DOUBTS, INEQUALITY, AND BUBBLES

IN-KOO CHO AND KENNETH KASA

Abstract. Two agents trade an indivisible asset. They are risk neutral and share a common benchmark dividend model. However, each has doubts about the specification of this model. These doubts manifest themselves as a preference for robustness (Hansen and Sargent (2008)). Robust preferences introduce pessimistic drift distortions into the benchmark dividend process. These distortions increase with the level of wealth, and give rise to endogenous heterogeneous beliefs. Belief heterogeneity allows asset price bubbles to emerge, as in Scheinkman and Xiong (2003). A novel implication of our analysis is that bubbles occur when wealth inequality increases. Empirical evidence supports this prediction. Detection error probabilities suggest that the implied degree of belief heterogeneity is empirically plausible.

JEL Classification Numbers: D84, G12

1. Introduction

Conventional wisdom blames the most recent financial crisis on a bubble in the housing market. Bubbles have been blamed for previous crises as well. Economists do not agree on what causes bubbles. Some argue the concept is meaningless, or that bubbles are a case of hindsight being 20/20. Others point out that bubbles can be perfectly rational in a world where current outcomes depend on expectations of future outcomes. Still others argue that bubbles are evidence of irrationality, and suggest that economists build models based on insights from psychology and sociology.

Each of these approaches to bubbles has problems. Denying that bubbles exist presents the challenge of explaining why prices rise so much and fall so quickly. Although Garber (2000) points to the possibility that real-time assessments of fundamentals could have justified price increases in several bubble episodes, these examples do not explain their magnitude, nor do they explain why some people buy while others sell, or why the bubble suddenly bursts. Theories of 'rational bubbles' show that the conditions supporting the existence of bubbles tend to be quite fragile (Santos and Woodford (1997)). Rational bubble theories also do not explain why bubbles get started in the first place, or why they are correlated with large trading volumes. Finally, models based on psychology or sociology raise questions about testability and consistency with the way economists interpret other market data.

Our paper is based on a model of bubbles proposed by Scheinkman and Xiong (2003). Their model builds on the previous work of Harrison and Kreps (1978). It has two key

---

Date: January, 2018.

Cho: Department of Economics, University of Illinois, `inkoocho@illinois.edu`. Kasa: Department of Economics, Simon Fraser University, `kkasa@sfu.ca`. We would like to thank Tom Sargent for suggesting this idea. Chris Gibbs, Xiaowen Lei, and Jose Scheinkman provided useful comments.

ingredients: (1) Heterogeneous beliefs, and (2) Short-selling constraints. Heterogeneous beliefs imply the coexistence of optimists and pessimists, while the presence of short-selling constraints means that the views of optimists are more fully expressed than those of pessimists. Hence, prices tend to be biased up on average, even when beliefs are unbiased.[1] When relative optimism fluctuates, an asset holder recognizes that at some point in the future, other traders may become relatively more optimistic. This creates an *option value* of selling to future optimists. Scheinkman and Xiong (2003) define a bubble as the value of this option, and then apply well known tools from the options pricing literature to formally calculate it as a function of the model's underlying parameters. The fact that they are able to quantify both the timing and magnitude of bubbles is what sets Scheinkman and Xiong's work apart from previous efforts to understand bubbles, because it makes their theory testable. For example, a key prediction of their model is that bubbles should be accompanied by large trading volume.

However, Scheinkman and Xiong (2003) is not without flaws, as Scheinkman (2014) himself acknowledges. One drawback concerns the underlying source of belief heterogeneity. In their model, belief heterogeneity is exogenous. Agents receive exogenous signals about future fundamentals, and are assumed to over-react to distinct signals. This makes relative optimism fluctuate. Hong and Scheinkman (2008) motivate this assumption with a model of financial advisors. Still, even though by construction beliefs are unbiased on average, one could argue that the magnitude and fluctuations in relative optimism lack discipline in the sense that they are not linked to observable data (other than the bubble itself). Another drawback is that Scheinkman-Xiong abstract from learning. Over-reaction in their model is caused by agents thinking a useless signal is in fact useful. Wouldn't agents eventually learn their signal is useless? Why don't beliefs merge (Morris (1996))?

Our paper addresses these shortcomings. Like Scheinkman-Xiong, we assume agents are risk-neutral. However, they are *not* 'uncertainty neutral'. They have doubts about the (common) benchmark model for fundamentals. In his survey, Xiong (2013) notes,

> *It is difficult to precisely identify a bubble, as any valuation requires a model of fundamentals, which would be subject to concerns about potential model errors.* (Xiong (2013, p. 2))

The agents in our model share Xiong's doubts about model specification. Their doubts are expressed as pessimistic drift distortions that are proportional to the marginal value of wealth (Hansen, Sargent, Turmuhambetova, and Williams (2006)). Hence, if wealth differs across agents, so do beliefs. Since wealth is endogenous, so are belief differences. As a corollary, an intriguing prediction of our model is that bubbles emerge during periods of increasing wealth inequality.

Like Scheinkman-Xiong, the key mechanism in our model is belief reversals, which create an option to resell to future optimists. Although utility functions are linear, value functions are not, as they embody the option value of selling the asset in the future. This resale option makes the value function convex. Convexity implies that drift distortions *increase* with wealth. Effectively, the agent who owns the asset becomes more pessimistic as prices rise and his wealth increases. This makes sense since he has more to lose following a price decline. This endogenously increasing pessimism is the source of belief reversals in

[1]This point was first made by Miller (1977). However, his model was static.

our model. Although both agents are pessimistic relative to the (untrusted) benchmark model, the agent who owns the asset gets progressively more pessimistic. Eventually, he becomes even more pessimistic than the agent who sold him the asset in the first place, and he exercises his option to sell.

Scheinkman (2014), Garber (2000), and Shiller (2015) all emphasize that historical bubble episodes coincided with the introduction of new technologies or new markets. They argue that novelty breeds disagreement. In our model, technology and markets don't change. However, a similar dynamic emerges. Rapid growth in fundamentals enriches asset owners, which raises concerns about downside risk, and owners eventually decide to cash out. Data generated by our model exhibit an apparently strong over-reaction to positive growth in fundamentals. This apparent over-reaction is not irrational. It simply refects the increased option value of selling the asset, which in turn is governed by the endogenously evolving belief differences of asset owners and nonowners.

Another advantage of our approach is that it imparts discipline on the degree of belief heterogeneity. Following Anderson, Hansen, and Sargent (2003), we link belief distortions to detection error probabilities. Asset owners become increasingly pessimistic as prices rise, but they do not become paranoid. They construct a defensive worst-case scenario that cannot be statistically rejected using observed data. Pessimism here means that agents overstate mean reversion. That is, they believe 'good times' are more transitory than they really are.[2] In our benchmark calibration, a modest decline in the fundamental's AR parameter from .98 to .97 is enough to trigger a large increase in resale option value and price, and this modest distortion would be very difficult to detect using historical data. Even with 150 years of data, the detection error probability exceeds 40%.

The remainder of the paper is organized as follows. Section 2 provides motivation by quickly reviewing data on bubbles. We show that bubbles are accompanied by growing inequality. Section 3 outlines the model. We characterize a competitive equilibrium by first solving a planner's problem. As in Scheinkman-Xiong, this is an optimal stopping problem. The new twist here is that Pareto weights are endogenous. We then discuss how the planner's optimal asset allocation sequence can be decentralized via asset trade. Section 4 presents simulation results. These simulations reveal that the resale option component of asset prices is typically less than 20% of the observed price, but that during bubble episodes it often exceeds 40% of the price. Section 5 provides simulation evidence on detection error probabilities. We argue that empirically plausible doubts about fundamentals can easily account for observed bubble episodes. Section 6 briefly discusses related literature. We compare our approach to models based on portfolio rebalancing and state-dependent risk aversion. We also argue that wealth dependent pessimism offers a potential explanation of the 'disposition effect' (Shefrin and Statman (1985)). Section 7 concludes, and offers a few suggestions for future work. A technical appendix contains proofs and derivations.

2. Motivation

Our model makes several testable predictions about wealth inequality, belief heterogeneity, and asset prices. We begin by briefly discussing some informal evidence in support of

[2]Note, given the option embedded in the price, distortions are not symmetric. What matters is mean reversion during good times.

these predictions. First, our model predicts that in the presence of short sale constraints belief heterogeneity increases asset prices. Since this prediction is shared by many heterogeneous beliefs models, it has been subject to many tests. Most use analyst forecast dispersion as a proxy for belief heterogeneity, and find that stocks with greater forecast dispersion do indeed earn lower future returns (Diether, Malloy, and Scherbina (2002)). Boehme, Danielsen, and Sorescu (2006) find that results are stronger and more robust when you control for short-selling costs. These results are only suggestive, of course, since what matters are the beliefs of investors. In particular, our model predicts that belief heterogeneity is determined by wealth inequality, which in turn depends on asset ownership. Vissing-Jorgensen (2003) reports evidence that is more closely related to this mechanism. She uses UBS/Gallup household survey data for the period 1998-2003. To qualify for the survey, households must hold a threshold level of financial assets. She finds that the cross-sectional standard deviation of expected returns increases with the level of stock prices. In our model, belief heterogeneity is both a cause and effect of higher stock prices.

Our model links belief heterogeneity to wealth inequality. Since belief heterogeneity increases resale option values, we should see evidence of asset market bubbles during periods of growing inequality. Figure 1 presents evidence in support of this prediction. It plots the top 0.1% wealth share in the USA (as reported in Saez and Zucman (2016)) against Shiller's cyclically adjusted Price/Earnings ratio for the period 1920-2012.

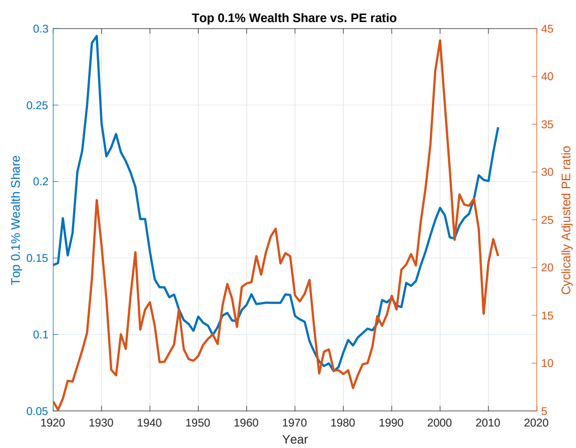

FIGURE 1. Top 0.1% Wealth Share vs. Price/Earnings Ratio (1920-2012)

Stock market bubbles are often defined as periods when the PE ratio rises significantly above its long-run mean of about 17. Based on this definition, there were three bubble episodes during this roughly 100 year period: (1) The late 1920s, when the PE ratio rose above 25, (2) The late 1960s, when the PE ratio again reached 25, and (3) The dotcom boom of the late 1990s, when the PE ratio hit its all-time high of nearly 45. Notice that during all three episodes inequality was increasing. During the late 1920s, the top 0.1% wealth share hit its all-time high of nearly 30%. It then fell steadily throughout the Depression and WWII, as the stock market declined, but then experienced a mild uptick during the stock market boom of the 1960s. During the 1970s, both inequality and the stock market fell. The top 0.1% wealth share hit its all-time low of 7% at the end of the 1970s. However, since the 1980s, inequality has been steadily increasing. By 2012, the top 0.1% wealth share had reached nearly 23%, close to its pre-Depression peak. For most of this period the stock market rose as well. However, the collapse of the dotcom bubble in 2000 along with the continuing rise of inequality, seems to be a break in the pattern.[3]

Bubbles are not confined to stock markets. Many argue that housing markets are especially prone to bubbles. Housing markets are a better fit to our model as well, since it assumes the underlying asset is indivisible. At the same time, however, housing is a more disaggregated and localized market. Although local house price indices are available, local measures of inequality are scarce. As a compromise, we report data at the US state level. Sommeiller, Price, and Wazeter (2016) report time series data on the top 1% income share. Unfortunately, state level wealth share data do not exist. For housing prices we use data from the Federal Housing Finance Agency, which constructs an All Transactions Index of state housing prices going back to the mid 1970s. To make the results more robust, we rank each state according to the percentage increase in housing prices it experienced from 1980 to the first quarter of 2007, at the peak of the most recent housing bubble. Likewise, we rank each state according to the percentage point increase it experienced in the top 1% income share during the same period. Figure 2 depicts the results for all 50 states plus Washington DC.

