

Extrapolation and Bubbles

Nicholas Barberis, Robin Greenwood, Lawrence Jin, Andrei Shleifer

September 2015*

Abstract

We present an extrapolative model of bubbles. In the model, many investors form their demand for a risky asset by weighing two signals—an average of the asset’s past price changes and the asset’s degree of overvaluation. The two signals are in conflict, and investors “waver” over time in the relative weight they put on them. The model predicts that good news about fundamentals can trigger large price bubbles. We analyze the patterns of cash-flow news that generate the largest bubbles, the reasons why bubbles collapse, and the frequency with which they occur. The model also predicts that bubbles will be accompanied by high trading volume, and that volume increases with past asset returns. We present empirical evidence that bears on some of the model’s distinctive predictions.

*The authors’ affiliations are Yale School of Management, Harvard Business School, California Institute of Technology, and Harvard University, respectively. Comments are welcome.

1 Introduction

In the classical account of a financial market bubble, the price of an asset rises dramatically over the course of a few months or even years, reaching levels that appear to far exceed reasonable valuations of the asset’s future cash flows. These price increases are accompanied by widespread speculation and high trading volume. The bubble eventually ends with a crash, in which prices collapse even more quickly than they rose. Bubble episodes have fascinated economists and historians for centuries (e.g., Mackay 1841, Bagehot 1873, Galbraith 1954, Kindleberger 1978, Shiller 2000), in part because human behavior in bubbles is so hard to explain, and in part because of the devastating side effects of the crash.

At the heart of the standard historical narratives of bubbles is the concept of extrapolation—the formation of expected returns by investors based on past returns. In these narratives, extrapolators buy assets whose prices have risen because they expect them to keep rising. According to Bagehot (1873), “owners of savings . . . rush into anything that promises speciously, and when they find that these specious investments can be disposed of at a high profit, they rush into them more and more.” These historical narratives are supported by more recent research on investor expectations, using both survey data and lab experiments. Case, Shiller, and Thompson (2012) show that in the U.S. housing market, homebuyers’ expectations of future house price appreciation are closely related to lagged house price appreciation. Greenwood and Shleifer (2014) present survey evidence of expectations of stock market returns and find strong evidence of extrapolation, including during the internet bubble. Extrapolation also shows up in data on expectations of participants in experimental bubbles, where subjects can be explicitly asked about their expectations of returns. Both the classic study of Smith, Suchanek, and Williams (1988) and more recent experiments such as Haruvy, Lahav, and Noussair (2007) find direct evidence of extrapolative expectations during a well-defined experimental price bubble.

In this paper, we present a new model of bubbles based on extrapolation. In doing so, we seek to shed light on two key features commonly associated with bubbles. The first is what Kindleberger (1978) called “displacement”—the fact that nearly all bubbles from tulips to South Sea to the 1929 U.S. stock market to the late 1990s internet occur on the back of good

fundamental news. Among other questions, we would like to understand which patterns of news are likely to generate the largest bubbles, and whether a bubble can survive once the good news comes to an end. Second, we would like to explain the crucial fact that bubbles feature very high trading volume (Galbraith 1954, Carlos, Neal, and Wandschneider 2006, Hong and Stein 2007). At first sight, it is not clear how extrapolation can explain this: if, during a bubble, all extrapolators have similarly bullish views, then they would not trade with each other.

To address these questions, we present a model in the spirit of earlier work by Cutler, Poterba, and Summers (1990), De Long et al. (1990), Barberis and Shleifer (2003), and Barberis et al. (2015), but with some significant new elements.¹ There is a risk-free asset and a risky asset that pays a liquidating cash flow at a fixed time in the future. Each period, news about the value of the final cash flow is publicly released. There are two types of investors. The first type is extrapolators, who form their share demand based on an extrapolative “growth signal”, which is a weighted average of past price changes. In a departure from prior models, extrapolators also put some weight on a “value signal” which measures the difference between the price and a rational valuation of the final cash flow. The two signals, which can be interpreted as “greed” and “fear”, give the extrapolator conflicting prompts. If prices have been rising strongly and the asset is overvalued, the growth signal encourages him to buy (“greed”) while the value signal encourages him to sell (“fear”).

Our second departure from prior models is to assume that, at each date, and independently of other extrapolators, each extrapolator slightly but randomly shifts the relative weight he puts on the two signals. This assumption, which we refer to as “wavering”, reflects extrapolators’ ambivalence about how to balance the conflicting signals they face. Importantly, the degree of wavering is constant over time. We show that wavering can plausibly account for a good deal of evidence other models have trouble with.

As in earlier models, extrapolators are met in the market by fundamental traders who

¹These earlier papers use models of return extrapolation to address fundamental patterns in asset prices such as excess volatility, return predictability, and nonzero return autocorrelations. They do not discuss bubbles. Glaeser and Nathanson (2015) analyze housing bubbles using a return extrapolation framework, albeit one very different from our own.

lean against the wind, buying the asset when its price is low relative to their valuation of the final cash flow and selling when its price is high. Both extrapolators and fundamental traders face short-sale constraints.

In line with Kindleberger's notion of displacement, a bubble forms in our model after a sequence of large positive cash-flow shocks. The bubble evolves in three stages. In the first stage, the cash-flow news pushes up the asset's price; extrapolators sharply increase their demand for the risky asset, buying from fundamental traders. In the second stage, the asset becomes sufficiently overvalued that the fundamental traders exit the market, leaving the asset in the hands of the exuberant extrapolators who trade with each other because of wavering. Once the good cash-flow news subsides, prices stop rising as rapidly, extrapolator enthusiasm abates, and the bubble begins its collapse. In the third stage, prices fall far enough that fundamental traders re-enter the market, buying from extrapolators.

Our model sheds light on the displacement process that gives rise to the bubble. The largest bubbles arise from sequences of cash-flow shocks that first increase in magnitude, and then decrease. Wavering can significantly increase the size of a bubble through a novel mechanism that we call a "price spiral". During a bubble, the asset can become so overvalued that even some extrapolators hit their short-sale constraints. The bubble selects only the most bullish investors as asset holders, which leads to an even greater overvaluation, causing even more extrapolators to leave. The bubble takes on a life of its own, persisting well after the end of the positive cash-flow news.

The model predicts substantial volume in the first and third stages of a bubble, as fundamental traders sell to extrapolators and vice-versa. But it predicts particularly intense trading during the height of the bubble as extrapolators, as a consequence of wavering, trade among themselves. During normal times, wavering has very little impact on trading volume because it is very slight. During bubbles, in contrast, the same small degree of wavering that generates little volume in normal times endogenously generates intense volume: the growth and value signals that extrapolators attend to are now so large in magnitude that even tiny shifts in their relative weightings lead to large portfolio adjustments. One manifestation of such adjustments, exemplified by Isaac Newton's participation in the South Sea bubble, is extrapolators getting in, out, and back in the market.

After presenting the model, we compare it to two standard approaches to modeling bubbles: rational bubbles (Blanchard and Watson 1982, Tirole 1985) and disagreement (Harrison and Kreps 1978, Scheinkman and Xiong 2003). Models of rational bubbles assume homogeneous investors and therefore cannot explain any volume, let alone highly specific patterns of volume documented in the literature. In addition, models of rational bubbles do not rely on Kindleberger’s displacement, which seems central to most historical episodes. Disagreement-based models can explain why volume is high during bubbles. In these models, the increase in volume during a bubble is due to an exogenous increase in disagreement. In our framework, in contrast, the increase in volume is due to disagreement that grows endogenously over the course of the bubble. Our framework is also more successful at matching the extrapolative expectations many investors hold during bubble periods, as well as the high correlation between volume and past returns observed during these episodes.

Finally, we examine empirically some of the model’s predictions. Using data from four historical bubbles, we document that increases in trading volume during a bubble are strongly predicted by high past returns. For the technology bubble of the late 1990s, we also show that, as the bubble progresses, it draws in new investors with extrapolator-like characteristics. Some of the less obvious predictions of the model are thus consistent with empirical evidence.

Some recent research has questioned whether bubble-like price episodes are actually irrational (Pastor and Veronesi 2006) or whether bubbles in the sense of prices undeniably and substantially exceeding fundamentals over a period of time ever exist (Fama 2014). Although the existence of bubbles in this sense appears uncontroversial in experimental (Smith, Suchanek, and Williams 1988) or some unusual market (Xiong and Yu 2011) settings, our paper does not speak to these controversies. Rather, we present a model of a market for a security in which price patterns commonly described as bubbles materialize infrequently, and other phenomena such as repeated good fundamental news prior to the bubble and very high trading volume also obtain. The model generates new predictions about the structure of price bubbles, such as the ability of past returns to predict volume. Some of these distinctive predictions are consistent with the data.

