Price Distortions in High-Frequency Markets*

Jakub Steiner[†]
Northwestern University and CERGE-EI

Colin Stewart[‡]
University of Toronto

March 15, 2013

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. Similar results hold if, instead of using a simplified model of the world, some agents overestimate the likelihood of small probability events, as in prospect theory.

*We thank Milo Bianchi, Philippe Jehiel, Willemien Kets, Filip Matejka, John Moore, Stephen Morris, Andreas Park, Rani Spiegler, Muhamet Yildiz, and seminar participants for helpful comments.
[†]email: j-steiner@kellogg.northwestern.edu
[‡]email: colinbstewart@gmail.com

1 Introduction

Recent advances in technology have led to dramatic increases in trading speed. These changes have generated considerable debate about the effect of high frequencies on market prices, with one side arguing that, by providing liquidity, high frequency traders help to improve market efficiency, and the other side claiming that high frequencies can destabilize market prices. We present a model that lends support to the latter view and show that agents' use of simplifying models of the world can generate large distortions in prices at high trading frequencies.

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 differing states. These agents group states into categories and form forecasts in each state that are correct on average for the category containing that state.

If enough agents use very fine categorizations of states, then prices are typically close to rational expectations prices when the period length is not too small. However, as trading opportunities become frequent, distortions become large and prices collapse across states. More precisely, whenever two states are categorized together by a positive mass of agents, the price in those two states becomes identical in the high-frequency limit. 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 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.

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 trading frequency becomes high. 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 at high trading frequency. 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. At high frequency, 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 coordination of individual expectations is most transparent in our baseline model in Section 4. There, we focus on a special case in which the market price is equal to the average of all agents' expected values of owning the asset. In that case, we can explicitly express the steady-state prices as a sum of higher-order expectations of future dividends. At high frequencies, prices are driven by very high orders of expectations, which depend only on those features of individuals' theories that are common to all.[1]

Our main result holds under general conditions on the relationship between market prices and individual forecasts. In Section 5, 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.

Throughout much of the paper, we take a reduced form approach to price formation that eschews modelling demand and supply explicitly, and rather expresses current market price directly in terms of forecasts of future prices. The reduced form approach is sufficient to highlight the mechanism underlying the distortionary effect, and the main reason for taking it is tractability. In our baseline model, the pricing equation we apply directly is similar to that in a standard CARA-Normal overlapping generations framework. However, a key assumption of the standard framework is that agents live for only two periods, which is seemingly incompatible with our focus on vanishingly small period length. Accordingly, we examine robustness of our results in a model with explicit supply and demand from risk averse agents whose expected lifespan—in terms of real time—is fixed as the period length varies. Although such a model appears to be intractable in general, it is tractable in some special cases. In Section 6.1, we study a simple class of examples in which fundamentals depend on two variables but individual agents base expectations on only one of those variables. The results are in accordance with the reduced form analysis: as the period length vanishes, prices collapse across states, including pairs of states that no individual agent categorizes together. Although we cannot say whether the result extends to more general settings, this class of examples demonstrates that the effect we highlight persists even if agents face substantial risk over their lifetimes.

The coarse pricing result in this class of examples can be understood in terms of market demand. Learning using categories leads to outcomes as if agents form rational expectations based

[1] The setting of Section 4 is closely related to a dynamic version of a Morris and Shin (2002) beauty contest. 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).

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.

Results similar to our main result hold if, instead of categorizing states together, some agents overweight the likelihood of small probability events, as in prospect theory. Standard weighting functions used in the prospect theory literature (see, e.g., Prelec 1998, Gonzalez and Wu 1999) have the property that, relative to the true probabilities, weighted probabilities grow large as the probability vanishes. By increasing the weight placed on other states when forming price forecasts, this property has essentially the same effect as the use of categorization. Thus we obtain a similar result exhibiting coarse prices.

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 high frequencies 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 at high frequencies when the heterogeneity in beliefs about transition rates between those states is relatively high. Section 7 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 2009). 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 (2012), 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 our paper. They focus on how the composition of theories in the market affects prices and agents' individual performance at 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 shortening horizons 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 high frequencies 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)

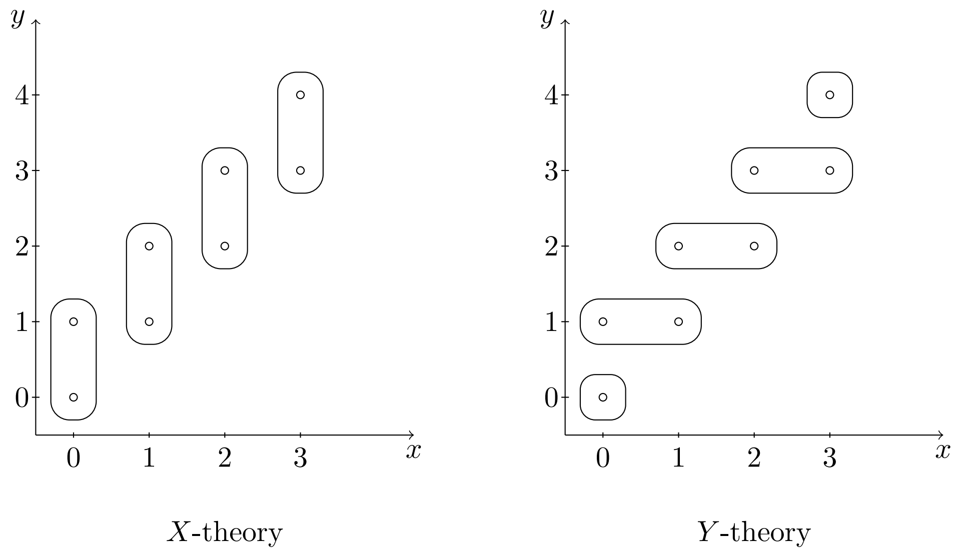

Figure 1: State space and categorizations for Section 3 when $S = 3$.

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 Example

We begin with a simple example to illustrate how categorization can lead to large distortions in prices when trading is frequent.

A continuum of investors $i \in [0,1]$ trade a single risky asset with fixed positive supply at times $t = 0, \Delta, 2\Delta \ldots$. The asset pays a flow dividend of $d(\omega_k)$ per unit of time that depends on the current state $\omega_k \in \Omega$ in period $k = t/\Delta$, where

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

as depicted in Figure 1. The state follows a continuous-time Markov process with transition rates $q(\omega,\omega') = 1/(2S+2)$ from state ω to state $\omega' \neq \omega$. This process can be thought of as drawing a new state at times that are distributed according to a Poisson process with arrival rate one, with the new state drawn uniformly from the entire state space. Let $q_\Delta(\omega, \omega')$ denote the single-period transition probability, that is, the probability that ω' is the state at time $t + \Delta$ given that ω is the state at time t. The flow dividend in state $\omega = (x, y)$ is

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

Thus flow dividends range from 0 to 1. The state is publicly observed prior to trade in each period.

For the purpose of this example, the price of the asset is determined as in Harrison and Kreps (1978). All agents are risk neutral and short selling is not possible. In each period, each agent chooses her demand to maximize her expected gains from buying the asset and reselling it in the following period. Risk neutrality ensures that the optimal demand is independent of the agent's lifespan; the demand of an agent with a long horizon is identical to that of an agent who maximizes single-period gains.

Agents are divided into a finite number of groups that differ in their forecasts of future prices; all agents within a group are identical. Agent i's willingness to pay for the asset is

$$\int_{k\Delta}^{(k+1)\Delta} d(\omega_k) e^{-(t-k\Delta)} dt + e^{-\Delta} E_k^i \left[P_{k+1}\right] = (1 - e^{-\Delta}) d(\omega_k) + e^{-\Delta} E_k^i \left[P_{k+1}\right],$$

where $E_k^i \left[P_{k+1}\right]$ denotes agent i's forecast in period k of the price in period $k + 1$, and agents have a common discount factor normalized to $1/e$. For markets to clear in period k, the price P_k must be equal to the maximal willingness to pay across agents, that is,

$$P_k = \max_i \left\{(1 - e^{-\Delta}) d(\omega_k) + e^{-\Delta} E_k^i \left[P_{k+1}\right]\right\}.$$

At any higher price, aggregate demand for the asset would be 0, and at any lower price aggregate demand would be infinite. We focus on steady-state prices $P_\Delta(\omega)$ that, for each Δ, depend only on the current state ω.

