Existence of ε_Δ satisfying (i) and (ii) follows from Corollary 1.2 of Hartfiel (1998). The corollary implies that if there exist $\delta_\Delta \geq 0$ and L such that $\tilde{m}^L_\Delta(\omega, \omega') \geq \delta_\Delta$ for all ω and ω', then

$$\|\tilde{m}^k_\Delta(\omega, \cdot) - \phi^C_\Delta\| \leq 2(1-\delta_\Delta)^{\frac{k}{L}-1}$$

for every $k > 0$.

Inequality (11) follows by taking $L = |C|$. Notice that $\tilde{m}^{|C|}_\Delta(\omega, \omega')$ is bounded from below by a constant δ independent of Δ. Thus we can choose ε_Δ to be $1-(1-\delta)^{\frac{1}{L}}$.

$\square$

Proof of Proposition 4. Notice that equations (8) and (1) imply

$$P(\omega) = d(\omega)(1-e^{-\Delta}) + e^{-\Delta} E_{m_\Delta(\omega,\omega')}[P(\omega')]$$

Let $\omega \in C$ and rewrite the last equation as

$$\begin{aligned} P(\omega) = d(\omega)(1-e^{-\Delta}) + e^{-\Delta} &\underbrace{E_{m_\Delta(\omega,\omega')}[P(\omega')|\omega' \notin C]}_{f(\omega,\Delta)} \underbrace{\Pr_{m_\Delta(\omega,\omega')}[\omega' \notin C]}_{\varepsilon_\omega\Delta + O(\Delta^2)} \\ &+ e^{-\Delta} \underbrace{(1-\Pr_{m_\Delta(\omega,\omega')}[\omega' \notin C])}_{1-\varepsilon_\omega\Delta+O(\Delta^2)} E_{\tilde{m}_\Delta(\omega,\omega')}[P(\omega')] \end{aligned}$$

where $\varepsilon_\omega = \lim_{\Delta\to 0} \Pr_{m_\Delta(\omega,\omega')}[\omega' \notin C]/\Delta$.[8]

Using the approximation $d(\omega)(1-e^{-\Delta}) = d(\omega)\Delta + O(\Delta^2)$, the last equation can be rewritten as

$$P(\omega) = d(\omega)\Delta + f(\omega,\Delta)\varepsilon_\omega\Delta + e^{-\Delta}(1-\varepsilon_\omega\Delta)E_{\tilde{m}_\Delta(\omega,\omega')}[P(\omega')] + O(\Delta^2).$$

This can be interpreted as the pricing equation of a process in which the asset pays a dividend $d(\omega)\Delta$, with some probability $\varepsilon_\omega\Delta$ the process terminates giving a final payoff $f(\omega, \Delta)$, and with the remaining probability $1-\varepsilon_\omega\Delta$ the process continues to the next trading period, in which the state will be $\omega' \in C$.

[8] The fact that $\Pr_{m_\Delta(\omega,\omega')}[\omega' \notin C] = \varepsilon_\omega\Delta + O(\Delta^2)$ holds because $m_\Delta(\omega,\omega') = q_\Delta(\omega,\omega')$ whenever ω and ω' lie in different aggregate categories.

Iterating the last equation for K periods gives

$$P(\omega) = \sum_{k=0}^{K-1} e^{-k\Delta} E\left[\left(\prod_{k'=0}^{k-1}(1-\varepsilon_{\omega_{k'}}\Delta)\right)(d(\omega_k)\Delta + f(\omega_k,\Delta)\varepsilon_{\omega_k}\Delta)\right] \\ + e^{-K\Delta} E\left[\left(\prod_{k'=0}^{K-1}(1-\varepsilon_{\omega_{k'}}\Delta)\right)P(\omega_K)\right] + O(K\Delta^2),$$

where ω_k is the state in period k of the Markov process $\tilde{m}_\Delta$ on C starting from $\omega_0 = \omega$. For any K, we may rewrite the last equation as

$$\frac{1}{K}P(\omega) = \frac{1}{K}\sum_{k=0}^{K-1} E\left[(d(\omega_k) + f(\omega_k,\Delta)\varepsilon_{\omega_k})\,\Delta\right] \\ + \frac{1}{K}(1-K\Delta)E\left[\left(1-\sum_{k'=0}^{K-1}\varepsilon_{\omega_{k'}}\Delta\right)P(\omega_K)\right] + O(K\Delta^2). \quad (12)$$

Taking $K = K(\eta,\Delta)$ from Lemma 2 and K_0 as in the proof of the lemma, we have

$$\begin{aligned} E\left[\left(1-\sum_{k'=0}^{K-1}\varepsilon_{\omega_{k'}}\Delta\right)P(\omega_K)\right] &= E\left[\left(1-\sum_{k'=0}^{K-K_0-1}\varepsilon_{\omega_{k'}}\Delta - \sum_{k'=K-K_0}^{K-1}\varepsilon_{\omega_{k'}}\Delta\right)P(\omega_K)\right] \\ &= E\left[\left(1-\sum_{k'=0}^{K-K_0-1}\varepsilon_{\omega_{k'}}\Delta\right)E\left[P(\omega_K)\mid\omega_{K-K_0}\right]\right] + O(\eta K\Delta) \\ &= E\left[\left(1-\sum_{k'=0}^{K-1}\varepsilon_{\omega_{k'}}\Delta\right)E\left[P(\omega_K)\mid\omega_{K-K_0}\right]\right] + O(\eta K\Delta). \end{aligned}$$