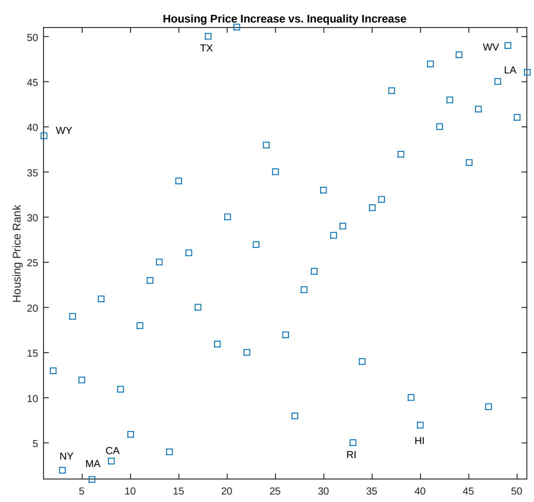

FIGURE 2. Housing Price Increase vs. Increase in Top 1% Income Share

The bottom left shows states that experienced both the largest increase in house prices and the largest increase in inequality. Not surprisingly, New York, California, and Massachussetts are here. The top right shows states that experienced the smallest increase in

[3]It is worth noting that if the sample were extended to the current date (late 2017), recent trends would look better, since the PE ratio has again reached 30.

house prices and the smallest increase in inequality. Again not surprisingly, West Virginia and Louisiana are here. Overall, the pattern is clear. Those states with the largest increase in inequality also experienced the largest increases in house prices. The Spearman rank correlation is .51, with a standard error of .09. There are a few notable outliers, however. Wyoming actually experienced the largest increase in inequality (although in terms of levels it is still below New York and Connecticut). At the same time, it experienced a relatively mild increase in house prices. Similarly, Texas experienced a relatively rapid increase in inequality, but as is well known, it largely escaped the housing bubble. On the other side, Hawaii and Rhode Island both experienced rapid increases in house prices, but growth in income inequality was relatively mild.

Finally, the reason wealth inequality induces belief heterogeneity is that rising wealth increases an agent's doubts about fundamentals. A preference for robustness makes asset owners increasingly 'nervous' as the value of their asset increases, simply because they now have more to lose. Owners respond to this by injecting a more pessimistic drift distortion into fundamentals. Is there any evidence for this? Later we shall briefly discuss how this mechanism is related to the widely documented 'disposition effect', but for now we simply report a suggestive piece of evidence based on Google Trends data.[4] Figure 3 plots USA Google Trends data for the keyword search "housing bubble" against Shiller's USA (real) housing price index. The house price data begin in 2002, but Google Trends data do not begin until 2004.

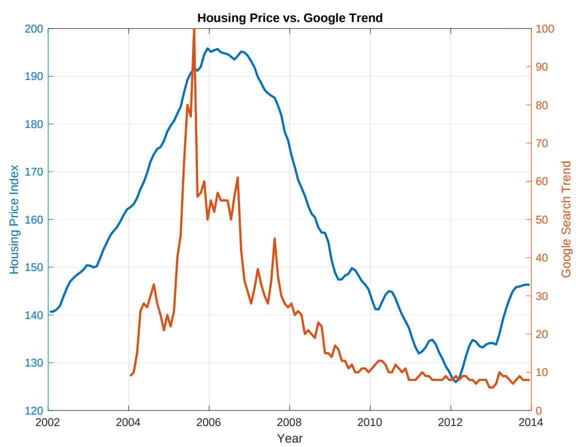

Figure 3. USA Housing Prices vs. Google Trends Data on "Housing Bubble"

Notice that the peak in "housing bubble" searches coincides almost exactly with the peak of the housing bubble. Of course, this proves nothing, since we have no idea who is doing these searches or why. Still, it is at least suggestive.

This section has shown that bubbles appear to coincide with rising wealth inequality. In a sense, there is no mystery here. An exogenous bubble will likely increase inequality

[4]We thank Chris Gibbs for suggesting this. In a conference discussion of our paper, he presented similar data for Sydney, Australia.

simply because asset owners tend to be relatively wealthy already. What's new here is to suggest that causation might also go the other way. Wealth inequality by itself can inflate bubbles since it creates belief heterogeneity, which in turn creates a resale option value. To sort out this bidirectional causality we need an explicit, quantifiable model.

3. The Model

3.1. **Fundamentals.** Two (groups of) risk neutral agents trade claims on an indivisible asset. The asset is in fixed supply, normalized to one. The asset yields a stream of nonstorable dividends governed by the diffusion process:[5]

$$dx = \alpha(b - x) \cdot dt + \sigma\sqrt{x} \cdot dB \tag{3.1}$$

This process generates a stationary distribution, with Gamma density

$$f_\infty(x) = \frac{\omega^z}{\Gamma(z)} x^{z-1} e^{-\omega x} \qquad z = \frac{2\alpha b}{\sigma^2} \qquad \omega = \frac{2\alpha}{\sigma^2}$$

The long-run mean is b, and α governs the strength of mean reversion. An attractive feature of this process is that fundamentals (and equilibrium prices) remain non-negative. When $2\alpha b/\sigma^2 \geq 1$, 0 is an entrance boundary and the process remains strictly positive.

Unlike in Scheinkman-Xiong, this stream of dividends is *observable*. Heterogeneous beliefs are instead determined by (Knightian) uncertainty about the dividend process in (3.1). That is, (3.1) is viewed as merely a useful benchmark model.

3.2. **Preferences.** Each agent has the following risk-neutral preferences:

$$V(x_0) = \hat{E}_0 \int_0^\infty 2\rho x e^{-\rho t} dt \tag{3.2}$$

where ρ is the rate of time preference. The scaling by 2ρ is just a convenient normalization. The key aspect of these preferences is the hat over the expectations operator, $\hat{E}$. This reflects the fact that agents use a distorted probability measure. Rather than commit to a single model/prior, agents entertain a *set* of alternative models, and then optimize against the worst-case model. To prevent agents from being unduly pessimistic, in the sense that they hedge against empirically implausible alternatives, the hypothetical 'evil agent' who selects the worst-case model must pay a penalty that is proportional to the relative entropy between the benchmark model and the worst-case model.

To operationalize this, let q_t^0 be the probability measure defined by the Brownian motion process in the benchmark model (3.1), and let q_t be some alternative probability measure, defined by some competing model. The (discounted) relative entropy between q_t and q_t^0 is then defined as follows:[6]

$$\mathcal{R}(q) = \int_0^\infty e^{-\rho t} \left[\int \log \left(\frac{dq_t}{dq_t^0} \right) dq_t \right] dt \tag{3.3}$$

---

[5]This process is alternatively known as a 'square-root process', or a 'Cox-Ingersoll-Ross' process.

[6]See Hansen, Sargent, Turmuhambetova, and Williams (2006) for a detailed discussion of robust control in continuous-time models, and in particular, on the role of discounting in the definition of relative entropy.

Hence, $\mathcal{R}(q)$ is just an expected log-likelihood ratio statistic, with expectations computed using the distorted probability measure. It can also be interpreted as the Kullback-Leibler 'distance' between q_t and q_t^0. From Girsanov's Theorem we have

$$\int \log\left(\frac{dq_t}{dq_t^0}\right) dq_t = \frac{1}{2}\hat{E}\int_0^t |h_s|^2 ds$$

where $\hat{E}$ denotes expectations with respect to the distorted measure q_t, and h_s represents a square-integrable process that is progressively measureable with respect to the filtration generated by q_t. It provides a convenient parameterization of alternative models.

Later, when solving the planner's problem, it will be convenient to evaluate expectations using the true (but unknown to the agents) probability measure, q_t^0, and to instead view agents' beliefs as being subject to 'preference shocks', z_t. From Girsanov's Theorem,

$$\hat{E}(x) = \int x dq = E(zx) = \int zx dq^0$$

where

$$z_t = 1 + \int_0^t z_s h_s dB_s$$

is a martingale under the true probability measure, q^0. This allows us to restate the agent's preferences using the undistorted probability measure

$$V(x_0) = \min_h E_0 \int_0^\infty z_t \left[2\rho x_t + \frac{1}{2\varepsilon}h_t^2\right] e^{-\rho t} dt \tag{3.4}$$

subject to

$$dz_t = z_t h_t dB_t \qquad z_0 = 1 \tag{3.5}$$

The key parameter here is ε. It penalizes the evil agent's distortions. If ε is small, the evil agent pays a big cost when distorting the benchmark model. In the limit, as $\varepsilon \to 0$, the evil agent sets $h = 0$, and agents no longer have doubts about the model.

3.3. **Single-Agent Competitive Equilibrium.** As a reference, it is useful to first compute the single-agent/buy-and-hold equilibrium, in which there is no possibility of trade. In this case, the equilibrium price is just the agent's value function.[7] The Hamilton-Jacobi-Bellman (HJB) equation for the problem in (3.4)-(3.5) is

$$\rho V(z,x) = \min_h \left\{ z\left[2\rho x + \frac{1}{2\varepsilon}h^2\right] + \alpha(b-x)V_x + \frac{1}{2}\sigma^2 x V_{xx} + \frac{1}{2}(zh)^2 V_{zz} + \sigma\sqrt{x} z h V_{xz}\right\} \tag{3.6}$$

By inspection, it is clear that V is homogeneous in z and we can write $V(z,x) = z\hat{V}(x)$, and then solve an ODE for $\hat{V}(x)$. Unfortunately, this ODE is still nonlinear, and cannot be solved in closed-form. However, ε is a natural perturbation parameter, since we know the solution when $\varepsilon = 0$. The details are relegated to the Appendix, and we merely state the result,

[7]The could also be interpreted as the outcome when the asset is divisible, and agents are able to achieve and maintain a perfectly pooled equilibrium. In this case, wealth and relative beliefs are invariant.