Steady-state prices are determined by agents' forecasts of future prices. As a benchmark, suppose that all agents' forecasts are based on the true process q_Δ, so that steady-state prices satisfy

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

The solution to this system of linear equations is given by

$$P_\Delta(\omega) = \frac{2d(\omega) + e^{-\Delta}}{2(1 + e^{-\Delta})}.$$

In this case, as usual, the price in each state is equal to the sum of the expected discounted future dividends.[2]

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 value of x. This coarse theory, which we call the X-theory, is represented by a partition Π_X of Ω into sets $\{(x, x), (x, x + 1)\}$ for $x = 0, \ldots, S$, as depicted in Figure 1. We refer to the elements of Π_X as categories. 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

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

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

As in the benchmark case, the prices $P^X_\Delta(\omega)$ given these coarse forecasts satisfy (1) 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, in each state ω,

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

When all agents use the same theory, the magnitude of price distortions 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

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

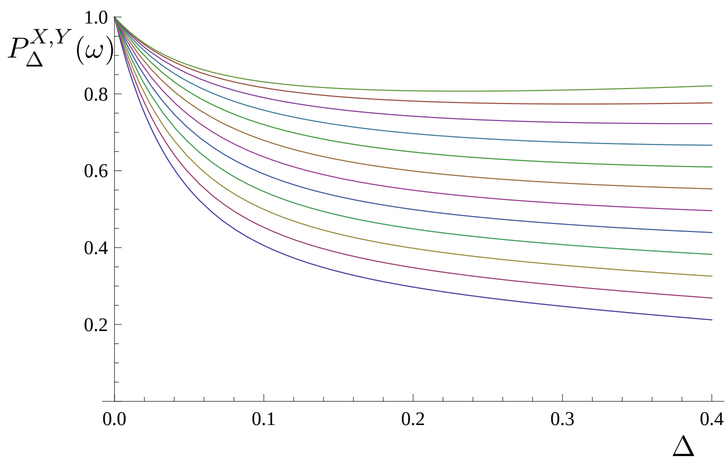

Figure 2: Prices as a function of the period length for $S = 5$ when both the X-theory and the Y-theory are present in the market. Each curve depicts the price in a particular state.

partition Π_Y of Ω into categories $\{(x, y') \in \Omega \mid y' = y\}$ for $y = 0, \dots, S+1$, as depicted in Figure 1.[3] For $Z \in \{X, Y\}$, agents using the Z-theory form expectations as if they believe the process has transition probabilities $q_\Delta^Z(\omega, \omega')$ equal to the average of $q_\Delta(\tilde{\omega}, \omega')$ over all $\tilde{\omega}$ lying in the same element of Π_Z as ω.

One can show that, for any Δ, the steady-state prices $P_\Delta^{X,Y}(\omega)$ 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) 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 $P_\Delta^{X,Y}(\omega)$ satisfy (1) except with the true process q_Δ replaced by q_Δ^X if ω lies on the diagonal, and replaced by q_Δ^Y if ω is off the diagonal.

Figure 2 depicts, for each state ω, the price $P_\Delta^{X,Y}(\omega)$ as a function of Δ when $S = 5$. 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 do high-frequency 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

[3] 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.

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 $\omega = (x, x)$ and let $\omega' = (x, x+1)$. 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 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 entire state space).

High frequency pricing is coarse in much more general settings, including settings in which market prices are not determined according to the Harrison and Kreps (1978) framework. We prove that under a mild condition relating market prices to individual reservation prices, agents perfectly coordinate their reservation prices in the limit.

To summarize, as the period length vanishes, agents place increasing weight on the resale price and diminishing weight on the dividend. As a result, in the limit, reservation prices end up being perfectly coordinated across agents. Since agents' actions can only be coordinated if they condition only on information used by all agents, their actions must be measurable with respect to the finest common coarsening of all theories.

4 Baseline model

4.1 Model setup

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 ω'. 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 $q_\Delta(\omega, \omega')$. The state affects the flow dividend $d(\omega_k)$

of the asset, which is paid at a constant rate from time $k\Delta$ to $(k+1)\Delta$.[4]

A continuum of agents of measure one trades the asset in each trading period k. We focus on steady-state prices $P : \Omega \longrightarrow \mathbb{R}$ that depend only on the current state. The market price $P(\omega)$ 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 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 $P(\omega)$ and any category $C \in \Pi_{n(i)}$, the forecasts E^i satisfy

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

where ϕ denotes the stationary distribution of states with respect to the true process q.[5] It follows that E^i is identical to the expectation with respect to the *modified process* $m_\Delta^{n(i)}$ given by

$$m_\Delta^n(\omega, \omega') = \sum_{\omega'' \in \Pi_n(\omega)} \phi\left(\omega'' \left| \Pi_n(\omega) \right.\right) q_\Delta(\omega'', \omega'). \tag{2}$$

In this section, we focus on a simple reduced-form setting in which market prices directly aggregate individual reservation prices. Section 5 extends the analysis to more general pricing rules. Focusing on this simple model offers two advantages: first, because of its relative simplicity, we believe the main insights are more transparent in this case, and second, we can fully characterize high-frequency prices in a way that clearly illustrates exactly how prices reflect a coarse model of the market.

Market prices are determined as follows. Given a price vector P, each agent in group n forms

[4]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.

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

demand $\alpha_n(\omega)$ proportional to her net expected profit from holding one unit of the asset for one period; that is, in each state ω, her demand is

$$\alpha_n(\omega) = (1 - e^{-\Delta})d(\omega) + e^{-\Delta} E_n^\omega \left[P(\omega') \right] - P(\omega),$$

where $E_n^\omega [P(\omega')]$ denotes the expected resale price $E_{m_\Delta^n(\omega,\cdot)} [P(\cdot)]$ in the next period given that the current state is ω. Assuming zero net supply, the market-clearing price is

$$P(\omega) = (1 - e^{-\Delta})d(\omega) + e^{-\Delta} \sum_n \pi_n E_n^\omega \left[P(\omega') \right]. \tag{3}$$

In the supplementary appendix, we show that the steady-state prices defined by (18) and (3) 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 extends the framework in two directions. First, we allow 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. Second, we extend the price forecasting rule to a general class of similarity-based rules in which the weights assigned to different states may vary according to the perceived degree of similarity. Our main result extends in both cases.

4.2 High frequency

Focusing on high frequencies (i.e. vanishing Δ) leads to a striking result: prices generally fail to distinguish among states that may differ substantially in terms of fundamentals.

Let Π denote the meet of $\Pi_1, \ldots, \Pi_N$.[6] 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 such that $\omega = \omega_1$, $\omega' = \omega_r$, and for each $\ell = 1, \ldots, r-1$, $\omega_{\ell+1} \in \Pi_n(\omega_\ell)$ for some $n \in \{1, \ldots, N\}$. 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.

The main result of this paper shows that, in the limit as Δ vanishes, prices are constant on aggregate categories. Moreover, prices 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 *rational expectations prices* with respect to a continuous-time Markov process $\tilde{q}$ and a dividend function $\tilde{d}$ to be the unique solution P to the system of equations

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

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 category with weights determined by the stationary distribution ϕ.

Proposition 1. *As Δ vanishes, prices in the baseline model become constant on each aggregate*

[6] The meet of a collection of partitions is their finest common coarsening.

category; that is, for any ω, ω' such that $\Pi(\omega) = \Pi(\omega')$,

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

Moreover, for each ω, $\lim_{\Delta \to 0} P_\Delta(\omega)$ 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.

For high trading frequencies, this result indicates that 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 these 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.

4.3 Discussion

Our baseline model can be interpreted as a dynamic version of a Morris and Shin (2002) beauty contest in which each agent i chooses an action $P^i \in \mathbb{R}$ in each period. Agent i's flow payoff from choosing P^i in period k is

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

where P_{k+1} is the average action in period $k + 1$. Thus each agent wants to match her action to both the current dividend and the average action in the next period, with increasing weight on the average action term as the period length vanishes. Steady-state prices in our baseline model correspond to average actions in a stationary equilibrium of this game.

The connection to Morris and Shin (2002) helps for understanding the role of higher order expectations, which are central to the proof of Proposition 1.

Given prices P, define the *reservation price* of agent i in state ω to be

$$P^i(\omega) = (1 - e^{-\Delta})d(\omega) + e^{-\Delta} E^{\omega}_{n(i)} \left[P(\omega') \right]. \tag{4}$$

Thus the reservation price is the expected value from holding one unit of the asset for one period given the current state. In our baseline model, this is the price at which agent i has zero demand.