In the next section, we present our model. Sections 3 and 4 describe circumstances under which bubbles occur and present our findings for price patterns and volume. Section 5

considers the possibility of negative bubbles. Section 6 compares our model to other models of bubbles while Section 7 presents the empirical evidence. Section 8 concludes. All proofs are in the Appendix.

2 A model of bubbles

We consider an economy with $T + 1$ dates, $t = 0, 1, \dots, T$. There are two assets: one risk-free and one risky. The risk-free asset earns a constant return which we normalize to zero. The risky asset, which has a fixed supply of Q shares, is a claim to a dividend \tilde{D}_T paid at the final date, T . The value of \tilde{D}_T is given by

$$\tilde{D}_T = D_0 + \tilde{\varepsilon}_1 + \dots + \tilde{\varepsilon}_T, \quad (1)$$

where

$$\tilde{\varepsilon}_t \sim N(0, \sigma_\varepsilon^2), \text{ i.i.d. over time.} \quad (2)$$

The value of D_0 is public information at time 0, while the realized value of $\tilde{\varepsilon}_t$ becomes public information at time t . The price of the risky asset, P_t , is determined endogenously.

There are two types of traders in the economy: fundamental traders and extrapolators. The time t per-capita demand of fundamental traders for shares of the risky asset is

$$\frac{D_t - \gamma\sigma_\varepsilon^2(T - t - 1)Q - P_t}{\gamma\sigma_\varepsilon^2}, \quad (3)$$

where $D_t = D_0 + \sum_{j=1}^t \varepsilon_j$ and γ is fundamental traders' coefficient of absolute risk aversion.

In the Appendix, we show how this expression can be derived from utility maximization. In brief, it is the demand of an investor who, at each time, maximizes a constant absolute risk aversion (CARA) utility function defined over next period's wealth, and who is boundedly rational: he uses backward induction to determine his time t demand, but, at each stage of the backward induction process, he assumes that, in future periods, the other investors in the economy will simply hold their per-capita share of the risky asset supply. In other words, he does not have a detailed understanding of how other investors in the economy form their share demands. For this investor, the expression $D_t - \gamma\sigma_\varepsilon^2(T - t - 1)Q$ in the numerator of (3) is the expected price of the risky asset at the next date, date $t + 1$. The

numerator is therefore the expected price *change* over the next period, and the fundamental trader’s demand is this expected price change scaled by the trader’s risk aversion and by his estimate of the risk he is facing. If all investors in the economy were fundamental traders, then, setting the expression in (3) equal to the risky asset supply of Q , the equilibrium price of the risky asset would be

$$D_t - \gamma\sigma_\varepsilon^2(T - t)Q. \quad (4)$$

We call this the “fundamental value” of the risky asset and denote it by P_t^F .²

The second type of traders in the economy is extrapolators. There are I types of extrapolators, indexed by $i \in \{1, 2, \dots, I\}$; we explain below how one type of extrapolator differs from another. We build up our specification of extrapolator demand for the risky asset in three steps. An initial specification of per-capita extrapolator share demand at time t is

$$\frac{X_t}{\gamma\sigma_\varepsilon^2}, \quad \text{where } X_t \equiv (1 - \theta) \sum_{k=1}^{t-1} \theta^{k-1} (P_{t-k} - P_{t-k-1}) + \theta^{t-1} X_1, \quad (5)$$

and where $0 < \theta < 1$.

In the Appendix, we show that this is the demand of an investor who, at each time, maximizes a CARA utility function defined over next period’s wealth, and whose belief about the expected price change of the risky asset over the next period is a weighted average of past price changes, with more recent price changes weighted more heavily. The parameter X_1 is a constant that measures extrapolator enthusiasm at time 1; in our numerical analysis, we assign it a neutral, steady-state value. The specification in (5) is similar to that in previous models of extrapolative beliefs, which have been used to shed light on excess volatility in asset prices, the predictability of returns using the dividend-price ratio, and momentum and reversals in asset returns (Cutler, Poterba, and Summers 1990, De Long et al. 1990, Hong and Stein 1999, Barberis and Shleifer 2003, Barberis et al. 2015).³

²We assume, for simplicity, that fundamental traders’ estimate of the risk they are facing is given by fundamental risk σ_ε^2 rather than by the conditional variance of price changes. When fundamental traders are the only traders in the economy, this approximation is exact.

³The form of bounded rationality we have assumed for fundamental traders means that these traders expect the price of the risky asset to revert to fundamental value within one period. This, in turn, means that they trade aggressively against any mispricing—more aggressively than if they were fully rational. In the

In this paper, we modify the specification in (5) in two quantitatively small but conceptually significant ways. First, we make extrapolators pay at least some attention to how the price of the risky asset compares to its fundamental value. Specifically, we change the demand function in (5) so that the demand of extrapolator i takes the form

$$w_i \left(\frac{D_t - \gamma\sigma_\varepsilon^2(T-t-1)Q - P_t}{\gamma\sigma_\varepsilon^2} \right) + (1 - w_i) \left(\frac{X_t}{\gamma\sigma_\varepsilon^2} \right). \quad (6)$$

Extrapolator i 's demand is now a weighted average of two components. The second component is the expression we started with in (5), while the first component is the fundamental trader demand from (3); w_i is the weight on the first component. Our framework accommodates any $w_i \in (0, 1]$, but we maintain $w_i < 0.5$ for all i so that the extrapolative component is weighted more heavily. Moreover, since we want to modify the traditional extrapolator demand function in (5) in a quantitatively minor way, we think of w_i as taking a low value. In our numerical work, its value is approximately 0.1. The motivation for (6) is that even extrapolators worry about how the price of the risky asset compares to its fundamental value. A high price relative to fundamental value exerts some downward pressure on their demand, counteracting the extrapolative component.

In what follows, we often refer to the two components of the demand function in (6) as “signals”: the first component, the expression in (3), is a “value” signal, which tells the trader to buy if the price is low relative to fundamental value; the second component, the expression in (5), is a “growth” signal, which tells the trader to buy if prices have recently been rising. These signals typically point in opposite directions. If the price of the risky asset is well above fundamental value, it has probably also been rising recently. The value signal then takes a large negative value, telling the investor to reduce his position, while the growth signal takes a large positive value, telling the investor to increase his position. The signals can be given an informal interpretation in terms of “fear” and “greed”. If the price has recently been rising, the value signal captures extrapolators’ fear that it might fall back to fundamental value, while the growth signal captures greed, their excitement at the prospect of more price rises. If the price has recently been falling, the growth signal latter case, they would recognize that extrapolator demand is persistent and would trade more conservatively against it. Under some conditions, they may even trade in the same direction as the extrapolators (De Long et al. 1990, Brunnermeier and Nagel 2004).

captures extrapolators’ fear of further price declines, and the value signal, their greed—their excitement at the thought of prices rebounding toward fundamental value.

Our second modification is to allow the weight w_i to vary slightly over time, and independently so for each extrapolator type, so that the demand function for extrapolator i becomes

$$w_{i,t} \left(\frac{D_t - \gamma\sigma_\varepsilon^2(T - t - 1)Q - P_t}{\gamma\sigma_\varepsilon^2} \right) + (1 - w_{i,t}) \left(\frac{X_t}{\gamma\sigma_\varepsilon^2} \right), \quad (7)$$

where (7) differs from (6) only in the t subscript added to $w_{i,t}$. Since the demand function in (6) is a weighted average of two signals that often point in opposite directions, the investor is likely to be unsure of what to do—and, in particular, unsure about how much weight to put on each signal at any point in time. As we model it, the weight an extrapolator puts on each signal shifts or “wavers” over time, to a small extent.

To model this wavering, we set

$$\begin{aligned} w_{i,t} &= \bar{w}_i + \min(1 - \bar{w}_i, \bar{w}_i) \tilde{u}_{i,t} \\ \tilde{u}_{i,t} &\sim N(0, \sigma_u^2), \text{ i.i.d. over time and across extrapolators.} \end{aligned} \quad (8)$$

Here, $\bar{w}_i \in (0, 1)$ is the *average* weight that extrapolator i puts on the value signal as opposed to the growth signal; as noted above, we think of \bar{w}_i as a small, positive number. The *actual* weight that extrapolator i puts on the value signal at time t is given by (8).⁴ Note that $\min(1 - \bar{w}_i, \bar{w}_i)$ measures how far \bar{w}_i is from the edge of the $[0, 1]$ interval. The expression in (8) therefore says that, the further \bar{w}_i is from the boundary of the $[0, 1]$ interval, the more he wavers. We think that this is psychologically plausible. If $\bar{w}_i = 0.5$, for example, so that \bar{w}_i is the furthest it can be from the boundary, the investor is paying an equal amount of attention to the two signals, on average. Since these signals are conflicting, he is likely to be particularly unsure of what to do, and therefore particularly unsure what the relative weight on the signals should be. In contrast, an extrapolator whose \bar{w}_i is close to the $[0, 1]$ boundary—for whom $\bar{w}_i = 0.1$, say—is very confident that the growth signal is the one he should pay most attention to; the weight he puts on this signal therefore varies less over time. Aside from its psychological plausibility, the $\min(1 - \bar{w}_i, \bar{w}_i)$ term is also useful

⁴Since $\tilde{u}_{i,t}$ is Normally distributed, $w_{i,t}$ could, in principle, take a value outside the $(0, 1)$ interval. To prevent this from happening, we truncate $\tilde{u}_{i,t}$ at -0.9 and at 0.9 .

because it allows the fundamental trader demand in (3) to be a special case of the more general demand function in (7) and (8)—specifically the case where $\bar{w}_i = 1$. However, our predictions in Sections 3 and 4 do not depend heavily on the $\min(1 - \bar{w}_i, \bar{w}_i)$ term; we obtain similar results if the degree of wavering is independent of \bar{w}_i .