Substituting into (12) and applying Lemma 2 gives

$$\frac{1}{K(\eta,\Delta)}P(\omega) = \left(d^C + [f\varepsilon]^C\right)\Delta + (1-K(\eta,\Delta)\Delta)\left(1-K(\eta,\Delta)\varepsilon^C\Delta\right)\frac{1}{K(\eta,\Delta)}P^C \\ + O(\eta/K(\eta,\Delta) + \eta\Delta + K(\eta,\Delta)\Delta^2), \quad (13)$$

where x^C denotes the average of $x(\omega)$ with respect to the stationary distribution of the process

$\tilde{m}_\Delta$.[9] Rearranging yields

$$P(\omega) - P^C = \left(d^C + [f\varepsilon]^C\right) K(\eta,\Delta)\Delta + \left(-K(\eta,\Delta) - K(\eta,\Delta)\varepsilon^C + \varepsilon^C K(\eta,\Delta)^2\Delta\right)\Delta P^C$$
$$+ O(\eta + K(\eta,\Delta)\eta\Delta + K(\eta,\Delta)^2\Delta^2).$$

Since $K(\eta,\Delta)\Delta \to 0$ as $\Delta \to 0$, this last equation implies that $\lim_{\Delta\to 0}\left(P(\omega) - P^C\right) = 0$. More precisely,

$$\begin{aligned} P(\omega) - P^C &= O(K(\eta,\Delta)\Delta + \eta + K(\eta,\Delta)\eta\Delta + K(\eta,\Delta)^2\Delta^2) \\ &= O(K(\eta,\Delta)\Delta + \eta). \end{aligned}$$

Taking the average of (13) with respect to the stationary distribution of $\tilde{m}_\Delta$ (which amounts to replacing $P(\omega)$ with P^C), dropping a term of order $K(\eta,\Delta)\Delta^2$, and simplifying leads to

$$\frac{1}{K(\eta,\Delta)}\left(K(\eta,\Delta)\Delta + K(\eta,\Delta)\varepsilon^C\Delta\right)P^C = \left(d^C + [f\varepsilon]^C\right)\Delta + O(\eta/K(\eta,\Delta) + \eta\Delta + K(\eta,\Delta)\Delta^2),$$

and hence

$$P^C = \frac{d^C + [f\varepsilon]^C}{1+\varepsilon^C} + O\left(\frac{\eta}{\Delta K(\eta,\Delta)} + \eta + K(\eta,\Delta)\Delta\right).$$

Notice that, as $\Delta \to 0$, the stationary distribution of $\tilde{m}_\Delta$ approaches ϕ_C, where ϕ_C is the stationary distribution of the true process q restricted to C.[10] It suffices to show that there exists $\Delta(\eta)$ such that $\eta/\left(\Delta(\eta)K\left(\eta,\Delta(\eta)\right)\right)$ and $K\left(\eta,\Delta(\eta)\right)\Delta(\eta)$ vanish as $\eta \to 0$. Given $a, b \in (0,1)$ such that $a < b$, it suffices to take $\Delta(\eta)$ such that $\eta^b < \Delta(\eta)K\left(\eta,\Delta(\eta)\right) < \eta^a$. By Lemma 2, the upper bound is satisfied for sufficiently small Δ. If the lower bound is not satisfied for any $\Delta > 0$ then we can simply replace $K(\eta,\Delta)$ with a larger value for a particular Δ in order to satisfy both bounds. □

[9] Note that, although it is omitted from the notation, each of these averages depends on Δ.

[10] To see this, note that for each individual i categorizing ω and ω' together, the limiting transition probability $m^i_\Delta(\omega,\omega')$ from ω to ω' is proportional to the stationary distribution mass assigned to ω'. Hence the stationary distribution of q restricted to C is also stationary with respect to each individual belief $m^i_\Delta(\omega,\omega')$. Aggregating across individuals gives the claim.