Equation (3) can be rewritten as

$$P(\omega) = \int_0^1 P^i(\omega) di;$$

that is, the market price is the population average of the reservation prices. Substituting for $P^i(\omega)$ from (4) and iterating yields an expression for the reservation price as a sum of higher order expectations of dividends, namely,

$$P^i(\omega) = \left(1 - e^{-\Delta}\right) d(\omega) + \left(1 - e^{-\Delta}\right) \sum_{k=0}^{\infty} e^{-(k+1)\Delta} E^i \left(\overline{E}^{\omega}\right)^k [d(\cdot)], \tag{5}$$

where $\overline{E}^{\omega} = \sum_n \pi_n E^{\omega}_n$ 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 (5) 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 in Morris and Shin (2002) the high order expectations 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 trade is frequent, it follows that reservation prices converge across groups.

The proof of Proposition 1 explicitly characterizes high order expectations in the baseline model. In more general settings, such a characterization is not feasible; however, convergence of reservation prices holds under much weaker conditions. When combined with the observation that, as Δ vanishes, a group's reservation prices become constant within each of its own categories, this result implies that reservation prices are constant on aggregate categories.

5 General model

The constant price result of the preceding section extends to a much more general class of pricing equations. All elements of the model are the same except for the way in which individual forecasts are aggregated to determine market prices.

Assume that there exists $\mu \in (0,1]$ such that for all ω and Δ, the steady-state price $P_\Delta(\omega)$ satisfies

$$(1-\mu)\min_n P^n_\Delta(\omega) + \mu \max_n P^n_\Delta(\omega) \le P_\Delta(\omega) \le \max_n P^n_\Delta(\omega), \tag{6}$$

where $P^n_\Delta(\omega)$ denotes the reservation price of agents from group n defined by (4). Roughly speaking, this assumption states that the price is never so high as to make every agent receive an expected loss from buying the asset, and 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). This assumption is satisfied with $\mu = \min_n \pi_n$ by the pricing rule in the previous section, and with $\mu = 1$ by the Harrison and Kreps (1978) pricing rule used in the example in Section 3. In particular, like Section 3, the current section allows for investors with long horizons.

Proposition 2. *As Δ vanishes, prices become constant on each aggregate category; that is,*

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

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

Lemma 1. *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 establishes that the agent would increase her reservation price in that state, establishing a contradiction.

Proof of Lemma 1. Suppose for contradiction that the lemma 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

$$(\omega^*, n^*) \in \underset{(\omega,n)\in R}{\arg\min} P^n(\omega). \tag{7}$$

By (6), 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

$$P^n(\omega) \geq P^{n^*}(\omega^*) \tag{8}$$

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 (7).

Inequalities (6) and (8) 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. □

Proof of Proposition 2. From (4), 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. □

6 Extensions

6.1 Risk averse agents with long horizons

For the sake of tractability, we have, for much of the paper, taken a reduced form approach to price formation. Solving for market-clearing prices in a model with long-lived risk averse agents appears to be intractable in general; we can, however, solve some special cases. Our basic insight extends to these cases: prices collapse across any two states that a positive mass of agents categorize together.

In this section, we return to the setting of Section 3, except that instead of Harrison and Kreps (1978) pricing, prices are determined by market clearing given demands of risk averse agents with short-selling allowed. The state space, dividends, and Markov process are the same as in Section 3. There are two groups X and Y of equal size; group X categorizes states according to the X-theory, group Y according to the Y-theory (see Figure 1). Each agent enters the market at some time and remains in the market for an exponentially distributed length of time with a common probability parameter (in particular, the expected lifespan is independent of Δ, and as Δ vanishes, each agent trades in the market for a number of periods that grows on the order of $1/\Delta$). Agents learn their time of exit only upon exiting. The rate of entry and exit in the market is the same, so that the size of the population is constant.

All agents have constant absolute risk aversion with Bernoulli utility $-e^{-w}$, where w is their wealth at the time of exit from the market. Agents enter the market with zero wealth and can freely borrow or save at a constant (continuous-time) interest rate normalized to 1.

In this setting, each agent's demand depends on her belief about her future change in wealth, which in turn may depend on her belief about the entire future path of states and her level of sophistication regarding her own future trading behavior. We do not specify a particular form for these beliefs; instead, we allow them to depend on prices and period length in a general way. For

$Z \in \{X, Y\}$, given steady-state prices P_Δ, we introduce a random variable W_Δ^Z. The distribution of W_Δ^Z is the subjective belief of each agent using theory Z in any period k about the present value of her future change in wealth from period $k+1$ until she exits the market.

An agent using theory Z chooses her demand $\alpha_\Delta^Z(\omega)$ to maximize

$$E_Z^\omega\left[-\exp\left(-W_\Delta^Z - \alpha_\Delta^Z(\omega)R_\Delta(\omega,\omega')\right)\right] = -\sum_{\omega'} m_\Delta^Z(\omega,\omega')V_\Delta^Z(\omega')\exp\left(-\alpha_\Delta^Z(\omega)R_\Delta(\omega,\omega')\right),$$

where

$$R_\Delta(\omega,\omega') = -P_\Delta(\omega) + (1-e^{-\Delta})d(\omega) + e^{-\Delta}P_\Delta(\omega')$$

is the gain from holding one unit of the asset for one period, and

$$V_\Delta^Z(\omega') = E[\exp(-W_\Delta^Z) \mid \omega'],$$

with ω' denoting the state in the following period.

Prices P_Δ are *equilibrium prices* if there are optimal demands α_Δ^Z given P_Δ that satisfy the market-clearing conditions

$$\alpha_\Delta^X(\omega) + \alpha_\Delta^Y(\omega) = 0 \tag{9}$$

for each ω.

We assume that, for each Z, V_Δ^Z depends continuously on P_Δ and, given any equilibrium prices $P_\Delta \in [0,1]^\Omega$, V_Δ^Z is bounded and bounded away from 0 as Δ vanishes. This assumption effectively means that agents do not expect arbitrarily large future losses, or, with probability approaching one, arbitrarily large future gains at high frequency.

Proposition 3. *As Δ vanishes, equilibrium prices become constant across states; that is, for any equilibrium prices P_Δ,*

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

for all $\omega, \omega' \in \Omega$.

Intuitively, one can think of the effect as follows. Suppose for contradiction that prices differ within a category used by some agents. Consider the state within that category at which the price is minimal (and suppose that there are no other agents categorizing that state together with one where the price is even lower). Then agents using this category act as if they believe the state is likely to transition to another state within the category, meaning they act as if the price is likely to go up very soon. Consequently, those agents will demand a large quantity at the current price. This demand drives up the price, causing agents who do not expect a price increase to view the asset as overpriced, and hence to sell. Note, however, that as the trading frequency shrinks, so do the per-period dividend and the true transition probabilities, causing the incentive to sell the asset to decrease in proportion to the period length. This means that, when trading is frequent, agents using the given categorization expect a much larger gain from buying than do other agents from selling, causing them to demand more and drive the price up further. As the period length vanishes, these two pressures on the market price will be balanced only if the differences in prices within the category also vanish.

6.2 Weighted probabilities

The preceding section indicates that coarsening of steady-state prices arises because agents who categorize two given states together effectively overestimate the probabilities of transitions between them. In this section we analyze a setting in which a similar result arises not from coarse thinking but from the use of weighted probabilities, as in prospect theory. In particular, we assume here that some agents overweight the likelihood of small probability events.

The sets of agents, the discount rate, and the Markov process are the same as in the previous sections. Reservation and market prices are determined as in Section 5. The key difference is in agents' forecasts of future prices. Agents from group n form forecasts using weighted transition probabilities

$$m_\Delta^n(\omega,\omega') = \frac{\lambda_n\left(q_\Delta(\omega,\omega')\right)}{\sum_{\tilde{\omega}} \lambda_n\left(q_\Delta(\omega,\tilde{\omega})\right)},$$

where the weighting function $\lambda_n : [0,1] \longrightarrow \mathbb{R}_+$ is increasing and bounded. In particular, λ_n may be the identity function, in which case agents in group n form correct forecasts.

We make the following assumptions:

A1: For some group n, $\lim_{p\to 0^+} \frac{\lambda_n(p)}{p} = \infty$.