Proposition 3.1. *A first-order perturbation approximation of the buy-and-hold price is given by*

$$P(z,x) = \frac{2z}{\rho+\alpha}\left[(\alpha b+\rho x)\left(1-\varepsilon\frac{\sigma^2}{4\rho}A_1^2\right)\right]+O(\varepsilon^2) \tag{3.7}$$

where

$$A_1 = \frac{2\rho}{\rho+\alpha}$$

Proof. See Appendix A. □

Notice that doubts about fundamentals lower equilibrium prices, $\partial P/\partial\varepsilon < 0$. This is because doubts make the agent defensively pessimistic about dividend growth. The worst-case model takes the form

$$dx = [\alpha(b-x)+\sigma\sqrt{x}h]\cdot dt+\sigma\sqrt{x}\cdot d\hat{B} \tag{3.8}$$

where $\hat{B}(t) = B(t) - \int_0^t h_s ds$ is a Brownian motion with respect to the distorted measure, $\hat{q}$. The evil agent's first-order condition implies

$$h(x) = -\varepsilon\sigma\sqrt{x}\hat{V}'(x) \tag{3.9}$$

which is negative as long as $\hat{V}'(x) > 0$. Notice also that prices decrease with σ^2. If the agent trusted his model, σ^2 would be irrevelant. However, as σ^2 increases, it becomes statistically more difficult to discriminate among models, and this gives the evil agent more latitude to distort the benchmark model.

Interestingly, we shall see that with *multiple* agents doubts will typically *increase* asset prices. With multiple agents, doubts create heterogeneous beliefs, and belief heterogeneity creates valuable trading opportunities.

3.4. **Comment on Rational Bubbles.** The approximate equilibrium price given by eq. (3.7) is the solution of a pair of linear 2nd-order ODEs. The homogeneous parts of these ODEs are the same,

$$\frac{1}{2}\sigma x\hat{V}''(x)+\alpha(b-x)\hat{V}'(x)-\rho\hat{V}(x) = 0 \tag{3.10}$$

The solution of this ODE can be expressed in terms of the confluent hypergeometric function. Froot and Obstfeld (1991) call this solution an 'intrinsic bubble'.[8] They show that intrinsic bubbles generate apparent over-reaction and excess volatility in asset prices. However, as noted in the Introduction, these kinds of Rational Bubbles do not explain the observed correlation between bubbles and trading volume. Nor do they explain why bubbles get started in the first place, or why they ultimately collapse. Hence, we rule them out, and the price given by eq. (3.7) is assumed to be a particular solution of the nonhomogeneous equation. Still, once we consider multiple agents and heterogeneous beliefs, the solution of this homogeneous equation will play a crucial role, since it will determine the resale option value.

[8]Their model assumes fundamentals follow a geometric Brownian motion. Hence, the functional form of the bubble component in their model is slightly different.

3.5. **Planner's Problem.** As in Scheinkman-Xiong, we consider an economy consisting of two groups of agents. We refer to these groups as Agent 1 and Agent 2. Beliefs are identical within groups, but differ across groups. Competition within and across groups forces prices to equal valuations. Again following Scheinkman-Xiong, we assume agents are aware that beliefs fluctuate. However, with endogenous belief fluctuations, relative beliefs become a new source of doubt. In a decentralized competitive equilibrium these doubts reflect uncertainty about future selling prices. To capture this trading uncertainty, we solve a planner's problem where the planner simultaneously distorts the beliefs of each agent, subject to separate (but equal) relative entropy penalties. The doubts of one agent therefore influence the doubts of the other.

The optimal allocation of the indivisible asset is given by the solution of the following pair of linked stopping time problems[9]

$$\begin{aligned} V(z_{1,t}, z_{2,t}, x_t, \iota) \quad = \quad & \max_{\tau} \min_{h_1,h_2} E_t \left\{ \int_t^{t+\tau} 2\rho[(\iota z_{1,s} + (1-\iota)z_{2,s})x_s] + \frac{1}{2\varepsilon}(z_{1,s}h_{1,s}^2 + z_{2,s}h_{2,s}^2)]e^{-\rho s}ds \right. \\ & \left. + e^{-\rho(t+\tau)}[V(z_{1,t+\tau}, z_{2,t+\tau}, x_{t+\tau}, 1-\iota) - \varepsilon \cdot c] \right\} \qquad (3.11) \end{aligned}$$

where $\iota = 1$ if Agent 1 holds the asset, and $\iota = 0$ if Agent 2 holds it. Notice that whenever the planner reallocates the asset he must pay a transfer cost of $\varepsilon \cdot c$. It is natural to scale the cost by ε, since this determines the magnitude of belief differences. The planner forecasts dividends using the true undistorted law of motion, but in recognition of the agents' doubts, he distorts their beliefs using the relative entropy increments, $\{h_1, h_2\}$.

$$\begin{aligned} dx &= \alpha(b-x)dt + \sigma\sqrt{x} \cdot dB \\ dz_i &= z_i h_i \cdot dB \qquad i = 1, 2 \end{aligned} \qquad (3.12)$$

The martingale state variables, (z_{1t}, z_{2t}), can be interpreted as endogenous Pareto weights. The less pessimistic agent receives a higher weight.[10] Again, one can easily see that V is homogeneous of degree one in $(z_1 + z_2)$. This allows us to reduce the number of state variables by one. In particular,

$$V(z_1, z_2, x, \iota) = (z_1 + z_2)\hat{V}(\theta, x, \iota)$$

where

$$\theta = \frac{z_1}{z_1 + z_2}$$

represents the relative optimism of Agent 1. As θ increases, Agent one becomes relatively more optimistic (or more accurately, becomes relatively less pessimistic). Using eq. (3.12), Ito's lemma implies θ follows the diffusion,

$$d\theta = \theta(1-\theta)(h_1 - h_2)\left\{-(\theta h_1 + (1-\theta)h_2)dt + dB\right\} \qquad (3.13)$$

Notice that when $h_1 > h_2$, Agent 1 becomes relatively more pessimistic and θ decreases.

[9] Miao and Wang (2011) discuss the influence of ambiguity on option values and option exercise. They show that it is important to distinguish between ambiguity about continuation payoffs and ambiguity about termination payoffs. Our linked stopping time problem features both.

[10] Anderson (2005), Borovicka (2016), and Bhandari (2013) study related problems, where heterogeneity among agents induces endogenous Pareto weights. In Anderson (2005), preferences are heterogeneous. In Borovicka (2016), beliefs are heterogeneous. As in Scheinkman and Xiong (2003), this belief heterogeneity is exogenous. In Bhandari (2013), agents' benchmark models are heterogeneous.

While Agent 1 holds the asset his valuation obeys the following HJB equation (a completely symmetric equation applies when Agent 2 holds the asset)

$$\begin{aligned}\rho\hat{V}(\theta,x,1) \quad = \quad \min_{h_1,h_2}\Big\{ & 2\rho\theta x + \hat{V}_x[\alpha(b-x)] + \frac{1}{2}\sigma^2 x\hat{V}_{xx} + \frac{1}{2\varepsilon}\left(\theta h_1^2 + (1-\theta)h_2^2\right) + \frac{1}{2}\theta^2(1-\theta)^2\hat{V}_{\theta\theta}(h_1-h_2)^2 \\ & + \theta h_1\sigma\sqrt{x}[\hat{V}_x + (1-\theta)\hat{V}_{\theta x}] + (1-\theta)h_2\sigma\sqrt{x}[\hat{V}_x - \theta V_{\theta x}]\ \Big\} \end{aligned} \tag{3.14}$$

The first-order conditions deliver the following approximate characterization of distorted beliefs

Proposition 3.2. *To an $O(\varepsilon)$ approximation, robust beliefs are given by*

$$h_1(\theta,x) \quad = \quad -\varepsilon\sigma\sqrt{x}(\hat{V}_x + (1-\theta)\hat{V}_{x\theta}) \tag{3.15}$$

$$h_2(\theta,x) \quad = \quad -\varepsilon\sigma\sqrt{x}(\hat{V}_x - \theta\hat{V}_{x\theta}) \tag{3.16}$$

Proof. See Appendix B. □

The first term in these expressions depicts the usual result that robustness produces a negative drift distortion when the marginal value of the state is positive. The new element here arises from the second terms. Since the marginal value of x increases with the optimism of the current owner, we expect $\hat{V}_{x\theta} > 0$ when Agent 1 owns, and $\hat{V}_{x\theta} < 0$ when Agent 2 owns.[11] In other words, fluctuations in relative beliefs contribute to doubts, and therefore cause larger drift distortions.

If the h_i policy functions are substituted back into the HJB equation in (3.14), we get a second-order nonlinear PDE that must be solved subject to the usual Value-Matching and Smooth-Pasting boundary conditions. Since the only thing differentiating Agents 1 and 2 are their beliefs and current ownership status, we can write these conditions as follows

$$\hat{V}(\theta^o,x) \quad = \quad \hat{V}(\theta^{no},x) - \varepsilon\cdot c/(z_1+z_2) \tag{3.17}$$

$$\hat{V}_x(\theta^o,x) \quad = \quad \hat{V}_x(\theta^{no},x) \tag{3.18}$$

$$\hat{V}_{\theta^o}(\theta^o,x) \quad = \quad \hat{V}_{\theta^o}(\theta^{no},x) \tag{3.19}$$

where θ^o denotes the relative optimism of the current owner, and $\theta^{no} = 1-\theta^o$ denotes the relative optimism of the current nonowner. The Value-Matching condition in (3.17) says that at the switching point the valuation of the nonowner must be sufficiently greater than the valuation of the current owner to compensate for the transfer cost. The Smooth-Pasting conditions in (3.19)-(3.20) require the valuations to meet smoothly at the switch point. This ensures that neither party has an incentive to wait a little longer to see what happens.

As in the single-agent case, the nonlinear HJB equation in (3.14) cannot be solved in closed-form, so we seek $O(\varepsilon)$ perturbation approximations,

$$\hat{V}(\theta,x,\iota) \approx \hat{V}^0(\theta,x,\iota) + \varepsilon\hat{V}^1(\theta,x,\iota) \qquad \iota = 1,2$$

As noted earlier, we want to rule out rational bubbles. When $\varepsilon = 0$, agents have no doubts, beliefs are homogeneous, and there is no resale option value. Therefore, when solving for $\hat{V}^0$ we only consider the particular/fundamentals solution of the second-order

[11] These expectations will be verified below.

linear ODE characterizing $\hat{V}^0$. However, when $\varepsilon > 0$, doubts emerge and beliefs fluctuate. Therefore, when solving for $\hat{V}^1$ we include the bubble term arising from the homogeneous part of the ODE. However, it must be emphasized that this bubble component represents a resale option value, not a typical rational bubble. Here its value must satisfy additional boundary conditions that are lacking in the rational bubble interpretation of Froot and Obstfeld (1991).