Under our assumptions, the I types of extrapolator differ only in the relative weight $w_{i,t}$ that they put on the value signal as opposed to the growth signal in each period. The values of the two signals themselves are identical across extrapolators.

We also impose short-sale constraints, so that the final risky asset share demand of the fundamental traders, N_t^F , and of extrapolator $i \in \{1, 2, \dots, I\}$, $N_t^{E,i}$, are given by

$$N_t^F = \max \left[\frac{D_t - \gamma \sigma_\varepsilon^2 (T - t - 1) Q - P_t}{\gamma \sigma_\varepsilon^2}, 0 \right] \quad (9)$$

and

$$N_t^{E,i} = \max \left[w_{i,t} \left(\frac{D_t - \gamma \sigma_\varepsilon^2 (T - t - 1) Q - P_t}{\gamma \sigma_\varepsilon^2} \right) + (1 - w_{i,t}) \left(\frac{X_t}{\gamma \sigma_\varepsilon^2} \right), 0 \right]. \quad (10)$$

Short-sale constraints are a realistic feature of financial markets for many investors. In Sections 3 and 4, we explain carefully how they affect our predictions.

The proposition below lays out the equilibrium price of the risky asset in our economy. The fundamental traders make up a fraction μ_0 of the overall population, while extrapolators of type i make up a fraction μ_i , so that $\sum_{i=0}^I \mu_i = 1$.⁵

Proposition 1. In the economy described above, a market-clearing price always exists and is determined as follows. Let \bar{P}_i , $i \in \{0, 1, \dots, I\}$, be the risky asset price at which trader i 's short-sale constraint starts to bind. Let $N_{\bar{P}_i}$ be the aggregate risky asset share demand across all traders when the price equals \bar{P}_i . If $\max_{i \in \{0, 1, \dots, I\}} N_{\bar{P}_i} < Q$, then, in equilibrium, all traders have a positive demand for the risky asset and the asset's price equals

$$P_t = D_t + \frac{\sum_{i=1}^I \mu_i (1 - w_{i,t})}{\mu_0 + \sum_{i=1}^I \mu_i w_{i,t}} X_t - \gamma \sigma_\varepsilon^2 Q \frac{(\mu_0 + \sum_{i=1}^I \mu_i w_{i,t})(T - t - 1) + 1}{\mu_0 + \sum_{i=1}^I \mu_i w_{i,t}}. \quad (11)$$

Otherwise, let i^* be the value of $i \in \{0, 1, \dots, I\}$ for which $N_{\bar{P}_i}$ exceeds Q by the smallest amount, and let I^* be the set of $i \in \{0, 1, \dots, I\}$ such that trader i has strictly positive demand for the risky asset at price \bar{P}_{i^*} . In this case, in equilibrium, only the traders in I^*

⁵Here and elsewhere, we index fundamental traders by the number “0”.

have a positive demand for the risky asset and the asset's price equals⁶

$$P_t = D_t + \frac{\sum_{i \in I^*} \mu_i (1 - w_{i,t})}{\sum_{i \in I^*} \mu_i w_{i,t}} X_t - \gamma \sigma_\varepsilon^2 Q \frac{(\sum_{i \in I^*} \mu_i w_{i,t})(T - t - 1) + 1}{\sum_{i \in I^*} \mu_i w_{i,t}}. \quad (12)$$

■

The first term on the right-hand side of each of (11) and (12) shows that the price of the risky asset is anchored to the expected value of the final cash flow. The second term reflects the impact of extrapolator demand: if past price changes have been high, so that X_t is high, extrapolator demand at time t is high, exerting upward pressure on the asset price. The third term is a price discount that compensates the holders of the risky asset for the risk they are bearing.

In what follows, it will sometimes be helpful to use, as a benchmark, the “steady state” of our economy, by which we mean the state that the economy converges to after many periods in which the cash-flow shocks equal zero. It is straightforward to show that, after many such periods, the economy reaches a state in which: the fundamental traders and *all* the extrapolators are in the market, with each trader holding the risky asset in proportion to his weight in the population; the price of the risky asset equals the fundamental value in (4); the change in price from one date to the next is constant and equal to $\gamma \sigma_\varepsilon^2 Q$; and the growth signal X_t , defined in (5), is also equal to $\gamma \sigma_\varepsilon^2 Q$.

In Sections 3 and 4, we explore the predictions of our model through both analytical propositions and numerical analysis. We now discuss the parameter values that we use in the numerical analysis. The asset-level parameters are D_0 , Q , σ_ε , and T . The investor-level parameters are I , μ_0 , μ_i and \bar{w}_i for $i \in \{1, \dots, I\}$, γ , θ , and σ_u .

We begin with θ , which governs the weight extrapolators put on recent as opposed to distant past price changes when forming beliefs about future price changes; as such, it determines the magnitude of the growth signal X_t . In setting θ , we are guided by the extensive survey evidence, recently analyzed by Greenwood and Shleifer (2014), on the beliefs of actual stock market investors about future returns. Barberis et al. (2015) use this evidence to formally estimate a parameter analogous to θ . They work in a continuous-time framework;

⁶If $i = 0$ is in the set I^* , the expression in (12) requires the value of $w_{0,t}$, in other words, the weight fundamental traders put on the value signal as opposed to the growth signal. By definition, $w_{0,t} = 1$.

to adapt their estimate for our analysis, we need to specify the length of time between consecutive dates in our model. We take this length of time to be a quarter. Given this, Barberis et al.’s (2015) estimate implies $\theta \approx 0.9$.⁷

We assign benchmark values to the other investor-level parameters as follows. We set μ_0 , the fraction of fundamental traders in the economy, to 0.3, so that fundamental traders make up 30% of the population, and extrapolators, 70%. The survey evidence in Greenwood and Shleifer (2014) suggests that many investors in the economy are extrapolators. We have $I = 50$ types of extrapolators, where each type has the same population weight, so that $\mu_i = (1 - \mu_0)/I$, for $i = 1, \dots, I$. As discussed earlier, we set \bar{w}_i to the same low value of 0.1 for all extrapolators i . And we set γ to 0.1. We do not have strong priors about the value of σ_u , which controls the degree of wavering. We assign it a low value—specifically, 0.3—so as to show that even a small degree of wavering can generate interesting results. Given our specification of wavering in (8), $\sigma_u = 0.3$ and $\bar{w}_i = 0.1$ together imply that, about two-thirds of the time, the weight $w_{i,t}$ extrapolator i puts on the value signal as opposed to the growth signal will be in the interval (0.07, 0.13), a very small degree of wavering.

As for the asset-level parameters, we set the initial expected dividend D_0 to 100, the standard deviation of cash-flow shocks σ_ε to 3, the risky asset supply Q to 1, and the number of dates T to 50. Since the interval between dates is a quarter, this value of T means that the life span of the risky asset is 12.5 years.