A2: The underlying process is fully connected: $q(\omega,\omega') > 0$ for all distinct states ω and ω'.

Assumption A1 holds for most weighting functions commonly used in prospect theory (e.g., Prelec 1998, Gonzalez and Wu 1999). For small Δ, A1 allows for smaller distortions of transition probabilities than those introduced by the use of categorization. While categorization makes the weighted transition probabilities $m^n_\Delta(\omega, \omega')$ bounded away from 0 for any pair of states in the same category, A1 allows for transition probabilities that converge to 0 as $\Delta \to 0$.

Assumption A2 is made for simplicity. Without this assumption, the same result holds under more restrictive conditions on λ. For example, Proposition 4 holds without A2 for weighting functions of the form $\lambda(p) = \exp\left(-\zeta(-\ln p)^\xi\right)$ with $\zeta > 0$ and $\xi \in (0, 1)$, as axiomatized by Prelec (1998).

Proposition 4. *Assume* (6)*, A1, and A2. Then as Δ vanishes, prices become constant across states; that is,*

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

for all $\omega, \omega' \in \Omega$.

The outline of the proof is simple. Suppose that prices do not collapse across states. To establish a contradiction, consider a state ω^* with the minimal price. Since some agents overweight small probabilities, the probability they assign to a discrete price increase has order exceeding Δ, while the dividend is of order Δ. Therefore, their reservation price exceeds the market price at ω^* by order exceeding Δ, driving the market price up above its actual value.

7 Concluding remarks

In order to highlight the effect of high frequencies, 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 high frequencies. 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 our results hold as long as there is continual entry of a nonvanishing mass of new traders using coarse categories.

In the variant of our model that incorporates individual demands in Section 6.1, risk aversion limits the size of the position taken by each trader. 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 at a fixed trading frequency we expect that increasing risk aversion should reduce price distortions. However, as Section 6.1 indicates, risk aversion alone does not overturn the constant price result in the high frequency limit.

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, at high frequency, 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.

Another simplifying assumption of our model is that all agents are able to trade at the same frequency. If instead some agents are forced to maintain their positions for some fixed time that is independent of the period length then these agents may drive prices back toward fundamentals. On the one hand, in the baseline model of Section 4, the addition of such traders would reduce the impact of coarse categorization, leading to prices lying between rational expectations and constant prices on categories. On the other hand, explicitly considering individual demands may mitigate the influence of these agents since agents who can trade at higher frequency tend to take larger positions.

Appendices

A Proof of Proposition 1

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)

$$\|\tilde{m}^k_\Delta(\omega, \cdot) - \phi^C_\Delta\| \leq 2(1-\varepsilon_\Delta)^{k-1} \tag{10}$$

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 (10) 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 $1 - (1-\delta)^{\frac{1}{L}}$.

□

Proof of Proposition 1. Notice that equations (3) and (18) 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

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

where $\varepsilon_\omega = \lim_{\Delta\to 0} \Pr_{m_\Delta(\omega,\omega')}[\omega' \notin C]/\Delta$.[7]

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 $\omega' \in C$.

Iterating the last equation for K periods gives

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

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

$$\begin{aligned} \frac{1}{K}P(\omega) = \frac{1}{K}\sum_{k=0}^{K-1} E\left[(d(\omega_k) + f(\omega_k,\Delta)\varepsilon_{\omega_k})\Delta\right] \\ + \frac{1}{K}(1-K\Delta)E\left[\left(1-\sum_{k'=0}^{K-1}\varepsilon_{\omega_{k'}}\Delta\right)P(\omega_K)\right] + O(K\Delta^2). \end{aligned} \quad (11)$$

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

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

[7] 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.

Substituting into (11) and applying Lemma 2 gives

$$\frac{1}{K(\eta,\Delta)}P(\omega) = \left(d^C + [f\varepsilon]^C\right)\Delta + (1 - K(\eta,\Delta)\Delta)\left(1 - K(\eta,\Delta)\varepsilon^C\Delta\right)\frac{1}{K(\eta,\Delta)}P^C$$
$$+ O(\eta/K(\eta,\Delta) + \eta\Delta + K(\eta,\Delta)\Delta^2), \quad (12)$$

where x^C denotes the average of $x(\omega)$ with respect to the stationary distribution of the process $\tilde{m}_\Delta$.[8] 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 (12) 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.[9] 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 $a, b \in (0,1)$

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

[9] 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.

such that $a < b$, it suffices to take $\Delta(\eta)$ such that $\eta^b < \Delta(\eta) K(\eta, \Delta(\eta)) < \eta^a$. 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. □

B Proof of Proposition 3

The first-order condition for this problem for an agent using theory Z is

$$\sum_{\omega'} m_\Delta^Z(\omega, \omega') V_\Delta^Z(\omega') R_\Delta(\omega, \omega') \exp\left(-\alpha_\Delta^Z(\omega) R_\Delta(\omega, \omega')\right) = 0. \tag{13}$$

Lemma 3. *For every Δ and ω, and all steady-state prices $P_\Delta(\omega)$, we have $P_\Delta(\omega) \in [0, 1]$.*

Proof. The first-order condition (13) can only hold if, for each Δ, $R_\Delta(\omega, \omega')$ is nonnegative for some ω' and nonpositive for some other ω'. Hence in the state $\overline{\omega}$ with the highest price, we must have $R_\Delta(\overline{\omega}, \overline{\omega}) \geq 0$, which implies that $P_\Delta(\overline{\omega}) \leq d(\overline{\omega})$. Similarly, we must have $P(\underline{\omega}) \geq d(\underline{\omega})$ in the lowest price state $\underline{\omega}$. The result follows since $d(\omega) \in [0, 1]$ for every ω. □

In case there are multiple steady-state price vectors for some Δ, fix an arbitrary choice.

We want to show that $\lim_{\Delta \to 0} (P_\Delta(\omega) - P_\Delta(\omega')) = 0$ whenever $\omega' \in \Pi(\omega)$. Note that, because the state space is finite, it suffices to consider vanishing sequences $(\Delta_\ell)_{\ell=1}^\infty$ for which there is an ordering $\omega^1, \ldots, \omega^S$ of all states such that $P_{\Delta_\ell}(\omega^1) \leq P_{\Delta_\ell}(\omega^2) \leq \cdots \leq P_{\Delta_\ell}(\omega^S)$ for each Δ_ℓ. From this point on, we fix such a subsequence and restrict attention to Δ such that $\Delta = \Delta_\ell$ for some ℓ.

Lemma 4. *Given any state ω, suppose there exists $\varepsilon \in (0, 1]$ such that for every $\omega' \in \Pi(\omega)$ for which $R_\Delta(\omega, \omega') < 0$ for sufficiently small Δ, we have $R_\Delta(\omega, \omega') = O(\Delta^\varepsilon)$. Then $R_\Delta(\omega, \omega'') = O(\Delta^{\varepsilon/2})$ whenever $\omega'' \in \Pi_Z(\omega)$ for some Z.*

Proof. Let ω, $\overline{\omega}$, and Z be such that the premise holds for ω and $\overline{\omega} \in \Pi_Z(\omega)$. We need to show that $R_\Delta(\omega, \overline{\omega}) = O(\Delta^{\varepsilon/2})$. If $R_\Delta(\omega, \overline{\omega}) < 0$ for sufficiently small Δ then we are done. Accordingly, suppose $R_\Delta(\omega, \overline{\omega}) \geq 0$.

The proof proceeds as follows. First we find a lower bound on the left-hand side of the first-order condition for Z in state ω. The first-order condition is a sum of terms, one for each ω', with each term having the same sign as $R_\Delta(\omega,\omega')$. We find lower bounds on the sum of the negative terms and on the term for $\omega' = \varpi$. For the first-order condition to hold, the sum of these two lower bounds cannot be greater than 0 (since all other terms are nonnegative). This gives an inequality that can be solved for the demand $\alpha^Z_\Delta(\omega)$ to obtain a lower bound on $\alpha^Z_\Delta(\omega)$ that depends on $R_\Delta(\omega,\varpi)$. We then compute a similar lower bound on the demand of agents using theory $\tilde{Z} \neq Z$. Market clearing requires that the sum of the two lower bounds on demand be at most 0. This gives an inequality that yields the desired bound on $R_\Delta(\omega,\varpi)$.