The homogeneous ODE determining the resale option value turns out to be the same confluent hypergeometric equation appearing in the single-agent case, eq. (3.10). To simplify the solution of this equation, and to simplify the fundamental solution, we make the following assumption

Assumption 3.3. *The parameters satisfy the following two restrictions*

$$\begin{aligned} \frac{1}{2}\sigma^2 &= \alpha\left(\frac{\alpha b}{\rho}\right) \\ \alpha &= \rho \end{aligned}$$

The first restriction implies the solution of the confluent hypergeometric equation takes a simple exponential form. The second restriction simplifies the fundamentals solution. Since ρ is a discount rate, it implies fundamentals exhibit weak mean reversion, which is empirically consistent for most markets. The first restriction just pins down the long-run mean of fundamentals. Combined with the second restriction, it implies $b = \sigma^2/2\rho$.

Given the model's symmetry, what matters are the beliefs of the owner relative to the nonowner, not the identities of the agents themselves. The proof of this, along with the remaining details, are relegated to the Appendix. Here we merely state the result.

Proposition 3.4. *Let θ^o denote the relative optimism of the current asset owner. An $O(\varepsilon)$ approximation of the owner's value function is*

$$\hat{V}(\theta^o, x) = \theta^o(b+x)\Phi + \varepsilon B(\theta^o)e^{x/b} \tag{3.20}$$

where

$$\Phi = 1 - \varepsilon\frac{\sigma^2}{4\rho} \tag{3.21}$$

$$B(\theta^o) = \frac{b\Phi}{2\varepsilon}\left[2\varepsilon e^{-A/2\varepsilon} + (1-2\theta^o)e^{-A/|1-2\theta^o|}\right] \tag{3.22}$$

and where $A = \frac{\varepsilon \cdot c_z}{b\Phi}$ and $c_z = c/(z_1+z_2)$. The nonowner's value function is $\hat{V}(1-\theta^o, x)$.

Proof. See Appendix C. □

The value function in eq. (3.20) reveals an interesting tension created by doubts about fundamentals. The coefficient $\Phi < 1$ captures the price depressing effect of pessimistic beliefs about dividend growth. This term is present even when beliefs are homogeneous. However, with heterogeneous beliefs, even though you are pessimistic on average, there is now hope that you will be able to profitably sell the asset in the future to someone who becomes less pessimistic than you. This resale option value is captured by the final exponential term. It raises prices, and as we shall see, it tends to dominate the first effect, so that on net doubts about fundamentals actually *increase* valuations and prices. Before describing these effects in more detail, a brief technical comment is warranted.

3.6. **Comment on Continuity at $\varepsilon = 0$.** When solving functional equations with perturbation approximations, we want the approximate solution to converge continuously to the (known) baseline solution as $\varepsilon \to 0$. To ensure this, the boundary conditions must hold for the baseline solution. Here this occurs due to the scaling of c by ε. As $\varepsilon \to 0$, transfer costs disappear, and $\theta = 1/2$ becomes a lower reflecting barrier. At the same time, since θ^o is $O_p(\varepsilon)$ with negative drift, θ^o effectively becomes 'pinned' at $\theta = 1/2$. Hence, $\lim_{\varepsilon\to 0} \varepsilon B(\theta) = 0$ and $\lim_{\varepsilon \to 0} \hat{V} = \frac{1}{2}(b + x)$.

3.7. **Trading Dynamics and Bubbles.** The Value-Matching and Smooth-Pasting conditions generate an option exercise frontier, $\hat{\theta}^o(x)$. The planner reallocates the asset when (θ^o, x) hits $\hat{\theta}^o(x)$ from above. Since a relative optimist would never sell, we know $\hat{\theta}^o(x) < 1/2$. In particular,

Corollary 3.5. *The planner reallocates the asset when (θ^o, x) hits the option exercise boundary*

$$\hat{\theta}^o(x) = \frac{1}{2}\left(1 - \frac{\varepsilon \cdot c_z}{\Phi x}\right) \qquad c_z = c/(z_1 + z_2) \tag{3.23}$$

from above.

Proof. See Appendix C. □

Not surprisingly, the boundary is more difficult to hit when transaction costs increase, $\partial\hat{\theta}^o/\partial c < 0$. More interestingly, notice that the boundary is *easier* to hit when fundamentals are strong. Even minor differences in beliefs are important when the stakes are large. This implies asset trading is procyclical, which is exactly what the data show.[12]

How do we know owners eventually sell? A key property of our endogenous beliefs model is that owners become defensively more pessimistic over time. If we combine the results in Proposition 3.4 with eqs. (3.15)-(3.16) we get

Corollary 3.6. *Robust beliefs are given by*

$$\begin{aligned} h_i(\theta, x) &= -\varepsilon\sigma\sqrt{x} \qquad \textit{if } i \textit{ owns} && (3.24)\\ &= 0 \qquad \textit{otherwise} && (3.25)\end{aligned}$$

Therefore, the relative optimism of the owner follows the diffusion

$$d\theta^o = -\varepsilon\sigma\sqrt{x}\theta^o(1-\theta^o)\left\{\varepsilon\sigma\sqrt{x}\theta^o \cdot dt + dB\right\} \tag{3.26}$$

Proof. First, note that the $O(\varepsilon)$ component of $\hat{V}(\theta^0, x)$ is negligible.[13] Second, note that $\hat{V}^0_x = \theta$ and $\hat{V}^0_{x\theta} = 1$ if Agent 1 owns. From (3.15)-(3.16) we get, $h_1 = -\varepsilon\sigma\sqrt{x}$ and $h_2 = 0$. When Agent 2 owns we get $\hat{V}^0_x = (1-\theta)$ and $\hat{V}^0_{x\theta} = -1$. Again from (3.15)-(3.16) we get $h_1 = 0$ and $h_2 = -\varepsilon\sigma\sqrt{x}$. Using these results in eq. (3.13) then gives eq. (3.26). □

Hence, the relative optimism of the owner drifts down stochastically, and eventually hits the option exercise boundary, $\hat{\theta}^o(x)$. If fundamentals are strong, it will hit it more quickly. Figure 4 depicts the implied trading dynamics.

[12] Although eq. (3.23) might suggest the existence of 'boundary layer' at $x = 0$, our parameterization implies that $x = 0$ is unattainable.

[13] Note, the convexity effect highlighted in the Introduction is $O(\varepsilon^2)$, whereas the volatility/detection error effect captured here is $O(\varepsilon)$.

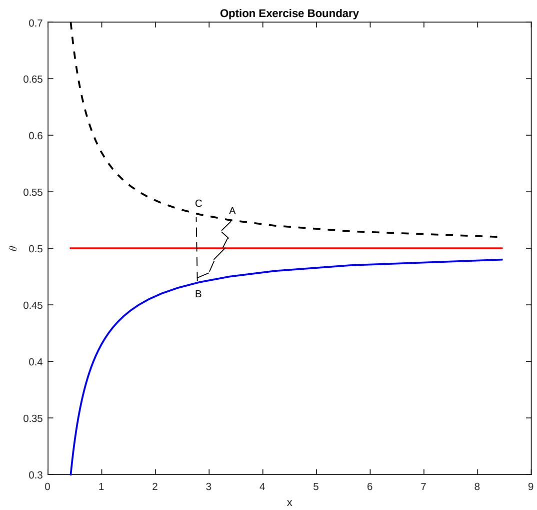

FIGURE 4. Option Exercise Boundary and Switching Dynamics

The parameter values have been set to: $\rho = .02$, $\sigma^2 = .09$, $\varepsilon = .10$, and $c = 2.0$. The implied long-run mean of fundamentals is $b = 2.25$, and the resulting 'normal' price is about 10, so that transaction costs are about 2% of price, but of course this declines during bubble episodes when prices increase. The solid blue line is the exercise boundary, $\hat{\theta}^o(x)$. The dashed black line above it plots $1 - \hat{\theta}^o(x)$. This represents the beliefs of a new owner. Starting from A, θ^o drifts down to the exercise boundary, point B. (The Figure shows θ^o declining monotonically, but this need not be the case). At B, the owner has become sufficiently pessimistic to warrant a change in ownership, and so θ^o jumps vertically to point C, and the process begins anew.

A key advantage of the Scheinkman-Xiong model is that it delivers a precise, quantitative measure of bubbles. Our model shares this advantage. From Proposition 3.4 we have

Corollary 3.7. *The bubble (resale option value) component of price is given by*[14]

$$\mathcal{B}(\theta^o, x) = \frac{1}{2} b\Phi \left[2\varepsilon e^{-A/2\varepsilon} + (1 - 2\theta^o) e^{-A/|1-2\theta^o|} \right] e^{x/b} \tag{3.27}$$

Therefore,

- **(i):** *Bubbles rise when fundamentals are strong.*
- **(ii):** *Strong fundamentals reduce the relative optimism of the current owner, which reinforces the bubble, since $\partial \mathcal{B}/\partial \theta^o < 0$.*
- **(iii):** *Bubbles and inequality are positively correlated. Bubbles rise as inequality increases, while at the same time bubbles increase inequality.*
- **(iv):** *Bubbles and trading volume are positively correlated, since $\partial \hat{\theta}^o / \partial x > 0$.*

[14]Although it is possible for $(1 - 2\theta^o) < 0$, especially for a new owner, one can verify that $\mathcal{B}(\theta^o, x)$ remains positive with arbitrarily high probability as $\varepsilon \to 0$.

Each of these predictions is consistent with the data. Although the rational/instrinsic bubble model of Froot and Obstfeld (1991) explains the first property, it cannot explain the other three. To gauge the quantitative significance of bubbles, Figure 5 plots the owner's resale option value along the exercise boundary, $\mathcal{B}(\hat{\theta}^o(x), x)$, against x. It also plots the buyer's resale option value and $\mathcal{B}(1-\hat{\theta}^o(x), x)$. For comparison, the fundamental value $b + x$ is also plotted. The parameter values are the same as in Figure 4.

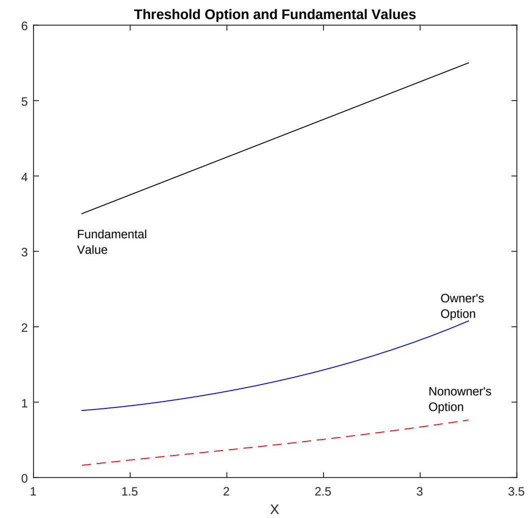

FIGURE 5. Threshold Fundamentals and Option Values

Notice that the resale option is more valuable to the owner than the nonowner, and that this value rises as fundamentals rise. Hence, when an exchange takes place, it must be the case that the owner's beliefs about fundamental value fall sufficiently below those of nonowners to offset the difference in resale option values. Figure 5 shows that the bubble component can be a significant fraction of asset values. For example, when a trade occurs at just beyond $x = 3$, the fundamental value is about 5.5, while the option/bubble component is about 2, so that the bubble component is more than 25% of valuation. Simulations presented below show that this bubble fraction can be even higher.