B Proof of Proposition 5

Proof of Proposition 5. It is straightforward to verify that each agent's expectation of the present value of the dividend stream within any trading period is

$$\int_t^{t+\Delta_n} e^{-(t'-t)} E[d(\omega_{t'}) \mid \omega_t] dt' = \int_0^{\Delta_n} e^{-t'} e^{-t'} \omega_t dt' = \frac{1 - e^{-2\Delta_n}}{2} \omega_t,$$

and that, if $P(\omega) = r\omega$ for all ω,

$$\begin{aligned} E_n \left[P\left(\omega_{t+\Delta_n}\right) \mid \omega_t = \omega \right] &= \frac{\int_{-\infty}^{\infty} \phi(\tilde{\omega}) \sigma_n(\omega, \tilde{\omega}) r E[\omega_{t+\Delta_n} \mid \omega_t = \tilde{\omega}] d\tilde{\omega}}{\int_{-\infty}^{\infty} \phi(\tilde{\omega}) \sigma_n(\omega, \tilde{\omega}) d\tilde{\omega}} \\ &= \frac{\int_{-\infty}^{\infty} \phi(\tilde{\omega}) \sigma_n(\omega, \tilde{\omega}) r e^{-\Delta_n} \tilde{\omega} d\tilde{\omega}}{\int_{-\infty}^{\infty} \phi(\tilde{\omega}) \sigma_n(\omega, \tilde{\omega}) \tilde{\omega} d\tilde{\omega}} \\ &= \frac{e^{-\Delta_n} r \omega_t}{1 + \rho_n^2}. \end{aligned}$$

Since $P(\omega) = \sum_n \pi_n P_n(\omega)$, it follows that

$$r\omega = \sum_n \pi_n \left(\frac{1 - e^{-2\Delta_n}}{2} + r \frac{e^{-\Delta_n} \omega_t}{1 + \rho_n^2} \right) \omega.$$

Solving for r gives (10).

It is immediate that r is decreasing in each ρ_n. To see that it is increasing in Δ_n, note that the sign of $\frac{\partial r}{\partial \Delta_n}$ is identical to that of

$$2\pi_n e^{-2\Delta_n} \left(1 - \sum_{n'} \pi_{n'} \frac{e^{-2\Delta_{n'}}}{1 + \rho_{n'}^2} \right) - \left(1 - \sum_{n'} \pi_{n'} e^{-2\Delta_{n'}} \right) \frac{2\pi_n e^{-2\Delta_n}}{1 + \rho_n^2},$$

which in turn has the same sign as

$$\sum_{n'} \left(1 - \frac{e^{-2\Delta_{n'}}}{1 + \rho_{n'}^2} - \frac{1}{1 + \rho_n^2} + \frac{e^{-2\Delta_{n'}}}{1 + \rho_n^2} \right).$$

This last expression is greater than

$$\sum_{n'} \left(1 - \frac{e^{-2\Delta_{n'}}}{1 + \rho_{n'}^2} - \frac{1}{1 + \rho_n^2} + \frac{e^{-2\Delta_{n'}}}{\left(1 + \rho_n^2\right)\left(1 + \rho_{n'}^2\right)} \right) = \sum_{n'} \left(1 - \frac{e^{-2\Delta_{n'}}}{1 + \rho_{n'}^2} \right) \left(1 - \frac{1}{1 + \rho_n^2} \right),$$

which is non-negative.

The characterization of the limit when all Δ_n vanish is immediate. $\square$

References

Al-Najjar, N. and M. Pai (2014). Coarse decision making and overfitting. *Journal of Economic Theory 150*, 467–486.

Allen, F., S. Morris, and H. Shin (2006). Beauty contests and iterated expectations in asset markets. *Review of Financial Studies 19*(3), 719–752.

Bacchetta, P. and E. Van Wincoop (2008). Higher order expectations in asset pricing. *Journal of Money, Credit and Banking 40*(5), 837–866.

Barberis, N. and A. Shleifer (2003). Style investing. *Journal of Financial Economics 68*(2), 161–199.

Bianchi, M. and P. Jehiel (2010). Bubbles and crashes with partially sophisticated investors. *PSE Working Papers*.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann (1990a). Noise trader risk in financial markets. *Journal of Political Economy 98*(4), 703–738.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann (1990b). Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance 45*(2), 379–395.

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

Fehr, E. and J. Tyran (2005). Individual irrationality and aggregate outcomes. *Journal of Economic Perspectives 19*(4), 43–66.

Gilboa, I. and D. Schmeidler (1995). Case-based decision theory. *Quarterly Journal of Economics 110*(3), 605–639.

Haltiwanger, J. and M. Waldman (1985). Rational expectations and the limits of rationality: An analysis of heterogeneity. *American Economic Review 75*(3), 326–340.

Haltiwanger, J. and M. Waldman (1989). Limited rationality and strategic complements: the implications for macroeconomics. *Quarterly Journal of Economics 104*(3), 463–483.

Harrison, J. and D. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *Quarterly Journal of Economics 92*(2), 323–336.

Hartfiel, D. (1998). *Markov Set-Chains*. Springer.

Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory 123*(2), 81–104.

Jehiel, P. and D. Samet (2007). Valuation equilibrium. *Theoretical Economics 2*(2), 163–185.

Morris, S. and H. S. Shin (2002). Social value of public information. *American Economic Review 92*(5), 1521–1534.

Mullainathan, S., J. Schwartzstein, and A. Shleifer (2008). Coarse thinking and persuasion. *Quarterly Journal of Economics 123*(2), 577–619.

Samet, D. (1998). Iterated expectations and common priors. *Games and Economic Behavior 24*(1-2), 131–141.

Steiner, J. and C. Stewart (2008). Contagion through learning. *Theoretical Economics 3*(4), 431–458.