Consider the first-order condition for Z in state ω. For each negative term of the sum in (13) (i.e. each term for which $R_\Delta(\omega,\omega') < 0$), either $\omega' \notin \Pi_Z(\omega)$, in which case $m^Z_\Delta(\omega,\omega') = O(\Delta)$ and $R_\Delta(\omega,\omega') = O(1)$, or $\omega' \in \Pi_Z(\omega)$, which, since $\Pi_Z(\omega) = \{\omega,\varpi\}$, implies that $\omega' = \omega$, in which case $R_\Delta(\omega,\omega') = O(\Delta)$. Therefore, there exists some $M > 0$ such that the sum of all negative terms is bounded below by

$$-M\Delta \exp(M\alpha^Z_\Delta(\omega)).$$

Now consider the term for $\omega' = \varpi$. Since $\varpi \in \Pi_Z(\omega)$, $m^Z_\Delta(\omega,\varpi)$ approaches a positive number as $\Delta \to 0$. Since $V^Z_\Delta(\omega)$ is bounded away from 0, there exists some $L > 0$ such that this term is bounded below by

$$LR_\Delta(\omega,\varpi)\exp\left(-\alpha^Z_\Delta(\omega)R_\Delta(\omega,\varpi)\right).$$

In particular, for (13) to hold, we must have

$$LR_\Delta(\omega,\varpi)\exp\left(-\alpha^Z_\Delta(\omega)R_\Delta(\omega,\varpi)\right) - M\Delta\exp(M\alpha^Z_\Delta(\omega)) \leq 0.$$

Isolating $\alpha^Z_\Delta(\omega)$ yields

$$\alpha^Z_\Delta(\omega) \geq \frac{\log\left(\frac{LR_\Delta(\omega,\varpi)}{M\Delta}\right)}{R_\Delta(\omega,\varpi) + M}. \tag{14}$$

Now consider the first-order condition in state ω for an individual using theory $\tilde{Z} \neq Z$. For each negative term in the first-order condition, $R_\Delta(\omega,\omega') = O(\Delta^\varepsilon)$ by assumption. Therefore, there exists some $M' > 0$ such that the sum of the negative terms in the first-order condition is bounded below by

$$-M'\Delta^\varepsilon \exp(M'\alpha^{\tilde{Z}}_\Delta(\omega)).$$

Now consider the term with $\omega' = \varpi$. Since $m_\Delta^{\tilde{Z}}(\omega, \varpi)$ is bounded below by a constant multiple of Δ, there exists some $L' > 0$ such that this term is bounded below by

$$L' \Delta R_\Delta(\omega, \varpi) \exp(-\alpha_\Delta^{\tilde{Z}}(\omega) R_\Delta(\omega, \varpi)).$$

By the first-order condition, the sum of these two bounds is nonpositive. Isolating $\alpha_\Delta^{\tilde{Z}}(\omega)$ in the corresponding inequality gives

$$\alpha_\Delta^{\tilde{Z}}(\omega) \geq \frac{\log\left(\frac{L'\Delta^{1-\varepsilon} R_\Delta(\omega,\varpi)}{M'}\right)}{R_\Delta(\omega, \varpi) + M'}. \tag{15}$$

Let $M'' = \max\{M, M'\}$ and note that we could replace M or M' with M'' in (14) and (15) and they would still hold (since we can always replace the constant with a larger one at the first step of each derivation).

Substituting (14) and (15) into the market-clearing condition yields

$$\frac{\log\left(\frac{L R_\Delta(\omega,\varpi)}{M''\Delta}\right)}{R_\Delta(\omega, \varpi) + M''} + \frac{\log\left(\frac{L'\Delta^{1-\varepsilon} R_\Delta(\omega,\varpi)}{M''}\right)}{R_\Delta(\omega, \varpi) + M''} \leq 0.$$

This implies that $R_\Delta(\omega, \varpi)^2 \Delta^{-\varepsilon}$ is bounded above (by $(M'')^2/(LL')$). Therefore, $R_\Delta(\omega, \varpi) = O(\Delta^{\varepsilon/2})$. ☐

The proposition follows by applying Lemma 4 inductively starting from the state with the lowest price. In such a state, the premise of the lemma holds with $\varepsilon = 1$. Then the lemma shows that the premise holds for every state that any group categorizes together with that state. Continuing in this fashion yields the conclusion for every state.

C Proof of Proposition 4

Proof. Assume for contradiction that the result does not hold. Then there exists a vanishing sequence Δ_ℓ such that the sequence $P_{\Delta_\ell}(\omega)$ converges for each ω, and there exist at least two states for which the limits of these sequences differ. For each ω, let $P(\omega) = \lim_\ell P_{\Delta_\ell}(\omega)$. Consider

$$\omega_\ell^* \in \arg\min_{\omega\in\Omega} P_{\Delta_\ell}(\omega).$$

Recall that, for sufficiently large ℓ, there exists $\hat{\omega}$ such that $P(\hat{\omega}) > P(\omega_\ell^*)$ and denote $P(\hat{\omega}) - P(\omega_\ell^*)$ by M, and $\min_{\omega\neq\omega'} q(\omega,\omega')$ by α. Note that α is positive by A2. Note also that there exist L and $\beta > 0$ such that

$$P_{\Delta_\ell}^n(\omega_\ell^*) = (1-e^{-\Delta_\ell})d(\omega_\ell^*) + e^{-\Delta_\ell} E_{m^n_{\Delta_\ell}(\omega_\ell^*,\omega)}\left[P^n_{\Delta_\ell}(\omega)\right] \geq P_{\Delta_\ell}(\omega_\ell^*) + L\Delta_\ell + \beta M \lambda_n(\alpha\Delta_\ell),$$

for every group n.

Furthermore, (6) implies that

$$P_{\Delta_\ell}(\omega_\ell^*) \geq (1-\mu)\min_n P^n_{\Delta_\ell}(\omega_\ell^*) + \mu \max_n P^n_{\Delta_\ell}(\omega_\ell^*) \geq P_{\Delta_\ell}(\omega_\ell^*) + L\Delta_\ell + \mu\beta M \lambda_{n^*}(\alpha\Delta_\ell),$$

where n^* is a group for which $\frac{\lambda_{n^*}(p)}{p}$ diverges as $p \to 0$ (which exists by A1). Hence $L\Delta_\ell + \mu\beta M\lambda_{n^*}(\alpha\Delta_\ell)$ is positive for sufficiently large ℓ. This in turn implies that $P_{\Delta_\ell}(\omega_\ell^*) > P_{\Delta_\ell}(\omega_\ell^*)$ when ℓ is large, yielding the desired contradiction. □

D Supplementary material

D.1 Dynamic process

For convenience, we first restate the relevant elements of the model. A single asset pays a dividend that 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 $q(\omega,\omega')$. Trading occurs at discrete times $t = 0, \Delta, 2\Delta, \ldots$. We write ω_k for $\omega(k\Delta)$ and $q_\Delta(\omega,\omega')$ for the transition probabilities between trading periods. A constant per-unit flow dividend of $d(\omega_k)$ is paid from time $k\Delta$ to $(k+1)\Delta$.

A continuum of agents indexed by $i \in [0,1]$ trades the asset in each period. Trading decisions are based on the current dividend and on agents' forecasts of the prices in the following period. Agents form these forecasts as follows. Each agent i categorizes states according to a partition Π^i of Ω that is fixed across all periods. For each state ω, let $\Pi(\omega)$ denote the element of the partition

Π containing ω. In period k, agent i forms a forecast Q^i_{k+1} of the price in period $k+1$ according to

$$Q^i_{k+1} = \frac{\sum_{s<k-1:\omega_s \in \Pi^i(\omega_k)} p_{s+1}}{\sum_{s<k-1:\omega_s \in \Pi^i(\omega_k)} 1}$$

whenever the denominator is nonzero (otherwise take the forecast to be some arbitrary fixed number), where p_s denotes the market price in period s as described below. Thus the price forecast Q^i_{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 Π^i).