3.8. **Decentralization.** So far we have studied the problem of a planner who by fiat optimally allocates an indivisible asset between two agents who have doubts about the underlying dividend process. Here we address the following question - Under what conditions can this allocation be decentralized as a competitive equilibrium? An unconventional feature of the above planner's problem is that the planner *jointly* selected pessimistic drift distortions. As a result, drift distortions were allowed to feedback on each other. However, in a competitive equilibrium each agent confronts his doubts alone. This raises questions about whether the above allocation can be decentralized.[15]

[15]Our motivation here is to study competitive pricing and trading under relatively frictionless market conditions. A recent literature debates whether frictionless financial markets are desirable when agents

The key result is Corollary 3.6. This result says that an agent only distorts cash flows while he holds the asset. When he doesn't, his beliefs remain invariant. This is precisely the outcome in a competitive equilibrium. If your wealth is not exposed to fluctuations in dividends, then you have no reason to doubt them. That doesn't mean uncertainty is no longer relevant, since uncertainty may be what is preventing you from acquiring the asset. That is, the *level* of the nonowner's martingale belief distortion, z_t, continues to matter, even though it no longer changes when the asset is not held.

Following Scheinkman and Xiong (2003), we assume agents operate in relatively frictionless financial markets. As is well known, trading volume depends on the structure of financial markets. Securities that are more complex and state-contingent need to be traded less frequently. For empirical reasons, we posit a simple financial market structure, in which agents can trade an (indivisible) equity claim to the dividend-yielding asset along with a risk-free bond. In particular, we make the following assumptions

Assumption 3.8. *Assume the following,*

(1) *There is a constant, exogenous, risk-free interest rate with $r = \rho$.*
(2) *Short-sales are prohibited.*
(3) *The asset supply is fixed, and agents always have sufficient wealth to purchase it.*
(4) *There are no commodity or asset rental markets. If you want to consume, you must buy the asset.*

These are not innocuous assumptions. Relaxing any one of them would substationally complicate the analysis. Dumas, Kurshev, and Uppal (2009) incorporate risk aversion into the Scheinkman-Xiong model, under the assumption of complete markets. This allows the short-selling constraint to be relaxed. However, as in Scheinkman-Xiong, belief heterogeneity in their model is exogenous. Endogenizing the interest rate would be feasible but nontrivial.[16] As usual, it can be defended as a 'small open-economy' assumption. The constraint on short sales is of course crucial, although as noted by Scheinkman (2014) and Xiong (2013), all that really matters is that short-sales are costly.[17] As discussed by Scheinkman (2014), limited capital and endogenous asset float tend to limit both the magnitude and duration of asset market bubbles. He also emphasizes that they appeared to be important factors in several historical bubble episodes.[18]

We can now state the following result,

have distorted beliefs. See *inter alia*, Brunnermeier, Simsek, and Xiong (2014), Gilboa, Samuelson, and Schmeidler (2014), and Blume, Cogley, Easley, Sargent, and Tsyrennikov (2014).

[16]Dumas (1989) studies an economy consisting of two agents with different degrees of relative risk aversion, who must share the returns of a single dividend-yielding asset. He shows that endogenous interest rate fluctuations are an important feature of the equilibrium.

[17]Simsek (2013) notes that relaxing the first assumption may endogenously relax the second. In his model, borrowing and lending must be collateralized by the value of the risky asset. Endogenous short-sale constraints emerge when belief disagreements concern upside states, since these are the default states for a short-seller.

[18]Nutz and Scheinkman (2018) develop a model with (asymmetric) quadratic asset holding costs, in which asset float influences bubble dynamics.

Proposition 3.9. *Given Assumption 3.8, the planner's allocation sequence described in Corollary 3.5 is an $O(\varepsilon)$ approximation of a competitive equilibrium. The market price is*

$$P(z_1, z_2, x) = (z_1 + z_2)\hat{V}(\theta^o, x) \tag{3.28}$$

where $\hat{V}(\theta^o, x)$ is given by Proposition 3.4, θ^o is given by eq. (3.26), z_i are given by eq. (3.12), and h_i are as given in Corollary 3.6.

Proof. See Appendix D. □

Observe that the beliefs of nonowners continue to exert an influence on the current market price. This is due to the fact that owners have the option to sell the asset back to them at some point in the future.

4. Simulations

Corollary 3.5 shows that trading dynamics are governed by the hitting time of a 2-dimensional diffusion process, (x, θ^o), where x follows the exogenous CIR process in (3.1) and θ^o follows the x-dependent logistic process in (3.26). The boundaries are given by $\hat{\theta}^o$ and $1 - \hat{\theta}^o$, as described in Corollary 3.5. In principle, one could derive analytical approximations of the resulting trading dynamics, e.g., mean times between trades and their autocorrelations. Unfortunately, this is likely to be a challenging exercise, since the hitting time densities of even relatively simple univariate processes with relatively simple boundaries, like Ornstein-Uhlenbeck processes, tend to be quite complicated. Moreover, matching these properties to observed data would be difficult in any case, given that the model only contains two classes of traders.

Hence, in this section we merely illustrate the qualitative properties of the price and trading dynamics. Our interest is to see how important the bubble component is, how it covaries with the state of the economy, and how it correlates with trading. We continue to impose the parameter restrictions described in Assumption 3.3. The remaining parameter values are given in Table 1

TABLE 1
Benchmark Parameter Values

σ^2	r	ε	c
.09	.02	.10	2.0

The implied long-run mean of x is $b = \sigma^2/2r = 2.25$, and the implied long-run mean Rational Expectations price is $p = b + E(x) = 2b = 4.5$. Hence, trading costs would be ($\varepsilon \cdot c/4.5 = .044$) 4.4% of the Rational Expectations price. Of course, in a Rational Expectations equilibrium there would be no trading, and so the importance of trading costs depends on whether robust prices are higher or lower than in a Rational Expectations equilibrium.

Figures 6-13 contain the results.

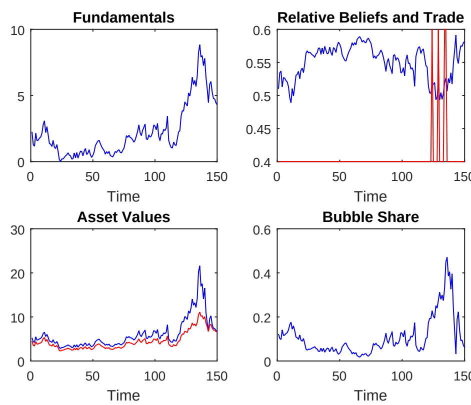

FIGURE 6. Simulation 1

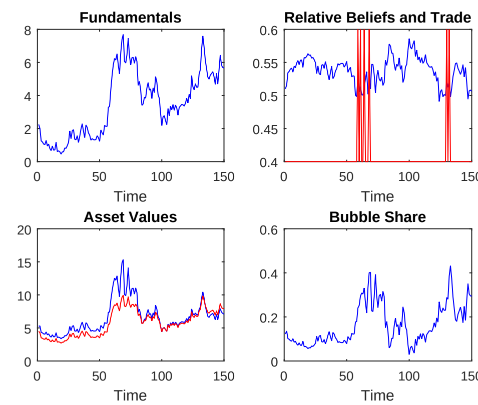

FIGURE 7. Simulation 2

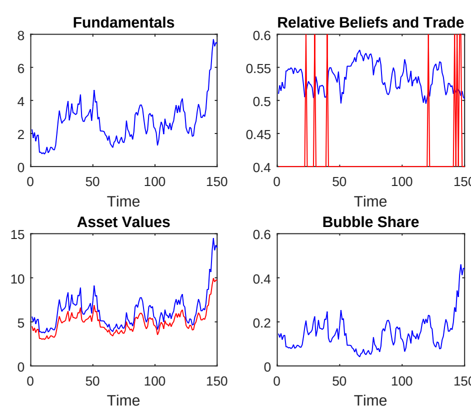

FIGURE 8. Simulation 3

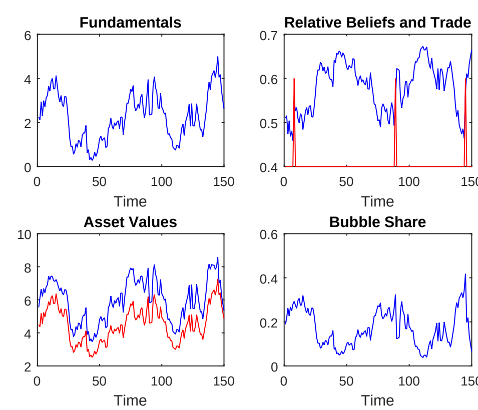

FIGURE 9. Simulation 4

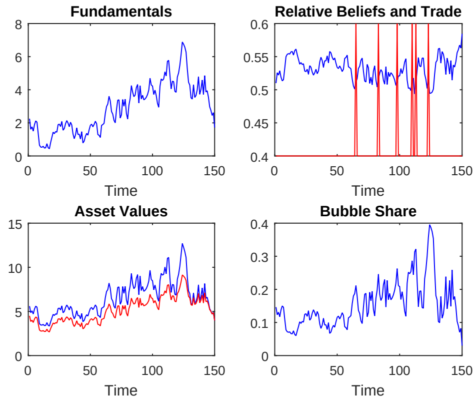

FIGURE 10. Simulation 5

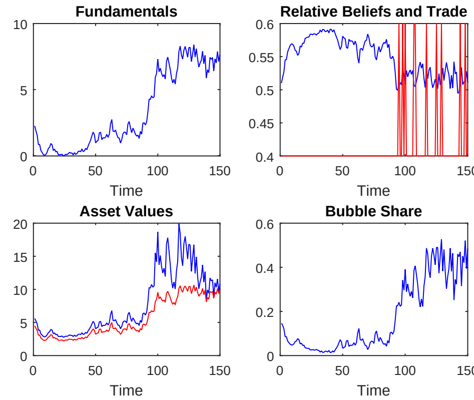

FIGURE 11. Simulation 6

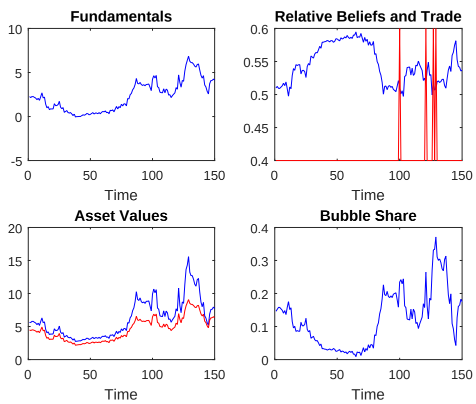

FIGURE 12. Simulation 7

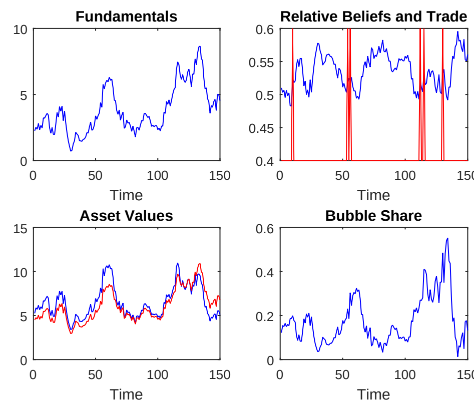

FIGURE 13. Simulation 8

The upper left panel displays the realized path of fundamentals, with a sample size of $T = 150$ (annual) observations, roughly equal to the number of observations in Shiller's U.S. data. The upper right panel plots the evolution of relative beliefs, θ^o. It is initialized so that the initial owner is mildly (relatively) optimistic, $\theta^o_0 = .51$. The vertical red lines correspond to asset trades. The bottom left panel plots the market price (in blue), along with the Rational Expectations price (in red). Finally, the bottom right panel plots the bubble component, expressed as a percentage of the market price.