3 Asset prices in a bubble

Our model can generate the most essential feature of a bubble, namely a large and growing overvaluation of the risky asset, where, by overvaluation, we mean that the price exceeds the fundamental value in (4). In our model, bubbles are initiated by a sequence of large, positive cash-flow shocks, which here are news about the future liquidating dividend. Figure 1 illustrates this. It uses the parameter values from Section 2 and equations (1), (4), (5), (11),

⁷Specifically, $\theta = \exp(-(0.5)(0.25)) \approx 0.9$, where 0.5 is Barberis et al.’s (2015) estimate of the extrapolation parameter in a continuous-time framework, and 0.25 corresponds to the one-quarter interval between consecutive dates in our framework.

and (12) to plot the price (solid line) and fundamental value (dashed line) of the risky asset for a particular 50-period sequence of cash-flow shocks, in other words, for a particular set of values of $\tilde{\varepsilon}_1, \tilde{\varepsilon}_2, \dots, \tilde{\varepsilon}_{50}$. The first ten shocks, $\tilde{\varepsilon}_1$ through $\tilde{\varepsilon}_{10}$, are all equal to zero. These are followed by four positive shocks, namely 2, 4, 6, and 6; these are substantial shocks: the last two are two-standard deviation shocks. These are followed by 36 more shocks of zero.⁸

Once the positive shocks arrive, a large and sustained overpricing follows. The positive cash-flow news pushes prices up, which leads the extrapolators to sharply increase their share demand in subsequent periods; this, in turn, pushes prices well above fundamental value. Over the four periods of positive cash-flow news, starting at date 11, the expected final dividend increases by 18, the sum of 2, 4, 6, and 6. The figure shows, however, that between dates 11 and 18 *prices* rise by more than double this amount. After the cash-flow shocks drop back to zero at date 15, prices stop rising as rapidly; this, in turn, cuts off the “fuel” driving extrapolator demand. These investors eventually start reducing their demand and the bubble collapses.⁹

The bubble generated by our model has three distinct stages defined by the composition of the investor base. In the first stage, the fundamental traders are still in the market: even though the risky asset is overvalued, the overvaluation is sufficiently mild that the short-sale constraint does not bind for the fundamental traders. In our example, this first stage spans just two dates, 11 and 12. Figure 1 shows that, during this stage, the overvaluation is mild: precisely because the fundamental traders are present in the market, they absorb much of the demand pressure from extrapolators by selling to them.

The second stage of the bubble begins when the risky asset becomes sufficiently over-

⁸We set the value of the growth signal at time 1, X_1 , equal to the steady-state value of X , namely $\gamma\sigma_\varepsilon^2Q = 0.9$. This, together with the fact that the first ten cash-flow shocks are equal to zero, means that the price of the risky asset equals the asset’s fundamental value for the first ten periods.

⁹Our pricing predictions depend, to some extent, on our assumption that fundamental traders cannot short: if they could short, the bubble would be smaller in magnitude. However, there are variants of our model in which the bubble will be large even if fundamental traders *can* short. For example, if fundamental traders can short but are also aware of how extrapolators form their demand, they will recognize that this demand is persistent and will take a less aggressive position against it, thereby allowing for a large bubble (see Barberis et al. 2015, for a formal analysis of this point). Put differently, the short-sale constraint on fundamental traders is not essential for our predictions about prices during a bubble.

valued that the fundamental traders exit the market. In our example, this occurs at date 13. During this stage, extrapolators alone trade the risky asset, which becomes progressively more overvalued as the high past price changes make the extrapolators increasingly enthusiastic. In the absence of cash-flow news, however, the price increases eventually subside, extinguishing extrapolator enthusiasm and causing the bubble to start deflating.

The third stage of the bubble begins when the bubble has deflated to such an extent that the fundamental traders re-enter the market. In our example, this occurs at date 23. In this example, both the fundamental traders *and* the extrapolators are present in the market in this stage. For other cash-flow sequences, the price declines during the collapse of the bubble can be so severe as to cause the extrapolators to exit the market, leaving the asset in the hands of the fundamental traders for some period of time.

To understand more precisely why the bubble in Figure 1 bursts, note that, from price equations (11) and (12), the size of the bubble depends on the magnitude of the growth signal X_t , itself a measure of extrapolator enthusiasm. From equation (5), this growth signal evolves as

$$X_{t+1} = \theta X_t + (1 - \theta)(P_t - P_{t-1}). \quad (13)$$

The first term on the right-hand side, θX_t , indicates that the bubble has a natural tendency to deflate; recall that $0 < \theta < 1$. As time passes, the sharply positive price changes that excited the extrapolators recede further into the past; they are therefore downweighted, by an amount θ , reducing extrapolator enthusiasm. However, if the most recent price change, $P_t - P_{t-1}$, is sufficiently positive, both the growth signal and the bubble itself can maintain their size. Once the good cash-flow news subsides—after date 14, in our example—it becomes increasingly unlikely that the most recent price change is large enough to offset the bubble’s tendency to deflate, in other words, that the second-term on the right-hand side of (13) will dominate the first. As a consequence, the price level starts falling, sharply reducing extrapolator enthusiasm, and setting in motion the collapse of the bubble.¹⁰

Wavering does not play a significant role in the evolution of the price path in Figure 1. If we replaced the extrapolators in our model with extrapolators who all put the same,

¹⁰The use of leverage can amplify the effects of extrapolation, leading to larger bubbles and more dramatic collapses. See Jin (2015) for a detailed analysis of this point.

invariant weight of 0.1 on the value signal, we would obtain a bubble almost identical to that in Figure 1. The reason is that, for the particular sequence of cash-flow shocks used in Figure 1, virtually *all* of the extrapolators are present in the market during all three stages of the bubble. By the law of large numbers, the aggregate demand of $I = 50$ extrapolators whose weight on the value signal equals 0.1 is approximately equal to the aggregate demand of $I = 50$ extrapolators whose weight on the value signal is drawn from a distribution with *mean* 0.1. As a result, the pricing of the risky asset is very similar whether the extrapolators are homogeneous or exhibit wavering.

In some cases, wavering *can* affect the price of the risky asset—in particular, it can significantly amplify the overvaluation of the asset. This is due to a novel bubble mechanism that we call a “price spiral” which evolves in the following way. During the second stage of the bubble, when the fundamental traders are out of the market, the asset can become so overvalued that even some extrapolators exit the market—specifically, those extrapolators with the highest values of $w_{i,t}$ at time t , who put the largest weight on the value signal. Once these extrapolators leave the market, the asset is left in the hands of the most enthusiastic extrapolators, those who put more weight on the growth signal. This generates an even larger overvaluation, causing yet more extrapolators to hit their short-sale constraints, and leaving the asset in the hands of an even more enthusiastic group of extrapolators. Eventually, in the absence of cash-flow shocks, the price increases become less dramatic and extrapolator demand abates, causing the bubble to deflate. At this point, extrapolators who had previously exited the market begin to re-enter.

Figure 2 depicts a price spiral. The parameter values are the same as for Figure 1, but we now use the cash-flow sequence 2, 4, 6, 6, 12, and 10 in place of 2, 4, 6, and 6. The dashed line plots the asset’s fundamental value, while the solid line plots its price. For comparison, the dash-dot line plots the price in an economy where the extrapolators are homogeneous, placing the same, invariant weight of 0.1 on the value signal. The figure shows that, for this cash-flow sequence, wavering significantly amplifies the degree of overpricing: the solid line rises well above the dash-dot line. As explained above, this is due to some extrapolators exiting the market, starting at date 15; at the peak of the price spiral around date 20, about

half of the extrapolators are out of the market.¹¹

The price spiral we have just described can also result from a type of heterogeneity that is simpler than wavering, where extrapolators differ in the weight they put on the value signal, but this weight is constant over time, so that $w_{i,t} = w_i$ for all t . The stochasticity embedded in wavering is not required for price spirals to occur, but becomes crucial for the volume predictions we discuss in Section 4.¹²

So far, we have illustrated the pricing predictions of our model using specific numerical examples. We next present some general analytical results about the dynamics of overpricing. In Proposition 2 below, we calculate the overpricing generated by a sequence of cash-flow shocks. To allow for closed-form expressions, we assume a continuum of extrapolators rather than a finite number.

Proposition 2. Suppose that there is a continuum of extrapolators and that each extrapolator draws an independent weight $w_{i,t}$ at time t from a bounded and continuous density $g(w)$, $w \in [w_l, w_h]$, with mean \bar{w} and with $0 < w_l < w_h < 1$. Suppose that the economy has been in its steady state up to time $l - 1$ and that there is then a sequence of positive shocks $\varepsilon_l, \varepsilon_{l+1}, \dots, \varepsilon_n$ that move the economy from the first stage of the bubble to the second stage of the bubble at some intermediate date j with $l < j < n$. Also suppose that the economy has not moved from the second stage to the third stage by date $N > n$.

No price spiral. If all the extrapolators are in the market at all dates—we specify the condition for this below—the overpricing generated at time t by the cash-flow shocks ε_l ,

¹¹In our model, the price of the risky asset exhibits “excess volatility”: price changes are more volatile than changes in fundamental value. Figure 1 shows that this excess volatility need not have anything to do with wavering; it can be due simply to the presence of extrapolators in the economy. However, Figure 2 shows that, in some cases, wavering creates additional excess volatility.

¹²Price spirals typically deflate within a few periods. In some rare cases, however—specifically, for sequences of very large and positive cash-flow shocks—the price spiral can lead to extremely high prices that revert back only after many periods. We do not put much weight on this prediction. First, the cash-flow shocks required for such out-of-control spirals are so large as to be unlikely in reality. Second, factors absent from our model, such as equity issuance by firms, are likely to prevent these extreme spirals from arising.