Each agent i forms demand α^i_k in period k proportional to her net expected profit from holding the asset for one period:

$$\alpha^i_k = \left(1 - e^{-\Delta}\right) d(\omega_k) + e^{-\Delta} Q^i_{k+1} - p_k.$$

Assuming zero supply, the market clearing price is

$$p_k = \int_i p^i_k di, \tag{16}$$

where p^i_k is agent i's reservation price in period k, defined by

$$p^i_k = \left(1 - e^{-\Delta}\right) d(\omega_k) + e^{-\Delta} Q^i_{k+1}. \tag{17}$$

Let $\Pi_1, \ldots, \Pi_N$ denote those partitions belonging to a positive measure of agents, and denote by π_n the measure of agents using Π_n. Letting p^n_k denote the reservation price of each agent from group n, the market price p_k is

$$p_k = \sum_{n=1}^{N} \pi_n p^n_k.$$

D.2 Steady-state prices

Proposition 5 below shows that this learning process converges to steady-state prices $P : \Omega \longrightarrow \mathbb{R}$ that depend only on the current state. Steady-state prices turn out to be identical to rational expectations prices, not with respect to the true process, but with respect to a different process that reflects both the true process q_Δ and the categorizations used by agents.

Definition 1. Given any Δ, prices $P(\omega)$ are (steady-state) *rational expectations prices* with respect to a Markov process m on Ω and a dividend function d if

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

for every $\omega \in \Omega$.[10]

Let ϕ denote the stationary distribution of states with respect to the true process q. For given initial prices and a given realization of the sequence of states $(\omega_s)_{s=0}^{k-1}$, let $p_k(\omega)$ denote the price in period k that would obtain if $\omega_k = \omega$. Define the *modified process* by

$$m_\Delta(\omega, \omega') = \sum_{n=1}^{N} \pi_n \sum_{\omega'' \in \Pi_n(\omega)} \phi\left(\omega'' \left| \Pi_n(\omega)\right.\right) q_\Delta(\omega'', \omega'). \tag{18}$$

Proposition 5. *For each Δ, the sequence $p_k(\omega)$ almost surely converges to the vector $P_\Delta(\omega)$ of rational expectations with respect to the modified process m_Δ and the dividend function d.*

Proposition 5 is a corollary of Proposition 6 below.

To understand the modified process m_Δ, first consider the case in which all agents distinguish all states, i.e. $\Pi^i(\omega) = \{\omega\}$ for every ω and i. In this this case, $m_\Delta = q_\Delta$, and hence the long-run prices are precisely the rational expectations prices with respect to the true process. To see why, consider the forecasting procedure. In period k, each agent uses data from previous periods $s < k - 1$ in which the state was indistinguishable from the current state (according to her own categorization). For the finest categorization, these relevant periods are those s such that $\omega_s = \omega_k$. In the steady state, the agent's forecast is just the average of $P(\omega_{s+1})$ across the relevant periods s. In the long run, the forecast is equal to $\sum_{\omega'} q_\Delta(\omega_k, \omega')P(\omega')$, coinciding with the rational expectation of the price in the next period.

For general categorizations, a given agent's forecast is based on all previous periods s in which the state ω_s belonged to the current category $\Pi^i(\omega_k)$. In the long run, the average of $P(\omega_{s+1})$ for those values of s is equal to

$$\sum_{\omega'' \in \Pi^i(\omega)} \phi\left(\omega'' \left| \Pi^i(\omega)\right.\right) q_\Delta(\omega'', \omega')P(\omega'),$$

[10] Note that for any m and d, rational expectations prices exist and are unique.

where the term $\phi\left(\omega''\left|\Pi^i(\omega)\right.\right)$ captures the long-run frequency of state ω'' in the sample of relevant periods s. Taking the average across agents, the population-wide forecast is the expectation with respect to the modified process m_Δ in (18).[11]

The next section extends Proposition 5 in two directions. First, we extend the price forecasting rule to a general class of similarity-based rules in which agents forecast using data from similar past states. Unlike the categorization considered here, the weights assigned to different states may vary according to the perceived degree of similarity. Second, we allow for an arbitrary fraction of agents to form rational expectations knowing all parameters of the model, including other agents' forecasting procedures. In the long run, such agents have the same effect on prices as agents who categorize every state separately.

D.3 Proof and Generalizations

This section proves convergence of the above learning process, extends the result to a more general class of processes in which agents learn from similar past states, and shows that our results remain unchanged if we allow for some agents to form rational expectations. We start by describing learning by similarity, which includes categorization as a special case. We then consider an even more general class of processes that is sufficiently broad to allow for the inclusion of agents who form rational expectations about future states and other agents' behavior.

D.3.1 Learning by similarity

The categorization framework of Section D.1 is a special case of a model in which agents learn prices based on past prices in states similar to the current one, but do not necessarily apply equal weight to all similar states. Proposition 5 extends to this more general case.

Each agent i is endowed with a symmetric similarity function $g_i : \Omega \times \Omega \longrightarrow \mathbb{R}_+$ determining the weight assigned to various states in forming forecasts of future prices. We assume that for each i and ω, there exists some ω' such that $g_i(\omega, \omega') \neq 0$. Given a history of states and prices up to

[11] The modified process is closely related to the coarse expectation formation in Eyster and Piccione (2012). Indeed, applying Proposition 5 to a homogenous population gives convergence to Eyster's and Piccione's stationary price function with respect to the agents' categorization.

period $k-1$, agent i's forecast in period k of the price in period $k+1$ is

$$Q^i_{k+1} = \frac{\sum_{s<k-1} g_i(\omega_k, \omega_s) p_{s+1}}{\sum_{s<k-1} g_i(\omega_k, \omega_s)}$$

whenever the denominator is nonzero, and some fixed constant otherwise. Thus the forecast is formed by averaging the one-period-ahead prices in all past states, weighted according to the degree of similarity to the current state. The categorization of Section D.1 is a special case of this framework in which, for each i, g_i takes only the values 0 and 1.

For simplicity, we assume that only a finite number of different similarity functions are used by the agents. That is, there exists a finite partition of the population into groups of measures $\pi_1, \ldots, \pi_N$, and similarity functions $g_1, \ldots, g_N$ such that, for each n, every agent in n's group uses similarity function g_n. As before, each agent's action in each period is given by (17) and the market price in period k is the population-wide average action given by (16).

We show in the next subsection that Proposition 1 carries over directly to this setting except that the modified process m_Δ is defined more generally by

$$m_\Delta(\omega, \omega') = \sum_{n=1}^{N} \pi_n \frac{\sum_{\omega''} g_n(\omega, \omega'') \phi(\omega'') q_\Delta(\omega'', \omega')}{\sum_{\omega''} g_n(\omega, \omega'') \phi(\omega'')}. \tag{19}$$

Within each group, in the steady-state price forecasts, the weight given to each possible state ω' one period ahead is based on the likelihood of transitions to ω' from each state ω'' similar to the current state ω. The weight given to the transition from ω'' to ω' depends on the similarity between ω and ω'' together with the frequency $\phi(\omega'')$ with which state ω'' occurs. The aggregate distribution m_Δ is obtained by averaging the individual distributions across all agents. Note that, as before, agents should be interpreted as behaving, in the long-run, *as if* they believe (on average) that the state evolves according to m_Δ; agents do not literally hold these beliefs.

D.3.2 Proof of Proposition 5

The general learning process is as follows. The state space Ω and the true process q are as in the main text. Without loss of generality, let $\Delta = 1$. Let $\boldsymbol{\omega}^k = (\omega_s)_{s=0}^k$ denote the finite history of states up to period k, and $\mathbf{p}^k = (p_s)_{s=0}^{k-1}$ be the history of prices up to period $k-1$. We assume that all prices lie in a bounded interval $[\underline{p}, \overline{p}]$. The price p_k in period k is determined according to

$$p_k = (1-\rho)d(\omega_k) + \rho Q\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right), \tag{20}$$

where $Q : \bigcup_k \left(\Omega^k \times [\underline{p}, \overline{p}]^{k-1}\right) \longrightarrow [\underline{p}, \overline{p}]$ can be interpreted as the average forecast of the price in period $k+1$ and $\rho = e^{-1}$ is the discount factor.

We assume that Q satisfies the following condition.