The main points to notice in these Figures are the following: (1) Robust prices are typically, though not always, above their Rational Expectations counterparts. By itself, robustness depresses prices. However, it also creates trading opportunities, which increases them. The plots suggest that on net the resell option component dominates; (2) When fundamentals rise, current asset owners become relatively pessimistic, and θ^o declines. Eventually, θ^o hits the exercise boundary, $\hat{\theta}^o$, and a burst of trading occurs. Hence, as in the data, *trading is procyclical and positively autocorrelated*;[19] (3) The bubble component of prices is small on average (less than 20% of price), but becomes quite significant during bubble episodes, typically rising to about 40% of price. As a result, market prices appear 'overvalued' relative to the Rational Expectations benchmark; (4) Part of the reason trading becomes frenetic during bubbles is that prices roughly double on average, so that trading costs expressed as a percentage of price fall to only about 2% of value; (5) Bubble episodes occur irregularly, at random intervals. Without detailed knowledge of θ^o, they would be difficult to predict using observed data. The best predictor would be the current state of fundamentals. If fundamentals rise above $x = 5$, a significant bubble is likely to be present.

As emphasized by Scheinkman (2014) and Xiong (2013), each of these features are an important part of observed bubble episodes. The fact that bubbles are correlated with trading distinguishes our model of speculative bubbles from more traditional Rational Bubble models. The key new feature of our model relative to other models of speculative bubbles is that here belief heterogeneity is *endogenous.* It depends on the distribution of wealth. Of course, just because belief heterogeneity is endogenous, doesn't mean it is empirically plausible. However, another advantage of our approach is that the underlying doubts that create belief heterogeneity can be empirically assessed using Hansen and Sargent's notion of detection error probabilities. We turn to that task now.

5. Detection Error Probabilities

The key mechanism in our model is wealth dependent (relative) pessimism. Pessimism is operationalized using a max-min objective function, in which agents optimize against an endogenous worst-case scenario. A potential criticism of this approach is to ask why agents should expect the worst. Following Hansen and Sargent (2008), we answer this critique by constraining the set of models the agent considers. Agents only consider models that could have plausibly generated the observed data.

Plausibility is quantified using 'detection error probabilities'. Agents are viewed as statisticians, who attempt to discriminate among models using likelihood ratio statistics. When likelihood ratio statistics are large, detection error probabilities are small, and models are easy to distinguish. Detection error probabilities will be small when models are very different, or when there is a lot of data. Classical inference is based on adopting a null hypothesis, and fixing the probability of falsely rejecting this hypothesis. Detection error probabilities treat the null and alternative symmetrically, and average between Type

[19] Statman, Thorley, and Vorkink (2006) examine monthly NYSE/AMEX data for the period 1962-2002. They show that turnover rates are positively autocorrelated, and positively correlated with lagged returns, for both individual stocks and the market as a whole.

I and Type II errors. In particular,

$$DEP = \frac{1}{2}\text{Prob}(H_A|H_0) + \frac{1}{2}\text{Prob}(H_0|H_A)$$

Hence, a DEP is analogous to a p-value. Our results would therefore be implausible if the DEP is too small. Small DEP's imply agents are hedging against models that could easily be rejected using observed data.

Anderson, Hansen, and Sargent (2003) provide an analytical upper bound on DEP's. Unfortunately, this upper bound might not be very tight, so we instead follow the monte carlo simulation strategy outlined in Chapter 9 of Hansen and Sargent (2008). The null approximating model can be discretized as follows

$$x_t^n = rb + (1-r)x_{t-1}^n + \sigma\sqrt{x_{t-1}^n}B_t$$

where $B_t \sim N(0,1)$. The alternative worst-case model exhibits a negative drift distortion

$$x_t^a = rb + [(1-r) - \varepsilon\sigma^2\Phi]x_{t-1}^a + \sigma\sqrt{x_{t-1}^a}B_t$$

Notice that the worst-case model features less persistence. Let $\ell = \log\frac{L^n}{L^a}$ be the log-likelihood ratio in favor of the the null model. When in fact the null model generates the data we have

$$\ell|n = \sum_{t=0}^{T-1}\left[\frac{1}{2}(w_{t+1}^n)^2 - w_{t+1}^n B_{t+1}^n\right]$$

where $w_t^n = -\varepsilon\sigma\Phi\sqrt{x_{t-1}^n}$ and $B_{t+1}^n = \sqrt{x_t^n}B_{t+1}$. A detection error occurs if the realization of this statistic is negative. If instead the worst-case alternative model generates the data we have

$$\ell|a = -\sum_{t=0}^{T-1}\left[\frac{1}{2}(w_{t+1}^a)^2 + w_{t+1}^a B_{t+1}^a\right]$$

where $w_t^a = -\varepsilon\sigma\Phi\sqrt{x_{t-1}^a}$ and $B_{t+1}^a = \sqrt{x_t^a}B_{t+1}$. Now a detection error occurs if this is positive. We numerically approximate DEP's by averaging across $N = 1000$ replications. For each replication we draw T shocks, and use them to construct a sample path of dividends under the two models. DEPs are then computed by counting the proportion of detection errors in the N replications, and then averaging across the two types of DEPs.

Figure 14 displays the results

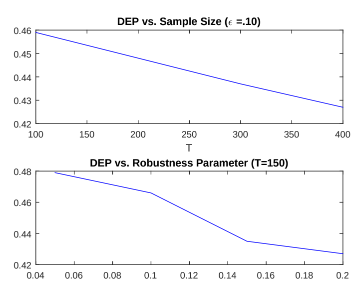

Figure 14. Detection Error Probabilities

The top panel plots DEP's as a function of sample size, T, assuming the robustness parameter equals its benchmark value in Table 1 ($\varepsilon = .10$). The bottom panel then plots DEP's as a function of the robustness parameter, assuming the sample size equals its benchmark value ($T = 150$).

As expected, DEP's decrease with both the sample size and the discrepancy between models, as parameterized by ε. However, notice that in all cases the DEP's are in excess of 0.40. Even with hundreds of years of data it would be quite difficult to discriminate between the benchmark and worst-case models. This is simply because the two models are not that different. In the benchmark model the AR parameter is 0.98, whereas in the worst-case model it is 0.97. It is very difficult to statistically distinguish these two models. After all, the approximate standard deviation of the shocks is $\sigma\sqrt{x} \approx .3\sqrt{2.25} = .45$.

What is more surprising is that this relatively minor discrepancy can generate significant belief differences and significant asset price bubbles. As fundamentals rise, current asset owners begin to think the good times will be somewhat less persistent than nonowners believe. (Remember, nonowners do not revise their persistence estimates). With infinite horizons and low interest rates, even minor differences in beliefs about persistence translate into significant valuation differences.[20]

6. Related Literature

Scheinkman (2014) argues that a crucial piece of evidence in favor of speculative bubbles (as opposed to Rational Bubbles) is that observed bubble episodes are correlated with large trading volumes. Our model of wealth dependent pessimism endogenously generates such a correlation. It also provides a fresh perspective on some other properties of asset trading. One of the most widely studied puzzles is the so-called 'disposition effect'. Shefrin and Statman (1985) showed that traders tend to sell 'winners' too soon, and hold 'losers' for too long. They argued that Prospect Theory provides a potential explanation of this puzzle. Since then, other behavioral explanations have been proposed. For example, Barberis and Xiong (2012) show that if traders derive utility from realized gains, then even if utility functions are linear they will appear to display a disposition effect.[21] Interestingly, wealth dependent pessimism also generates an apparent disposition effect. In our model, current asset owners sell before the peak, and appear to systematically forego capital gains. They do this because they become concerned about downside risk. Is this an irrational/behavioral explanation? Only if you believe agents have full confidence in their models of fundamentals. Our model also turns the conventional behavioral interpretation of procyclical trading on its head. According to Statman, Thorley, and Vorkink (2006), the reason trading is procyclical is that agents become 'overconfident' during a period of positive excess returns. This leads them to trade more. In contrast, the confidence of our agents *decreases* during booms, simply because there is now more to lose. With doubts, it is rational to be 'underconfident'.

[20]Hansen and Sargent (2010) show that a similar mechanism (albeit in a representative agent model) can explain cyclical variation in the price of risk.

[21]The structure of their model is quite similar to ours. Agents have linear utility, and must solve an optimal stopping problem subject to a fixed transactions cost. The crucial difference is that they study the problem of a single agent who has Rational Expectations. In contrast, endogenous heterogeneous beliefs are central to our story.

Hansen and Sargent (2010) note that endogenous pessimism can be interpreted as a form of state-dependent risk aversion. Grossman and Laroque (1990), Chetty and Szeidl (2007), and Stokey (2009) study models that are similar to the one in this paper, and show that fixed costs of adjusting durable assets or durable consumption goods also generate state-dependent risk aversion. However, in their models agents become *less* risk averse as the value of durable assets rises relative to total wealth. Here increasing pessimism makes asset owners appear to become *more* risk averse as asset values rise.

Finally, portfolio rebalancing effects can also generate trading dynamics and apparent state-dependent risk aversion. However, most existing models of portfolio rebalancing consider representative agent economies, where the effects show up in prices rather than in trading (e.g., Cochrane, Longstaff, and Santa-Clara (2008)).[22] Although we interpret our model as reflecting endogenous wealth dependent belief dynamics, a recurring theme in the robust control literature is that one can alternatively interpret robustness in terms of 'risk-sensitivity'. From this perspective, the dynamics in our model could be viewed in terms of portfolio rebalancing. We prefer the robustness interpretation because it offers a more interesting interpretation of the crucial parametere, ε. Rather than view it as an exogenous, time invariant, aspect of preferences, we view it as a potentially time varying feature of the environment, which can be disciplined using detection error probabilities.

7. Conclusion

When agents have doubts about fundamentals, these doubts typically depend on an agent's wealth. If agents have different wealth, they will have different beliefs about fundamentals. Different beliefs about fundamentals create resale option values that can be interpreted as asset price bubbles. Our paper has shown that empirically plausible doubts about fundamentals produce quantitatively significant bubbles that coincide with large trading volumes and growing wealth inequality.

Although suggestive, our model is too stylized to take to the data in a serious way. Doing so would likely require the following extensions. First, to explain observed data on trading, bubbles, and wealth inequality it would be desirable to have a model with many agents and many assets. Second, financial assets are better thought of as being divisible and diversifiable. A perfectly pooled equilibrium with identical wealth would eliminate belief heterogeneity and bubbles. However, we conjecture that such an equilibrium would be fragile. Any initial wealth differences would create belief and portfolio hetergeneity. Third, it would be desirable to endogenize the interest rate and short-selling costs, perhaps by requiring borrowing to be collateralized by the value of the underlying asset (Simsek (2013)). Finally, the bubbles generated by our model often last for a decade or more. Observed bubbles tend to have shorter durations. Endogenous asset float would likely remedy this discrepancy (Scheinkman (2014), Nutz and Scheinkman (2018)).