$\varepsilon_{l+1}, \dots, \varepsilon_n$ is

$$\mathcal{O}_t \equiv P_t - P_t^F = \begin{cases} \sum_{m=l}^{t-1} \mathcal{L}_1(t-m)\varepsilon_m & l \leq t < j \\ \sum_{m=j}^{t-1} \mathcal{L}_2(t-m)\varepsilon_m + \mathcal{O}_t^1 & j \leq t \leq N \end{cases}, \quad (14)$$

where \mathcal{O}_t^1 is the contribution to the time t overpricing generated by the shocks $\{\varepsilon_i\}_{i=l}^{j-1}$ that occurred during the first stage of the bubble and is given by

$$\mathcal{O}_t^1 = \begin{cases} (\alpha_2 + \theta\alpha_2\alpha_1^{-1})\mathcal{O}_{j-1} - \alpha_2\mathcal{O}_{j-2} + \alpha_2\varepsilon_{j-1} - \frac{\mu_0}{\bar{w}(1-\mu_0)}\gamma\sigma_\varepsilon^2 Q & t = j \\ (\alpha_2 + \theta)\mathcal{O}_{t-1}^1 - \alpha_2\mathcal{O}_{t-2}^1 - \frac{\mu_0(1-\theta)}{\bar{w}(1-\mu_0)}\gamma\sigma_\varepsilon^2 Q & j < t \leq N \end{cases}, \quad (15)$$

where $\alpha_1 \equiv (1-\theta)(1-\mu_0)(1-\bar{w})/[\mu_0 + (1-\mu_0)\bar{w}]$ and $\alpha_2 \equiv (1-\theta)(1-\bar{w})/\bar{w}$. The quantities $\{\mathcal{L}_i(j)\}_{j \geq 0}$ are determined as follows. If $\alpha_i < 2 - \theta - 2\sqrt{1-\theta}$ or $\alpha_i > 2 - \theta + 2\sqrt{1-\theta}$, then

$$\begin{aligned} \mathcal{L}_i(j) &= 2^{-j}\alpha_i[(\alpha_i + \theta)^2 - 4\alpha_i]^{-0.5} \times \\ &\quad \left[\left(\alpha_i + \theta + \sqrt{(\alpha_i + \theta)^2 - 4\alpha_i} \right)^j - \left(\alpha_i + \theta - \sqrt{(\alpha_i + \theta)^2 - 4\alpha_i} \right)^j \right]. \end{aligned} \quad (16)$$

If $2 - \theta - 2\sqrt{1-\theta} < \alpha_i < 2 - \theta + 2\sqrt{1-\theta}$, then

$$\mathcal{L}_i(j) = 2\alpha_i^{0.5j+1} [4\alpha_i - (\alpha_i + \theta)^2]^{-0.5} \sin(j\beta), \quad (17)$$

where $\beta = \cos^{-1}(0.5(\alpha_i + \theta)\alpha_i^{-0.5})$. If $\alpha_i = 2 - \theta + 2\sqrt{1-\theta}$ or $\alpha_i = 2 - \theta - 2\sqrt{1-\theta}$, then

$$\mathcal{L}_i(j) = j\alpha_i^{0.5(j+1)}. \quad (18)$$

Price spiral. If, at some date j' , $j \leq j' \leq N$, the overpricing $\mathcal{O}_{j'}$ computed using (14) is greater than $\bar{\mathcal{O}} \equiv [(1-w_h) + (1-\mu_0)(w_h - \bar{w})]\gamma\sigma_\varepsilon^2 Q / [(1-\mu_0)(w_h - \bar{w})]$, then a price spiral begins at j' . During the spiral, the time t overvaluation is

$$\mathcal{O}_t = \frac{1 - \bar{w}(X_t)}{\bar{w}(X_t)} X_t + \gamma\sigma_\varepsilon^2 Q - \frac{\gamma\sigma_\varepsilon^2 Q}{(1-\mu_0)\eta(X_t)\bar{w}(X_t)}, \quad t \geq j', \quad (19)$$

where

$$\bar{w}(X_t) = \int_{w_l}^{w^*} wg(w)dw / \int_{w_l}^{w^*} g(w)dw, \quad \eta(X_t) = \int_{w_l}^{w^*} g(w)dw, \quad (20)$$

$w^*(X_t)$ is determined by

$$w^*\gamma\sigma_\varepsilon^2 Q = X_t(1-\mu_0) \int_{w_l}^{w^*} (w^* - w)g(w)dw, \quad (21)$$

and the growth signal X_t evolves as

$$X_t = \begin{cases} \frac{\theta[\mu_0 + (1 - \mu_0)\bar{w}](\mathcal{O}_{t-1} - \gamma\sigma_\varepsilon^2 Q) + \theta\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} + (1 - \theta)(\mathcal{O}_{t-1} - \mathcal{O}_{t-2} + \varepsilon_{t-1} + \gamma\sigma_\varepsilon^2 Q) & t = j' \\ \theta X_{t-1} + (1 - \theta)(\mathcal{O}_{t-1} - \mathcal{O}_{t-2} + \varepsilon_{t-1} + \gamma\sigma_\varepsilon^2 Q) & j > j' \end{cases} \quad (22)$$

At each time t , extrapolators with $w_{i,t} < w^*(X_t)$ stay in the market, while those with $w_{i,t} \geq w^*(X_t)$ stay out of the market. If the price spiral ends before time N , equations (19), (20), and (22) still apply but with w^* set to w_h . \blacksquare

Equation (14) gives the magnitude of overvaluation in the absence of a price spiral. To understand it, suppose that, up until time $l - 1$, the economy has been in its steady state, and that, at time l , there is a unit cash-flow shock $\varepsilon_l = 1$, after which the cash-flow shocks revert to zero forever. The quantities $\mathcal{L}_1(1)$, $\mathcal{L}_1(2)$, $\mathcal{L}_1(3)$,... are equal to the overvaluation of the risky asset 1, 2, 3,... periods after the shock, in other words, at dates $l + 1$, $l + 2$, $l + 3$,..., conditional on the bubble staying in the first stage, so that fundamental traders and all extrapolators are in the market. The first row of equation (14) shows that our model has a linear structure, in the sense that, during the first stage of the bubble, the total overvaluation at time t caused by a *sequence* of shocks $\varepsilon_l, \varepsilon_{l+1}, \dots, \varepsilon_{t-1}$ is given by

$$\mathcal{L}_1(1)\varepsilon_{t-1} + \mathcal{L}_1(2)\varepsilon_{t-2} + \dots + \mathcal{L}_1(t-l)\varepsilon_l.$$

Now suppose that the bubble is in the second stage, but with no price spiral, so that the fundamental traders are not in the market but all extrapolators are. Suppose that there is a unit cash-flow shock at time j , $\varepsilon_j = 1$, after which the shocks equal zero forever. The quantities $\mathcal{L}_2(1)$, $\mathcal{L}_2(2)$, $\mathcal{L}_2(3)$,... measure how much *additional* overvaluation this shock creates 1, 2, 3,... periods later, in other words, at dates $j + 1$, $j + 2$, $j + 3$,..., relative to the case in which $\varepsilon_j = 0$, and conditional on all extrapolators staying in the market. The second row of equation (14) shows that, in this second stage of the bubble, the total overvaluation at time t caused by a *sequence* of shocks $\varepsilon_l, \varepsilon_{l+1}, \dots, \varepsilon_n$ has two components. The first is the overvaluation created by the cash-flow shocks that arise during the second stage of the bubble. This is again linear in structure and equals

$$\mathcal{L}_2(1)\varepsilon_{t-1} + \mathcal{L}_2(2)\varepsilon_{t-2} + \dots + \mathcal{L}_2(t-j)\varepsilon_j.$$

The second component of the overvaluation, \mathcal{O}_t^1 , is typically much smaller in magnitude. It is the overvaluation at time t caused by the lingering effect of the cash-flow shocks that occurred during the *first* stage of the bubble.

The quantities $\mathcal{L}_1(\cdot)$ and $\mathcal{L}_2(\cdot)$ can be computed explicitly, as in equations (16), (17), and (18). These expressions show that $\mathcal{L}_i(\cdot)$ can take one of four shapes. The two most common shapes are shown in Figure 3: a curve that rises and then falls monotonically (top panel) and a curve that oscillates with dampening amplitude (bottom panel). The other possibilities are a curve that oscillates with increasing amplitude and a curve that increases monotonically. For our particular parameter values, $\mathcal{L}_1(\cdot)$ is given by the curve in the top panel, while $\mathcal{L}_2(\cdot)$ is given by the curve in the bottom panel.

We learn a number of things from Proposition 2. First, comparing the top and bottom panels in Figure 3, and keeping in mind that the vertical scales are very different in the two panels, we see that a cash-flow shock has a much larger effect on prices in the second stage of the bubble than in the first—it generates an overvaluation that both builds up over a longer period of time and reaches a much higher peak. Because in the second stage the fundamental traders are out of the market, the aggregate demand of the remaining investors loads much more positively on past price changes. The immediate price increase caused by a cash-flow shock then leads to sustained price rises for several subsequent periods and an eventual overvaluation that is large in magnitude.