A1. There exists a continuous monotone function

$$\mathcal{E} : [\underline{p}, \overline{p}]^\Omega \longrightarrow [\underline{p}, \overline{p}]^\Omega$$

such that, for any $\underline{P}, \overline{P} \in [\underline{p}, \overline{p}]^\Omega$, any K, and any $\varepsilon > 0$, if

$$\Pr\left(p_k \in [\underline{P}, \overline{P}]\; \forall k > K\right) > 1 - \varepsilon, \tag{21}$$

then for any $\delta > 0$ there exists K' such that, for each ω,

$$\Pr\left(Q\left(\left(\boldsymbol{\omega}^k, \omega\right), \mathbf{p}^k\right) \in \left(\mathcal{E}(\underline{P})(\omega) - \delta, \mathcal{E}(\overline{P})(\omega) + \delta\right)\; \forall k > K'\right) > 1 - \varepsilon - \delta. \tag{22}$$

In the case of a homogeneous population using similarity function g, the learning process from Section D.3.1 (and hence also the categorization-based learning from Section D.1) is captured by

$$Q^{\text{sim}}\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right) = \begin{cases} \frac{\sum_{s<k-1} g(\omega_k, \omega_s) p_{s+1}}{\sum_{s<k-1} g(\omega_k, \omega_s)} & \text{if } \sum_{s<k-1} g(\omega_k, \omega_s) > 0, \\ p_0 & \text{otherwise,} \end{cases}$$

where p_0 is arbitrary. For a heterogeneous population, Q is obtained by aggregating the values of Q^{sim} across groups (see Lemma 6).

Lemma 5. *For any similarity function g, Q^{sim} satisfies A1 with*

$$\mathcal{E}(P)(\omega) = \sum_{\omega'} m_\Delta(\omega, \omega') P(\omega'),$$

where m_Δ is the modified process in (19).

Proof. We prove only the upper bound; the proof for the lower bound is similar. Suppose that for some K, $\varepsilon > 0$, and $\overline{P}$, $\Pr\left(p_k \leq \overline{P}\,\forall k > K\right) > 1 - \varepsilon$. Given any $\delta > 0$ and $\gamma > 0$, by the Law of Large Numbers, there exists some $K' > K$ such that, with probability greater than $1-\delta$, for every pair (ω', ω'') and every $k > K'$, the fraction of periods $s < k$ such that $(\omega_s, \omega_{s+1}) = (\omega', \omega'')$ lies in $(\phi(\omega')q_\Delta(\omega', \omega'') - \gamma, \phi(\omega')q_\Delta(\omega', \omega'') + \gamma)$. Since the process q_Δ is ergodic, we can choose K' such that this property holds regardless of the history $\boldsymbol{\omega}^K$. Furthermore, for $K' > K/\gamma$, $p_s \leq \overline{P}(\omega_s)$ for a fraction of at least $1 - \gamma$ periods $s \leq k$ with probability greater than $1 - \varepsilon$, in which case the average of the prices p_{s+1} across those periods s such that $(\omega_s, \omega_{s+1}) = (\omega', \omega'')$ is at most $(1-\gamma)\overline{P}(\omega'') + \gamma\overline{p}$. Hence for $k > K'$, we have

$$Q^{\text{sim}}\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right) \leq \frac{\sum_{\omega',\omega''} g(\omega_k, \omega'')(\phi(\omega'')q_\Delta(\omega'', \omega') + \gamma)((1-\gamma)\overline{P}(\omega') + \gamma\overline{p})}{\sum_{\omega''} g(\omega_k, \omega'')(\phi(\omega'') - \gamma)}$$

with probability greater than $1 - \varepsilon - \delta$. Given $\delta > 0$, we can choose $\gamma > 0$ sufficiently small so that the right-hand side of the preceding inequality is less than $\mathcal{E}(\overline{P})(\omega_k) + \delta$, as needed. □

In the main text, agents differ in their forecasting procedures and the price is determined by the average of agents' forecasts. The following lemma indicates that A1 aggregates across heterogeneous groups.

Lemma 6. *Suppose that a fraction π_n of the population use prediction rule Q^n, with $\sum_{n=1}^N \pi_n = 1$. Suppose moreover that all rules Q^n satisfy A1 with functions $\mathcal{E}^n(P)$ respectively. Finally assume that price evolution is governed by* (20) *with prediction rule $Q = \sum \pi_n Q^n$. Then Q satisfies A1 with $\mathcal{E}(P) = \sum_n \pi_n \mathcal{E}^n(P)$.*

Proof. Using the property of A1 with $\pi_n\delta$ for each subpopulation and taking the maximum of the K' needed for each process gives the result. □

Proposition 5 is a special case of the following convergence result.

Proposition 6. *If Q satisfies A1, prices are determined according to* (20)*, and the mapping* $(1-\rho)d + \rho\mathcal{E}$ *is a contraction (with respect to some metric) then prices almost surely converge to the unique fixed point of* $d + \rho\mathcal{E}$.

In the case of learning by similarity, $(1-\rho)d+\rho\mathcal{E}$ is a contraction with respect to the sup norm, and we therefore obtain convergence to a unique price profile, proving Proposition 5.

Proof of Proposition 6. The mapping $(1-\rho)d+\rho\mathcal{E}$ has extreme fixed points $\underline{P}^*$, $\overline{P}^*$: for every fixed point P^*, we have $\underline{P}^* \leq P^* \leq \overline{P}^*$. This follows immediately from Tarski's Fixed Point Theorem since $[\underline{p},\overline{p}]^\Omega$ is a complete lattice and $(1-\rho)d+\rho\mathcal{E}$ is continuous and monotone.

We will prove that for each ω, the set of cluster points of $(p_k(\omega))_k$ is almost surely contained in $[\underline{P}^*(\omega), \overline{P}^*(\omega)]$. The proposition follows immediately since the fixed point is unique when $(1-\rho)d+\rho\mathcal{E}$ is a contraction.

We prove only that the cluster points are almost surely at most $\overline{P}^*(\omega)$. The proof of the lower bound is similar.

Let $\overline{P}_0 = \overline{p}\mathbf{1}$, where $\mathbf{1}$ denotes the vector with a 1 in each component, and for $l \in \mathbb{N}_+$, let $\overline{P}_l = (1-\rho)d + \rho\mathcal{E}(\overline{P}_{l-1})$. Since $\overline{P}_l$ is nonincreasing in l, $\lim_l \overline{P}_l$ exists and is a fixed point of $(1-\rho)d+\rho\mathcal{E}$ (by continuity of $\mathcal{E}$).

Note that $p_k \leq \overline{P}_0(\omega_k)$ for each $k > 0$. Suppose for induction that, given any $\varepsilon > 0$, there exists K_l such that

$$\Pr\left(p_k < \overline{P}_l(\omega_k) + \varepsilon \text{ for all } k > K_l\right) > 1-\varepsilon.$$

We will show that the same condition holds when each l is replaced with $l+1$.

For any $\delta > 0$, combining A1 with the inductive hypothesis, there exists some K_{l+1} such that

$$\Pr\left(Q\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right) < \mathcal{E}(\overline{P}_l + \varepsilon\mathbf{1})(\omega_k) + \delta \;\forall k > K_{l+1}\right) > 1-\varepsilon-\delta.$$

Substituting for $Q\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right)$ using (20), we have

$$\Pr\left(p_k < (1-\rho)d(\omega_k) + \rho\mathcal{E}(\overline{P}_l + \varepsilon\mathbf{1})(\omega_k) + \delta \;\forall k > K_{l+1}\right) > 1-\varepsilon-\delta.$$

Given any $\gamma > 0$, since $\mathcal{E}$ is continuous, there exist some $\varepsilon, \delta \in (0,\gamma)$ such that, for each ω, $\rho\mathcal{E}(\overline{P}_l + \varepsilon\mathbf{1})(\omega) + \delta < \rho\mathcal{E}(\overline{P}_l)(\omega) + \gamma$. Since ε and δ are arbitrary, we have that, for some K_{l+1},

$$\Pr\left(p_k < (1-\rho)d(\omega_k) + \rho\mathcal{E}(\overline{P}_l)(\omega_k) + \gamma \;\forall k > K_{l+1}\right) > 1-\gamma.$$

Since $\overline{P}_{l+1} = (1-\rho)d + \rho\mathcal{E}(\overline{P}_l)$, this completes the proof of the inductive step. □

D.3.3 Presence of rational agents

We now consider a setting in which some agents form rational expectations. For simplicity, we assume that the population consists of two parts. A fraction π of agents are rational while the remaining $1-\pi$ are coarse thinkers who use a prediction rule Q^C satisfying A1. Rational agents know Q^C and the underlying Markov process, and form rational expectations of the forecasts formed by coarse thinkers in the next period. The rational agents' prediction rule Q^R satisfies

$$Q^R\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right) = E\left[(1-\rho)d(\omega_{k+1}) + \rho\left((1-\pi)Q^C\left(\boldsymbol{\omega}^{k+1}, \mathbf{p}^{k+1}\right) + \pi Q^R\left(\boldsymbol{\omega}^{k+1}, \mathbf{p}^{k+1}\right)\right)\middle|\boldsymbol{\omega}^k\right]. \tag{23}$$

This equation implies that rational agents correctly predict prices given the history to date and the prediction rules used by other agents.