[22]Dumas (1989), Wang (1994), and Dumas, Kurshev, and Uppal (2009) are notable exceptions. In Dumas (1989) trading is driven by heterogenous preferences. In Wang (1994) trading is driven by heterogeneous information. In Dumas, Kurshev, and Uppal (2009) trading is driven by (exogenous) heterogeneous beliefs.

Appendix A. Proof of Proposition 3.1

Expressed in terms of $\hat{V}(x)$, the agent's HJB equation becomes

$$\rho\hat{V}(x) = \min_h \left\{ 2\rho x + \frac{1}{2\varepsilon}h^2 + \alpha(b-x)\hat{V}'(x) + \frac{1}{2}\sigma^2 x\hat{V}''(x) + \sigma\sqrt{x}h\hat{V}'(x) \right\} \tag{A.29}$$

The evil agent's policy function is

$$h(x) = -\varepsilon\sigma\sqrt{x}\hat{V}'(x) \tag{A.30}$$

Substituting this back into the HJB equation we get

$$\rho\hat{V}(x) = 2\rho x + \alpha(b-x)\hat{V}'(x) + \frac{1}{2}\sigma^2 x\hat{V}''(x) - \frac{1}{2}\varepsilon\sigma^2 x(\hat{V}'(x))^2 \tag{A.31}$$

Except for the last term multiplying ε, this is linear, with a particular (no bubbles) solution of the form $\hat{V}(x) = A_0 + A_1 x$. This opens the door to a perturbation approximation around ε. To obtain this, we posit

$$\hat{V}(x) \approx V^0(x) + \varepsilon V^1(x)$$

Our goal is to solve for $V^0(x)$ and $V^1(x)$. Matching terms of equal order and dropping terms of order $O(\varepsilon^2)$ gives

$$\begin{aligned} \varepsilon^0: \quad \rho V^0 &= 2\rho x + \alpha(b-x)V_x^0 + \frac{1}{2}\sigma^2 x V_{xx}^0 \\ \varepsilon^1: \quad \rho V^1 &= \alpha(b-x)V_x^1 + \frac{1}{2}\sigma^2 x V_{xx}^1 - \frac{1}{2}\sigma^2 x(V_x^0)^2 \end{aligned}$$

Note this system is recursive. We can first solve the ε^0 equation for V^0, and then substitute this into the ε^1 equation. One can readily verify that the solution for V^0 is just

$$V^0(x) = A_0 + A_1 x \qquad A_1 = \frac{2\rho}{\rho+\alpha} \qquad A_0 = \frac{\alpha b}{\rho}A_1$$

Note that this implies $V_x^0 = A_1$. One can again readily verify that the solution for $V^1(x)$ is just

$$V^1(x) = B_0 + B_1 x \qquad B_1 = -\frac{\sigma^2}{2(\alpha+\rho)}A_1^2 \qquad B_0 = \frac{\alpha b}{\rho}B_1$$

Combining these results we get

$$V(x) = \frac{2}{\rho+\alpha}\left[\alpha b + \rho x - \varepsilon\frac{\sigma^2}{4\rho}A_1^2(\alpha b + \rho x)\right]$$

which is the same as equation (3.7) in Proposition 3.1. □

Appendix B. Proof of Proposition 3.2

The first-order conditions for the HJB equation in (3.15) are

$$\begin{aligned} \theta h_1 + \varepsilon[\theta^2(1-\theta)^2\hat{V}_{\theta\theta}(h_1-h_2) + \theta\sigma\sqrt{x}(\hat{V}_x + (1-\theta)\hat{V}_{x\theta})] &= 0 \qquad (B.32) \\ (1-\theta)h_2 - \varepsilon\theta^2(1-\theta)^2\hat{V}_{\theta\theta}(h_1-h_2) + \varepsilon(1-\theta)\sigma\sqrt{x}(\hat{V}_x - \theta\hat{V}_{x\theta}) &= 0 \qquad (B.33) \end{aligned}$$

Solving (B.33) for h_2 gives

$$h_2 = \varepsilon\left[\frac{\theta^2(1-\theta)^2\hat{V}_{\theta\theta}h_1 - (1-\theta)\sigma\sqrt{x}(\hat{V}_x - \theta\hat{V}_{x\theta})}{(1-\theta) + \varepsilon\theta^2(1-\theta)^2\hat{V}_{\theta\theta}}\right] \tag{B.34}$$

Substituting this into (B.32) gives

$$[\theta + \varepsilon\theta^2(1-\theta)^2\hat{V}_{\theta\theta}]h_1 - \varepsilon^2\theta^2(1-\theta)^2\hat{V}_{\theta\theta}\left[\frac{\theta^2(1-\theta)^2\hat{V}_{\theta\theta}h_1 - (1-\theta)\sigma\sqrt{x}(\hat{V}_x - \theta\hat{V}_{x\theta})}{(1-\theta) + \varepsilon\theta^2(1-\theta)^2\hat{V}_{\theta\theta}}\right] + \varepsilon\theta\sigma\sqrt{x}(\hat{V}_x + (1-\theta)\hat{V}_{x\theta}) = 0$$

The middle term is clearly $O(\varepsilon^2)$ and can be dropped. Solving for h_1 then gives

$$h_1 \approx \frac{-\varepsilon\theta\sigma\sqrt{x}(\hat{V}_x + (1-\theta)\hat{V}_{x\theta})}{\theta[1 + \varepsilon\theta(1-\theta)^2\hat{V}_{\theta\theta}]} \approx -\varepsilon\sigma\sqrt{x}(\hat{V}_x + (1-\theta)\hat{V}_{x\theta})$$

This verifies eq. (3.15). Substituting this back into eq. (B.34) and again dropping $O(\varepsilon^2)$ terms gives eq. (3.16). $\square$

Appendix C. Proof of Proposition 3.4

Substituting the $h_i(\theta, x)$ policy functions in (3.16)-(3.17) into the HJB eqs. in (3.14) and dropping $O(\varepsilon^2)$ terms gives

$$\rho\hat{V}(\theta,x,1) = 2\rho\theta x + \hat{V}_x[\rho(b-x)] + \frac{1}{2}\sigma^2 x\hat{V}_{xx} - \frac{1}{2}\varepsilon\sigma^2 x\hat{V}_x^2 - \frac{1}{2}\varepsilon\theta(1-\theta)\sigma^2 x\hat{V}_{x\theta}^2 \quad (C.35)$$

$$\rho\hat{V}(\theta,x,2) = 2\rho(1-\theta)x + \hat{V}_x[\rho(b-x)] + \frac{1}{2}\sigma^2 x\hat{V}_{xx} - \frac{1}{2}\varepsilon\sigma^2 x\hat{V}_x^2 - \frac{1}{2}\varepsilon\theta(1-\theta)\sigma^2 x\hat{V}_{x\theta}^2 \quad (C.36)$$

where we've imposed the parameter restriction $\alpha = \rho$. Let's focus first on the $\hat{V}(\theta, x, 1)$ equation. It will become clear that $\hat{V}(\theta, x, 2)$ can then be obtained by replacing θ with $1-\theta$ in $\hat{V}(\theta, x, 1)$. The last two terms make this a nonlinear PDE. Hence, we look for a solution of the form

$$\hat{V}(\theta,x,1) = V^0(\theta,x) + \varepsilon V^1(\theta,x)$$

Matching terms of equal order and dropping terms of order $O(\varepsilon^2)$ gives

$$\varepsilon^0: \quad \rho V^0 = 2\rho\theta x + \rho(b-x)V_x^0 + \frac{1}{2}\sigma^2 x V_{xx}^0$$

$$\varepsilon^1: \quad \rho V^1 = \rho(b-x)V_x^1 + \frac{1}{2}\sigma^2 x V_{xx}^1 - \frac{1}{2}\sigma^2 x (V_x^0)^2 - \frac{1}{2}\theta(1-\theta)\sigma^2 x (V_{x\theta}^0)^2$$

As discussed in the main text, we look for a fundamentals (no bubbles) solution to V^0. Hence, conjecture a solution of the form $V^0 = A\theta x + B\theta$. This delivers the following two fixed point conditions

$$\rho A = 2\rho - \rho A$$
$$\rho B = \rho b A$$

These imply $A = 1$ and $B = b$. Hence, $V^0(\theta, x) = \theta(b + x)$. Note this implies $V_x^0 = \theta$ and $V_{x\theta}^0 = 1$. Substituting these results into the V^1 equation gives

$$\rho V^1 = \rho(b-x)V_x^1 + \frac{1}{2}\sigma^2 x V_{xx}^1 - \frac{1}{2}\sigma^2 x\theta$$

The general solution of this linear nonhomogeneous ODE is the sum of a particular solution and a solution of the homogeneous equation, $V^1 = V_p^1 + V_h^1$. The solution of the homogeneous equation determines the resale option value. To obtain a particular solution, let's again posit a solution of the form $V_p^1 = C\theta x + D\theta$. Matching coefficients delivers the fixed point conditions

$$\rho C = -\rho C - \frac{1}{2}\sigma^2$$
$$\rho D = \rho b C$$

These imply $C = -\frac{\sigma^2}{4\rho}$ and $D = -\frac{b\sigma^2}{4\rho}$. Therefore, $V_p^1 = -\theta\frac{\sigma^2}{4\rho}(b + x)$. The homogeneous solution solves the hypergeometric equation

$$\rho V^1 = \rho(b-x)V_x^1 + \frac{1}{2}\sigma^2 x V_{xx}^1$$

Let's guess a solution of the form $V_h^1 = B(\theta)e^{\lambda x}$. Note that since we are solving a PDE, the constant of integration will in general be a function of θ. We denote this function, $B(\theta)$. Substituting in this guess gives

$$B(\theta)\left[\frac{1}{2}\sigma^2 x\lambda^2 + \rho\lambda(b-x) - \rho\right]e^{\lambda x} = 0$$

For this solution to work the expression in brackets must be identically zero. This requires

$$\rho\lambda b - \rho = 0$$
$$\frac{1}{2}\sigma^2\lambda^2 - \rho\lambda = 0$$

The first equation implies $\lambda = 1/b$. The second equation then implies $b = \sigma^2/2\rho$. This is the restriction stated in Assumption 3.3 in the main text. Therefore, given Assumption 3.3, the solution to the homogeneous equation is $V_h^1 = B(\theta)e^{x/b}$.Combining our results gives the following first-order perturbation approximation,

$$V(\theta, x, 1) = \theta(b+x)\Phi + \varepsilon B(\theta)e^{x/b} + O(\varepsilon^2) \qquad \Phi = 1 - \varepsilon\frac{\sigma^2}{4\rho}$$

which is the expression given in Propostion 3.4. By following the same steps for $V(\theta, x, 2)$, one can easily verify that $V(\theta, x, 2) = V(1-\theta, x, 1)$.