The proposition also sheds light on another interesting question: What kinds of cash-flow sequences are likely to generate the largest bubbles? It is easier to see the answer using the expressions for the case of no price spiral, although a similar logic applies in the case of spirals. In our model, bubbles attain their largest size during their second stage, when fundamental traders are out of the market. The degree of overvaluation during this second stage is largely determined by the first term in the second row of (14), namely

$$\sum_{m=j}^{t-1} \mathcal{L}_2(t-m)\varepsilon_m.$$

Suppose that there are eight cash-flow shocks that occur over eight dates during the second stage, $\varepsilon_{t-8}, \varepsilon_{t-7}, \dots, \varepsilon_{t-1}$. For the parameter values used in Figure 1, we can read the values of $\mathcal{L}_2(1), \mathcal{L}_2(2), \dots, \mathcal{L}_2(8)$ directly off the bottom panel in Figure 3. This, in turn, allows us

to write the (approximate) degree of overvaluation during the second stage, at time t , as

$$\begin{aligned} & \mathcal{L}_2(1)\varepsilon_{t-1} + \mathcal{L}_2(2)\varepsilon_{t-2} \dots + \mathcal{L}_2(7)\varepsilon_{t-7} + \mathcal{L}_2(8)\varepsilon_{t-8} \\ &= 0.9\varepsilon_{t-1} + 1.62\varepsilon_{t-2} + 2.11\varepsilon_{t-3} + 2.33\varepsilon_{t-4} + 2.3\varepsilon_{t-5} + 2.05\varepsilon_{t-6} + 1.61\varepsilon_{t-7} + 1.06\varepsilon_{t-8}. \end{aligned}$$

From the above expression, we can see what kinds of cash-flow sequences will generate the largest bubbles. In general terms, we want to associate the highest value of ε with the highest value of $\mathcal{L}_2(j)$, namely 2.33; the second highest value of ε with the second highest value of $\mathcal{L}_2(j)$, namely 2.3, and so on. This means two things. First, to produce a large bubble, the cash-flow shocks needs to be clumped together in time, just as the highest values of $\mathcal{L}_2(j)$ are for lags that are temporally close—specifically, lags 3, 4, 5, and 6—something that is clearly visible in the bottom panel of Figure 3. Second, the biggest cash-flow shock should be sandwiched between the two next biggest cash-flow shocks, just as the highest value of $\mathcal{L}_2(j)$, 2.33, is sandwiched between the two next highest values of $\mathcal{L}_2(j)$, 2.3 and 2.11. For example, if the eight possible values of ε_{t-8} through ε_{t-1} are equal to 1, 2, 3, 4, 5, 6, 7, 8, the above discussion suggests that the largest bubble will be generated by the sequence 2, 3, 5, 7, 8, 6, 4, 1—and this is indeed the case. In short, since, for the first few lags, $\mathcal{L}_2(j)$ rises to a peak and then declines, the largest bubble will be created by a sequence of cash flows that itself rises to a peak and then declines.

How common are large bubbles in our framework? To compute the frequency with which they occur, we use the cash-flow process in (1) and the price processes in (11) and (12) to simulate a $T = 40,000$ -period price sequence and record the number of bubbles for which the level of overvaluation—the bubble’s “size”—exceeds a threshold such as 10 or 20, and also the length of time for which this threshold is exceeded. To put these bubble sizes in context, recall that, in non-bubble times, a one-standard deviation cash-flow shock increases the asset’s price by approximately 3.

In our model, bubbles are rare. For our benchmark parameter values, a bubble whose size exceeds 10 occurs once every 17 years, on average, with the overvaluation exceeding 10 for approximately one year. A bubble of size 20 occurs just once every 50 years, on average, and maintains this size for approximately three quarters. Bubbles are rare for two reasons. First, for a bubble to occur, the cash-flow shocks need to be large enough to cause the

fundamental traders to exit; when fundamental traders absorb extrapolator demand, they prevent a bubble from forming. Second, as described above, for a large bubble to form, the cash-flow shocks need both to be large *and* to be clumped together in time. The probability of this happening is low.¹³

To conclude our analysis of prices, we verify, through simulations, that our model also captures the basic asset pricing patterns that the previous generation of extrapolation models was designed to explain. Specifically, we confirm that our model generates excess volatility (the standard deviation of price changes exceeds the standard deviation of changes in fundamental value); predictability (the difference between the dividend and the price at time t , $D_t - P_t$, predicts the change in price over the next 12 periods, $P_{t+12} - P_t$); and positive (negative) autocorrelations in price changes at short (long) lags. It is not surprising that our framework can explain these facts: while we modify the earlier extrapolation models in qualitatively significant ways, these modifications are *quantitatively* small.

4 Volume in a bubble

Bubbles feature very high trading volume (Ofek and Richardson 2003, Hong and Stein 2007). A central goal of our paper is to propose a way of understanding this fact.¹⁴

Figure 4 plots the share demand N_t^F of the fundamental traders (dashed line) and the share demands $N_t^{E,i}$ of the I types of extrapolator (solid lines) for the same 50-period cash-flow sequence that we used in Figure 1, namely 10 shocks of zero, followed by four positive shocks of 2, 4, 6, and 6, followed by 36 more shocks of zero. Recall from Figure 1 that this sequence of cash-flow shocks generates a large bubble between dates 15 and 21.

¹³If we raise θ to 0.95, bubbles are even rarer. In this case, extrapolators form their demand by averaging a very long sequence of prior price changes. Since only a few of these price changes are likely to be strongly positive, extrapolators only rarely become enthusiastic enough to trigger a bubble.

¹⁴A small fraction of bubbles, often those associated with debt securities, do not feature very high trading volume. Hong and Sraer (2013) explain this by noting that, if investors are over-optimistic about the value of the asset underlying a debt security and also differ in how optimistic they are, they overvalue the debt security but do not disagree about its value—its value does not depend on beliefs about good states of the world. Trading in the debt security is therefore muted.

Figure 4 shows that, during the bubble, the share demands of extrapolators become very volatile, suggesting a large increase in volume. Figure 5 confirms this. The solid line in this figure plots total trading volume at each of the 50 dates, and shows a dramatic increase in volume between dates 12 and 25. More specifically, Figure 5 shows that our model predicts a structure for volume during bubbles that consists of three “peaks” that correspond to the three bubble stages outlined in Section 3: there is a small peak centered around date 13 in the first stage, a much wider peak centered around date 17 in the second stage, and a thin but tall peak centered around date 23 in the third stage. To understand these three peaks, it is useful to look at Figure 5 in conjunction with Figure 4, but also at the dashed line in Figure 5. In our model, total volume is the sum of two components: trading that takes place within the set of I extrapolators, and trading that takes place between the extrapolators in aggregate and fundamental traders. The dashed line in Figure 5 plots the first component—trading volume within the set of I extrapolators. Between dates 14 and 22, the solid and dashed lines in Figure 5 lie on top of each other, indicating that all of the trading volume during this time takes place among the I extrapolators.

The first peak in Figure 5 centered around date 13 arises during the first stage of the bubble and reflects trading between the extrapolators in aggregate and fundamental traders. Arrival of the good cash-flow news pushes prices up, which, in turn, leads extrapolators to sharply increase their demand for the risky asset and fundamental traders to sharply reduce their demand. The volume during this stage of the bubble therefore consists of fundamental traders selling to extrapolators. Before long, however, all the fundamental traders are out of the market and this first wave of trading subsides.

After the fundamental traders exit the market, the bubble enters its second stage. During this stage, the bubble keeps growing and trading volume rises again, as indicated by the wide second peak centered around date 17 in Figure 5. Figure 4 and the dashed line in Figure 5 show that this trading occurs among the I extrapolators. This is the most interesting and potentially large volume generated by our model.

The high trading volume represented by the wide peak in Figure 5 is due to wavering—the assumption that, in each period, each extrapolator slightly shifts the relative weight he puts on the value as opposed to the growth signal. It is not surprising that, in general, wavering

leads to trading volume. What is more interesting is that, even though the degree of wavering remains *fixed* over time—the value of σ_u in equation (8) is constant—the model endogenously generates much greater volume during bubble periods than non-bubble periods. Figure 4 shows this most starkly. Even though the degree of extrapolator wavering is constant over the 50 dates in the graph, wavering generates dramatically higher variation in extrapolator share demand during the bubble period than non-bubble periods.