While the model in the main text does not include rational agents, it does allow for some agents to perfectly distinguish among states. In the present setting, these agents are not rational insofar as their price forecasts are based only on past data and do not explicitly account for other agents' forecasts. We show here that, in the long-run, the difference between these agents and rational agents is immaterial. Long-run prices are identical if we replace any share of agents using the finest categorization with agents who form rational expectations.

A2. For each $\omega \in \Omega$, $K \in \mathbb{N}$, and almost every $\boldsymbol{\omega} \in \Omega^{\mathbb{N}}$,

$$\lim_{k\to\infty}\left(Q^C\left(\left(\boldsymbol{\omega}^k, \omega\right), \mathbf{p}^k\right) - Q^C\left(\left(\boldsymbol{\omega}^{k+K}, \omega\right), \mathbf{p}^{k+K}\right)\right) = 0,$$

where, for each κ, $\boldsymbol{\omega}^\kappa$ denotes the projection of $\boldsymbol{\omega}$ onto its first κ components.

Roughly speaking, A2 says that data from a fixed finite number of recent periods eventually has little impact on forecasts once the total quantity of data is large. Note that A2 is satisfied by the similarity-based learning procedure of Section D.3.1.

Proposition 7. *Suppose that a fraction π of the population form rational expectations, and the remaining $1-\pi$ use a prediction procedure Q^C satisfying A1 with bound $\mathcal{E}^C(P)$ and A2. Suppose further that the mapping $(1-\rho)d + \rho\mathcal{E}^C(P)$ is a contraction. Then the price vector $P(\omega)$ almost* *surely converges to the unique solution of*

$$P(\omega_k) = (1-\rho)d(\omega_k) + \rho\left(\pi E\left[P\left(\omega_{k+1}\right) \mid \omega_k\right] + (1-\pi)\mathcal{E}^C(P)(\omega_k)\right).$$

Lemma 7. *If Q^C satisfies A1 with bound $\mathcal{E}^C$ and A2 then Q^R satisfies A1 with bound*

$$\mathcal{E}^R(P)(\omega_k) = E\left[\sum_{l=1}^{\infty}(\pi\rho)^{l-1}(1-\rho)d(\omega_{k+l}) + (1-\pi)\rho\sum_{l=1}^{\infty}(\pi\rho)^{l-1}\mathcal{E}^C(P)(\omega_{k+l})\middle|\,\omega_k\right]. \tag{24}$$

Proof of Lemma 7. Iterating (23) gives

$$Q^R\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right) = E\left[\sum_{l=1}^{\infty}(\pi\rho)^{l-1}(1-\rho)d(\omega_{k+l}) + (1-\pi)\rho\sum_{l=1}^{\infty}(\pi\rho)^{l-1}Q^C\left(\boldsymbol{\omega}^{k+l}, \mathbf{p}^{k+l}\right)\middle|\,\boldsymbol{\omega}^k\right].$$

We need to show that for any $\underline{P}, \overline{P} \in [\underline{p}, \overline{p}]^{\Omega}$, any K, and any $\varepsilon > 0$, if condition (21) holds, then for any $\delta > 0$ there exists K' such that (22) holds for Q^R. We prove only the upper bound; the proof of the lower bound is similar.

Accordingly, suppose that (21) holds for some $\varepsilon > 0$ and K. Fix $\delta > 0$. Since Q^C and $\mathcal{E}^C$ are bounded, there exists M such that, for every $\boldsymbol{\omega}^k$ and $\mathbf{p}^k$

$$\begin{aligned} &Q^R\left(\boldsymbol{\omega}^k, \mathbf{p}^k\right) < \\ &E\left[\sum_{l=1}^{\infty}(\pi\rho)^{l-1}(1-\rho)d(\omega_{k+l}) + (1-\pi)\rho\left(\sum_{l=1}^{M}(\pi\rho)^{l-1}Q^C\left(\boldsymbol{\omega}^{k+l}, \mathbf{p}^{k+l}\right) + \sum_{l=M+1}^{\infty}(\pi\rho)^{l-1}\mathcal{E}^C(\overline{P})(\omega_{k+l})\right)\middle|\,\boldsymbol{\omega}^k\right] \\ &\qquad + \delta/3. \end{aligned} \tag{25}$$

Since Q^C satisfies A1, there exists some K' such that, for each ω,

$$\Pr\left(Q^C\left(\left(\boldsymbol{\omega}^{k-1}, \omega\right), \mathbf{p}^k\right) < \mathcal{E}^C(\overline{P})(\omega) + \delta/3M \;\forall k > K'\right) > 1 - \varepsilon - \delta/2. \tag{26}$$

By A2, there exists some K'' such that, for each $l = 1, \ldots, M$,

$$\Pr\left(Q^C\left(\boldsymbol{\omega}^{k+l}, \mathbf{p}^{k+l}\right) < Q^C\left(\left(\boldsymbol{\omega}^{k-1}, \omega_{k+l}\right), \mathbf{p}^k\right) + \delta/3M \;\forall k > K'' \mid \boldsymbol{\omega}^k\right) > 1 - \delta/2M. \tag{27}$$

Combining (26) and (27) gives

$$\Pr\left(Q^C\left(\boldsymbol{\omega}^{k+l},\mathbf{p}^{k+l}\right) < \mathcal{E}^C(\overline{P})(\omega_{k+l}) + 2\delta/3M\ \forall k > \max\{K',K''\},\ \forall l = 1,\ldots,M \mid \boldsymbol{\omega}^k\right) > 1-\varepsilon-\delta.$$

Combining the last inequality with (25) gives

$$\Pr\left(Q^R\left(\boldsymbol{\omega}^k,\mathbf{p}^k\right) < \mathcal{E}^R(\overline{P})(\omega_k) + \delta\ \forall k > \max\{K',K''\}\right) > 1-\varepsilon-\delta,$$

as needed. □

Proof of Proposition 7. Combining Lemma 6, Lemma 7, and Proposition 6, the cluster points lie between the extremal solutions to

$$P = (1-\rho)d + \rho\left(\pi\mathcal{E}^R(P) + (1-\pi)\mathcal{E}^C(P)\right).$$

Substituting for $\mathcal{E}^R(P)$ using (24) leads to

$$\begin{aligned}
P(\omega_k) &= (1-\rho)d(\omega_k) + \rho\pi E\left[\sum_{l=1}^{\infty}(\pi\rho)^{l-1}(1-\rho)d(\omega_{k+l}) + (1-\pi)\rho\sum_{l=1}^{\infty}(\pi\rho)^{l-1}\mathcal{E}^C(P)(\omega_{k+l})\middle|\,\omega_k\right] \\
&\quad + \rho(1-\pi)\mathcal{E}^C(P)(\omega_k) \\
&= (1-\rho)d(\omega_k) + \rho\pi E\left[E\left[\sum_{l=1}^{\infty}(\pi\rho)^{l-1}(1-\rho)d(\omega_{k+l}) + (1-\pi)\rho\sum_{l=1}^{\infty}(\pi\rho)^{l-1}\mathcal{E}^C(P)(\omega_{k+l})\middle|\,\omega_{k+1}\right]\middle|\,\omega_k\right] \\
&\quad + \rho(1-\pi)\mathcal{E}^C(P)(\omega_k) \\
&= (1-\rho)d(\omega_k) + \rho\left(\pi E[P(\omega_{k+1})|\omega_k] + (1-\pi)\mathcal{E}^C(P)(\omega_k)\right),
\end{aligned}$$

where the second last equality follows from the Law of Iterated Expectations, and the final equality uses the first equality with ω_{k+1} in place of ω_k. □

References

Al-Najjar, N. and M. Pai (2009). Coarse decision making. *Northwestern Working Paper*.

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 (2012). An approach to asset-pricing under incomplete and diverse perceptions. Forthcoming in *Econometrica*.

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.

Gonzalez, R. and G. Wu (1999). On the shape of the probability weighting function. *Cognitive Psychology 38*(1), 129–166.

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. *The 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.

Prelec, D. (1998). The probability weighting function. *Econometrica 66*, 497–527.

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.