To determine $B(\theta)$, we need to impose the Value-Matching and Smooth-Pasting boundary condtions. These will also determine the exercise boundary, $\hat{\theta}^o(x)$. Substituting the solutions for V into the Value-Matching and Smooth-Pasting conditions in eqs. (3.17)-(3.19) gives

$$\begin{aligned} (1-2\theta^o)(b+x)\Phi + \varepsilon e^{x/b}[B(1-\theta^o) - B(\theta^o)] &= \varepsilon \cdot c_z && \text{(C.37)} \\ (1-2\theta^o)\Phi + \varepsilon(1/b)e^{x/b}[B(1-\theta^o) - B(\theta^o)] &= 0 && \text{(C.38)} \\ -2(b+x)\Phi + \varepsilon e^{x/b}[B(1-\theta^o) - B(\theta^o)]' &= 0 && \text{(C.39)} \end{aligned}$$

Substituting (C.38) into (C.37) and solving for θ^o gives the exercise boundary

$$\hat{\theta}^o(x) = \frac{1}{2}\left(1 - \frac{\varepsilon \cdot c_z}{\Phi x}\right)$$

Note that $\hat{\theta}^o(x) \to 1/2$ as $\varepsilon \to 0$. Inverting $\hat{\theta}^o(x)$ gives $x = \frac{bA}{1-2\theta^o}$, where $A = \frac{\varepsilon \cdot c_z}{b\Phi}$. Substituting this into (C.38) and solving for $B(\theta^o)$ gives

$$B(\theta^o) = B(1-\theta^o) + \frac{b\Phi}{\varepsilon}(1-2\theta^o)e^{-A/(1-2\theta^o)}$$

Imposing continuity at $\theta^o = 1/2$ implies

$$\begin{aligned} B(\theta^o) &= B(1/2) + \frac{1}{2}\frac{b\Phi}{\varepsilon}(1-2\theta^o)e^{-A/(1-2\theta^o)} && \theta^o < 1/2 \\ B(\theta^o) &= B(1/2) - \frac{1}{2}\frac{b\Phi}{\varepsilon}(2\theta^o-1)e^{-A/(2\theta^o-1)} && \theta^o \geq 1/2 \qquad \text{(C.40)} \end{aligned}$$

To determine $B(1/2)$ we need another boundary condition. If θ^o were $O_p(1)$, this boundary condition would be given by $B(1) = 0$, which says that the option to resell becomess worthless when the owner's relative optimism goes to one. However, since θ^o is only $O_p(\varepsilon)$, the relevant boundary condition becomes $B(1/2+\varepsilon) = 0$. To an $O(\varepsilon)$ approximation, the value of the resell option goes to zero as $\theta^o \to 1/2+\varepsilon$. Evaluating (C.39) at $\theta^o = 1/2+\varepsilon$ and solving for $B(1/2)$ gives

$$B(1/2) = b\Phi e^{-A/2\varepsilon}$$

The verifies the expression for $B(\theta^o)$ given in Proposition 3.4. Finally, one can readily verify that this expression also satisfies the Smooth-Pasting condition for θ^o given in (C.38). □

Appendix D. Proof of Proposition 3.9

We must show that the planner's allocation rule produces the same sequence of trades as in a competitive equilibrium, where agents take the price and the beliefs of other traders as given. The key result from the planner's problem that delivers this equivalence is that drift distortions only occur while an agent holds the asset (Corollary 3.6). Hence, although the *level* of other traders' beliefs matter, since they influence the potential future selling price, current asset owners do not need to forecast the evolution of other traders' beliefs.

Without loss of generality, suppose agent 1 is the current owner. Once again, we can exploit homogeneity to write $V(z_1, z_2, x, 1) = (z_1 + z_2)\hat{V}(\theta, x, 1)$. Notice that the other agent's beliefs are a state variable in the current owner's value function. Ceteris paribus, if other agents are more optimistic (less pessimistic), that increases the value of the asset to you, since it increases the potential selling price. While holding the asset, agent 1's HJB equation is

$$\rho\hat{V}(\theta,x,1)=\min_{h_1}\left\{2\rho\theta x+\hat{V}_x[\alpha(b-x)]+\frac{1}{2}\sigma^2x\hat{V}_{xx}+\frac{1}{2\varepsilon}\theta h_1^2+\frac{1}{2}\theta^2(1-\theta)^2\hat{V}_{\theta\theta}h_1^2+\theta h_1\sigma\sqrt{x}[\hat{V}_x+(1-\theta)\hat{V}_{\theta x}]\right\} \tag{D.41}$$

where the evolution of θ is given by eq. (3.13) with $h_2=0$. Note, this is the same as the planner's HJB equation in (3.14), except now, in a competitive equilbrium, the agent is only distorting his own beliefs. Solving the first-order condition for h_1 and dropping $O(\varepsilon^2)$ terms gives the same distortion as in the planner's problem

$$h_1(\theta,x)=-\varepsilon\sigma\sqrt{x}(\hat{V}_x+(1-\theta)\hat{V}_{x\theta})$$

Repeating the same steps for the nonowner (Agent 2) yields

$$h_2(\theta,x)=-\varepsilon\sigma\sqrt{x}(\hat{V}_x-\theta\hat{V}_{x\theta})$$

If h_1 is substituted back into (D.41) we get

$$\rho\hat{V}(\theta,x,1)=2\rho\theta x+\hat{V}_x[\alpha(b-x)]+\frac{1}{2}\sigma^2x\hat{V}_{xx}-\frac{1}{2}\varepsilon\theta\sigma^2x\left[V_x^2+2(1-\theta)V_xV_{x\theta}+(1-\theta)^2V_{x\theta}\right]$$

Again we look for a first-order perturbation approximation, $V=V^0+\varepsilon V^1$. Under the posited equilibrium, $V_x^0=\theta$ and $V_{x\theta}^0=1$. Note that if these are substituted into the HJB equation, the last term in brackets just equals one, and so the nonhomogeneous term in the V^1 equation reduces to $-\theta x\sigma^2$, which produces the same particular solution as in the planner's problem. Clearly, the homogeneous equation is the same as before. Hence, with the same continuation value function as before, and the same boundary conditions, we get the exact same exercise boundary and trading strategies as before. □

References

Anderson, E. W. (2005): "The Dynamics of Risk-Sensitive Allocations," *Journal of Economic Theory*, 125, 93–150.

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003): "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection," *Journal of the European Economic Association*, 1, 68–123.

Barberis, N., and W. Xiong (2012): "Realization Utility," *Journal of Financial Economics*, 104, 251–71.

Bhandari, A. (2013): "Doubts, Asymmetries, and Insurance," New York University.

Blume, L. E., T. Cogley, D. A. Easley, T. J. Sargent, and V. Tsyrennikov (2014): "The Case for Incomplete Markets," Cornell University.

Boehme, R. D., B. R. Danielsen, and S. M. Sorescu (2006): "Short-Sale Constraints, Differences of Opinion, and Overvaluation," *Journal of Financial and Quantitative Analysis*, 41, 455–87.

Borovicka, J. (2016): "Survival and Long-Run Dynamics with Heterogeneous Beliefs Under Recursive Preferences," New York University.

Brunnermeier, M. K., A. Simsek, and W. Xiong (2014): "A Welfare Criterion for Models with Distorted Beliefs," *Quarterly Journal of Economics*, 129, 1753–97.

Chetty, R., and A. Szeidl (2007): "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics*, 122, 831–77.

Cochrane, J. H., F. A. Longstaff, and P. Santa-Clara (2008): "Two Trees," *Review of Financial Studies*, 21, 347–85.

Diether, K. B., C. J. Malloy, and A. Scherbina (2002): "Differences of Opinion and the Cross Section of Stock Returns," *Journal of Finance*, 57, 2113–41.

Dumas, B. (1989): "Two-Person Dynamic Equilibria in the Capital Market," *Review of Financial Studies*, 2, 157–88.

Dumas, B., A. Kurshev, and R. Uppal (2009): "Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility," *Journal of Finance*, 65, 579–629.

Froot, K. A., and M. Obstfeld (1991): "Intrinsic Bubbles: The Case of Stock Prices," *American Economic Review*, 81, 1189–1214.

Garber, P. (2000): *Famous First Bubbles: The Fundamentals of Early Manias*. MIT Press.

Gilboa, I., L. Samuelson, and D. Schmeidler (2014): "No-Betting Pareto Dominance," *Econometrica*, 82, 1405–42.

Grossman, S. J., and G. Laroque (1990): "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58, 25–51.

Hansen, L. P., and T. J. Sargent (2008): *Robustness*. Princeton University Press.

——— (2010): "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics*, 1, 129–62.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2006): "Robust Control, Min-Max Expected Utility, and Model Misspecification," *Journal of Economic Theory*, 128, 45–90.

Harrison, J. M., and D. M. Kreps (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323–36.

Hong, H. J., and J. A. Scheinkman (2008): "Advisors and Asset Prices: A Model of the Origin of Bubbles," *Journal of Financial Economics*, 89, 268–87.

Miao, J., and N. Wang (2011): "Risk, Uncertainty and Option Exercise," *Journal of Economic Dynamics and Control*, 35, 442–61.

Miller, E. M. (1977): "Risk, Uncertainty and Divergence of Opinion," *Journal of Finance*, 32, 1151–68.

Morris, S. (1996): "Speculative Investor Behavior and Learning," *Quarterly Journal of Economics*, 111, 1111–33.

Nutz, M., and J. A. Scheinkman (2018): "Supply and Shorting in Speculative Markets," working paper, Columbia University.

Saez, E., and G. Zucman (2016): "Wealth Inequality in the United States Since 1913: Evidence from Capitalized Income Tax Data," *Quarterly Journal of Economics*, 131, 519–78.

Santos, M., and M. Woodford (1997): "Rational Asset Pricing Bubbles," *Econometrica*, 65, 19–57.

Scheinkman, J. A. (2014): *Speculation, Trading, and Bubbles.* Columbia University Press.
Scheinkman, J. A., and W. Xiong (2003): "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111, 1183–1219.
Shefrin, H., and M. Statman (1985): "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, 40, 777–90.
Shiller, R. J. (2015): *Irrational Exuberance.* Princeton University Press, 3rd edn.
Simsek, A. (2013): "Belief Disagreement and Collateral Constraints," *Econometrica*, 81, 1–53.
Sommeiller, E., M. Price, and E. Wazeter (2016): "Income Inequality in the U.S. by State, Metropolitan Area, and County," Economic Policy Institute Report.
Statman, M., S. Thorley, and K. Vorkink (2006): "Investor Overconfidence and Trading Volume," *Review of Financial Studies*, 19, 1531–65.
Stokey, N. L. (2009): "Moving Costs, Nondurable Consumption, and Portfolio Choice," *Journal of Economic Theory*, 144, 2419–39.
Vissing-Jorgensen, A. (2003): "Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence From Expectations and Actions," in *NBER Macroeconomics Annual, 2003*, ed. by M. Gertler, and K. Rogoff. MIT Press.
Wang, J. (1994): "A Model of Competitive Stock Trading Volume," *Journal of Political Economy*, 102, 127–68.
Xiong, W. (2013): "Bubbles, Crises, and Heterogeneous Beliefs," NBER Working Paper No. 18905.