To understand this, we write the share demand of extrapolator i in equation (10) more simply as $w_{i,t}V_t + (1 - w_{i,t})G_t$, where V_t and $G_t = X_t/\gamma\sigma_\varepsilon^2$ are the value and growth signals, respectively, at time t . We ignore the short-sale constraint because it is not important for the intuition. A unit change in the relative weight $w_{i,t}$ on the two signals changes the extrapolator’s share demand by $V_t - G_t$. In “normal” times, when the cash-flow shocks are neither abnormally high nor abnormally low, the value and growth signals are both small in absolute magnitude: if the risky asset is neither particularly overvalued nor undervalued, the value signal V_t is close to zero in absolute magnitude; and if prices have not been rising or falling particularly sharply in recent periods, the growth signal G_t is also close to zero in absolute magnitude. In this case, $V_t - G_t$ is itself low in absolute magnitude, implying that, in normal times, wavering does not induce much variation in extrapolator demand. This can be seen in the first 10 and last 25 periods of Figure 4.

During a bubble, the situation is very different. At that time, the value signal V_t is large and *negative* (the asset is highly overvalued), and the growth signal is large and *positive* (the asset’s price has been rising sharply in recent periods). As a result, $V_t - G_t$ is *very* large in absolute value, and the same degree of wavering that generates low trading volume in normal times now generates very high trading volume. This is the mechanism behind the high trading volume represented by the wide peak centered around date 17 in Figure 5.

To put this more simply, during the bubble, the extrapolators holding the risky asset are subject to two powerful but conflicting investment signals. On the one hand, they see that prices are far above fundamental value; this makes them fearful of a crash and encourages them to sell. On the other hand, prices have recently been rising sharply, which makes extrapolators expect continued price appreciation and encourages them to buy. These two signals are so strong that even small shifts in the relative weight extrapolators put on the

two lead to large portfolio adjustments, and hence trading volume.

Once the bubble starts collapsing, the second wave of trading volume begins to subside: as the bubble deflates, both the value and growth signals decline in absolute magnitude; the quantity $V_t - G_t$ then also declines in absolute magnitude, and the impact of wavering on extrapolator share demands is reduced. Figure 5 shows, however, that once the bubble's collapse is well under way, there is a third wave of trading, represented by the thin third peak centered around date 23. This third peak is associated with the third stage of the bubble and represents trading between the extrapolators in aggregate and the fundamental traders: the extrapolators sell the risky asset to fundamental traders, who now re-enter the market. The third peak is taller than the first peak. The reason is that the first peak consists of extrapolators shifting from moderate holdings of the risky asset to large holdings of the asset. The third peak, however, consists of extrapolators shifting from large holdings of the risky asset to *low* holdings of the asset; the low holdings are driven by the poor performance of the risky asset during the bubble's collapse, which leads extrapolators to form pessimistic beliefs about the future and hence to choose a low exposure to the risky asset. This third volume peak therefore represents more intense trading than the first one.

The key message in the above discussion is that a fixed amount of wavering can endogenously generate much higher trading volume during bubble periods. We formalize this idea next.

Suppose that we are in the first stage of the bubble, when the fundamental traders and *all* the extrapolators are in the market. Substituting the expression for price in (11) into the demand functions in (9) and (10), we see that, during this first stage, the change in fundamental trader demand between time t and time $t + 1$ is approximately

$$N_{t+1}^F - N_t^F = \frac{-(X_{t+1} - X_t)(1 - \mu_0)(1 - \bar{w})}{(\mu_0 + (1 - \mu_0)\bar{w})\gamma\sigma_\varepsilon^2}, \quad (23)$$

while the change in extrapolator i 's demand over the same time interval is approximately¹⁵

$$N_{t+1}^{E,i} - N_t^{E,i} = \frac{(X_{t+1} - X_t)(\mu_0 + (1 - \mu_0)\bar{w} - w_{i,t+1})}{(\mu_0 + (1 - \mu_0)w)\gamma\sigma_\varepsilon^2} + \frac{(w_{i,t+1} - w_{i,t})(\gamma\sigma_\varepsilon^2 Q - X_t)}{(\mu_0 + (1 - \mu_0)\bar{w})\gamma\sigma_\varepsilon^2}. \quad (24)$$

¹⁵The approximation is due to the fact that our model features a finite number of extrapolators; with a continuum of extrapolators, (23) and (24) hold exactly.

Equation (23) shows that, when the size of the bubble, proxied by the growth signal X_t , grows in magnitude, fundamental traders reduce their demand: they perceive the risky asset to be overvalued. Equation (24) shows that the change in extrapolator demand has two components. The first—the first term on the right-hand side—has nothing to do with wavering. It says that, as the bubble grows in size, extrapolators gradually increase their holdings: for most extrapolators, $w_{i,t+1}$ is smaller than $\mu_0 + (1 - \mu_0)\bar{w}$. The second component, in contrast, is driven entirely by wavering. It says that, even if the bubble does not increase in size between time t and time $t + 1$, in the sense that $X_{t+1} = X_t$, extrapolator i 's demand will nonetheless change because of wavering: if $w_{i,t+1} < w_{i,t}$, so that he puts more weight on the growth signal at time $t + 1$, then since, during a bubble, $X_t > \gamma\sigma_\varepsilon^2 Q$, his demand for the risky asset will go up.

Now suppose that we are in the second stage of the bubble, after the fundamental traders have left the market. Substituting the expression for price in (12) into the demand function in (10), we see that the change in share demand for an extrapolator who is in the market both at time t and at time $t + 1$ is approximately

$$N_{t+1}^{E,i} - N_t^{E,i} = \frac{(X_{t+1} - X_t)(\bar{w} - w_{i,t+1})}{\bar{w}\gamma\sigma_\varepsilon^2} + \frac{(w_{i,t+1} - w_{i,t})(\gamma\sigma_\varepsilon^2 Q - (1 - \mu_0)X_t)}{(1 - \mu_0)\bar{w}\gamma\sigma_\varepsilon^2}. \quad (25)$$

Once again, the change in extrapolator demand has two components. The first is independent of wavering. It says that, if the bubble grows in size, those extrapolators who put more weight on the growth signal than the typical extrapolator ($w_{i,t+1} < \bar{w}$) increase their demand for the risky asset, while those who put less weight decrease their demand. The second term, in contrast, is entirely driven by wavering. It indicates that, even if the bubble does not grow in size between time t and time $t + 1$, in the sense that $X_{t+1} = X_t$, the extrapolators who increase the weight they put on the growth signal ($w_{i,t+1} < w_{i,t}$) increase their demand: during the second stage of a bubble, $(1 - \mu_0)X_t > \gamma\sigma_\varepsilon^2 Q$. The second term on the right-hand side of each of equations (24) and (25) captures the key insight described earlier. As the bubble grows in size, $\gamma\sigma_\varepsilon^2 Q - X_t$ and $\gamma\sigma_\varepsilon^2 Q - (1 - \mu_0)X_t$ become larger in absolute magnitude. This, in turn, means that a fixed shift $w_{i,t+1} - w_{i,t}$ in the relative weight an extrapolator puts on the value and growth signals has a larger impact on his demand.

To formally verify this claim, namely that the larger the size of the bubble, the more

volume wavering will induce, we now explicitly compute the value of a quantity we call “wavering-induced” volume. This is defined as the trading volume generated in aggregate by the second component of the change in share demand—the second term on the right-hand side of equations (24) and (25).

Proposition 3. Suppose that there is a continuum of extrapolators and that each extrapolator draws an independent weight $w_{i,t}$ at time t from a bounded and continuous density function $g(w)$, $w \in [w_l, w_h]$, with mean \bar{w} and with $0 < w_l < w_h < 1$. The per-capita wavering-induced trading volume for extrapolators, denoted as $V^W(X_t)$, is

$$V^W(X_t) = \begin{cases} \frac{|X_t - \gamma\sigma_\varepsilon^2 Q|\bar{\Delta}_0}{(\mu_0 + (1 - \mu_0)\bar{w})\gamma\sigma_\varepsilon^2} - \frac{w_l\gamma\sigma_\varepsilon^2 Q}{\mu_0(1 - w_l) + (1 - \mu_0)(\bar{w} - w_l)} \leq X_t < \frac{\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \\ \frac{((1 - \mu_0)X_t - \gamma\sigma_\varepsilon^2 Q)\bar{\Delta}_0}{(1 - \mu_0)\bar{w}\gamma\sigma_\varepsilon^2} & \frac{\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \leq X_t \leq \frac{w_h\gamma\sigma_\varepsilon^2 Q}{(w_h - \bar{w})(1 - \mu_0)} \\ \frac{(\eta(X_t)X_t - \gamma\sigma_\varepsilon^2 Q(1 - \mu_0)^{-1})\bar{\Delta}(X_t)}{\eta(X_t)\bar{w}(X_t)\gamma\sigma_\varepsilon^2} + \frac{2(1 - \eta(X_t))Q}{1 - \mu_0} & X_t > \frac{w_h\gamma\sigma_\varepsilon^2 Q}{(w_h - \bar{w})(1 - \mu_0)} \end{cases} \quad (26)$$

where

$$\begin{aligned} \eta(X_t) &\equiv \int_{w_l}^{w_\eta(X_t)} g(w)dw, & \bar{w}(X_t) &\equiv \eta^{-1}(X_t) \int_{w_l}^{w_\eta(X_t)} g(w)wdw, \\ \bar{\Delta}_0 &\equiv \int_{w_l}^{w_h} g(w_1)dw_1 \int_{w_l}^{w_h} |w_1 - w_2|g(w_2)dw_2, \\ \bar{\Delta}(X_t) &\equiv \int_{w_l}^{w_\eta(X_t)} g(w_1)dw_1 \int_{w_l}^{w_\eta(X_t)} |w_1 - w_2|g(w_2)dw_2, \end{aligned} \quad (27)$$

and where $w_\eta(X_t)$ is determined as the implicit solution to

$$(1 - \mu_0)w_\eta \left(\int_{w_l}^{w_\eta} g(w)dw \right) X_t - w_\eta\gamma\sigma_\varepsilon^2 Q = (1 - \mu_0) \left(\int_{w_l}^{w_\eta} g(w)wdw \right) X_t. \quad (28)$$

■

The three rows of (26) correspond to: the first stage of the bubble, when the fundamental traders and *all* the extrapolators are in the market; the less extreme part of the second stage of the bubble, when the fundamental traders are out of the market but all the extrapolators are still in; and the more extreme part of the second stage when even some of the extrapolators are out of the market.

We can use the expressions in Proposition 3 to compute the sensitivity of wavering-induced volume to the size of the bubble, as proxied by X_t .

Corollary 1. The sensitivity of wavering-induced trading volume to the growth signal X_t , denoted by $\partial V^W(X_t)/\partial X_t$, is given by

$$\frac{\partial V^W(X_t)}{\partial X_t} = \begin{cases} \frac{\text{sign}(X_t - \gamma\sigma_\varepsilon^2 Q)\bar{\Delta}_0}{(\mu_0 + (1 - \mu_0)\bar{w})\gamma\sigma_\varepsilon^2} & -\frac{w_l\gamma\sigma_\varepsilon^2 Q}{\mu_0(1 - w_l) + (1 - \mu_0)(\bar{w} - w_l)} \leq X_t < \frac{\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \\ \frac{\bar{\Delta}_0}{\bar{w}\gamma\sigma_\varepsilon^2} & \frac{\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \leq X_t \leq \frac{w_h\gamma\sigma_\varepsilon^2 Q}{(w_h - \bar{w})(1 - \mu_0)} \end{cases} \quad (29)$$

If $X_t > w_h\gamma\sigma_\varepsilon^2 Q/[(w_h - \bar{w})(1 - \mu_0)]$, $\partial V^W(X_t)/\partial X_t$ may become smaller and even turn negative as extrapolators exit the market. ■

The Corollary confirms the intuition we outlined above. Most crucially, in the less extreme part of the second stage of the bubble, when all extrapolators are in the market—this corresponds to the second row of (29)—wavering induces more trading volume, the larger the size of the bubble: $\bar{\Delta}_0$ is a positive quantity. The same is true during the first stage of the bubble—see the first row of (29)—although, during this stage, wavering-induced volume is a relatively small part of overall trading volume. If, during its second stage, the bubble becomes so large that even some extrapolators exit the market, then wavering-induced volume increases more slowly as a function of X_t , and can even decrease, simply because there are fewer extrapolators available to trade.

The above analysis indicates that volume is typically increasing in X_t . Since X_t is an average of past price changes, this suggests the following testable prediction: during a bubble, volume will be positively related to the asset’s past return. To verify that this is a prediction of our model, we simulate a long time series of prices from the model and extract three subsamples—the subsample where the asset is overvalued by $\gamma\sigma_\varepsilon^2 Q = 0.9$; the subsample where it is overvalued by at least $10\gamma\sigma_\varepsilon^2 Q = 9$; and the subsample where it is overvalued by at least $20\gamma\sigma_\varepsilon^2 Q = 18$.¹⁶ We find that in these three subsamples, the correlation between volume at time $t + 1$ and the price change between $t - 4$ and t , a year-long interval, is 0.15,

¹⁶The quantity $\gamma\sigma_\varepsilon^2 Q$ is the degree of overvaluation that causes fundamental traders to exit the market; it is therefore a convenient “unit” of overvaluation.

0.31, and 0.36, respectively. In Section 7.1, we test the prediction that volume is related to the past year’s return for four bubble episodes.

We conclude our discussion of trading volume with two points. First, the short-sale constraint on fundamental traders is important for our predictions about volume during a bubble. If fundamental traders could short, then, even at the height of the bubble, a significant fraction of the volume would be trading between the fundamental traders and the extrapolators in aggregate. When fundamental traders cannot short, however, volume at the height of the bubble is entirely driven by trading within a group of *similar* traders, namely extrapolators.

Second, alternative sources of heterogeneity among extrapolators—sources other than wavering—do not generate nearly as much trading volume during the bubble period. Specifically, if we turn off wavering by setting σ_u in (8) to 0 and instead allow the base weights \bar{w}_i and the weighting parameter θ to differ across extrapolators, we no longer see a large second volume peak like the one in Figure 5. The reason is that, after a sequence of price increases, extrapolators who do not exhibit wavering would almost all like to increase their holdings of the risky asset, even if they differ in their values of \bar{w}_i and θ : regardless of the specific values of \bar{w}_i and θ , a high growth signal means that most extrapolators will find the risky asset more attractive. Since most extrapolators want to trade in the same direction, there is relatively little trading between them: it is prices that adjust, not quantities. It is not that *any* heterogeneity across extrapolators could generate heavy trading during a bubble; in our model, the specific type of heterogeneity induced by wavering appears uniquely able to do so.

5 Negative bubbles

We have also studied the behavior of prices and volume after a sequence of negative cash-flow shocks. The results in this case are *not* the “mirror image” of those for the case of positive cash-flow shocks.

First, our model does not generate “negative” bubbles: while the price of the risky asset falls when bad cash-flow news arrives, it does not fall much below fundamental value. After

disappointing cash-flow news push the price of the risky asset down several periods in a row, the extrapolators would, in principle, like to short the risky asset. If they *could* short, they would cause the risky asset to become very undervalued. However, since they are subject to a short-sale constraint, they stay out of the market, and there is no significant undervaluation. This result is consistent with the observation that, while, in actual financial markets, assets are sometimes perceived to be experiencing a bubble, they are rarely perceived to be experiencing a *negative* bubble.

Even though the risky asset does not become significantly undervalued in bad times, it is nonetheless undervalued by at least some amount. In bad times, the risky asset is held by just one segment of the investor population, namely fundamental traders. To hold the entire market supply, these traders need a price discount relative to fundamental value. It is straightforward to check that, when the asset is held only by the fundamental traders, its price is

$$D_t - \gamma\sigma_\varepsilon^2(T - t - 1)Q - \frac{\gamma\sigma_\varepsilon^2}{\mu_0}Q,$$

which differs from the fundamental value in (4) by $\gamma\sigma_\varepsilon^2Q(1 - \mu_0)/\mu_0$. For our parameter values, this wedge is approximately \$2.

Figure 6 illustrates these points. It plots the price of the risky asset (solid line) and the asset's fundamental value (dashed line) for a 100-period ($T = 100$) cash-flow path that was randomly simulated from the process in (1). The figure shows that, while, for this particular simulation, the asset experiences two bubbles over the 25-year interval—specifically, around dates 20 and 60—it never experiences a negative bubble; at most, it becomes mildly undervalued relative to fundamental value.

Our model predicts heavy trading during bubbles, but little trading during severe downturns. When bad cash-flow news arrives, there is some trading as extrapolators sell to fundamental traders. Once the extrapolators leave the market, however, the asset is held only by fundamental traders, a homogeneous group. There is no more trading until the market recovers and extrapolators re-enter. More broadly, our model predicts much higher trading volume during bull than bear markets, a prediction consistent with the available evidence (Statman, Thorley, and Vorkink 2006, Griffin, Nardari, and Stulz 2007).

6 Comparison with other bubble models

It is impossible to do justice here to all the important contributions in the literature on bubbles, recently surveyed by Brunnermeier and Oehmke (2013) and Xiong (2013). Instead, we focus on two classes of models—rational bubble models and disagreement-based models—and compare our framework to those. We pick out rational bubble models because of their simplicity and long tradition; and we pick out disagreement models because, like us, they deal with volume.

6.1 Rational bubble models

In models of rational bubbles, the price of a risky asset is given by

$$P_t = P_{D,t} + B_t, \tag{30}$$

where $P_{D,t}$ is the present value of the asset’s future cash flows and where B_t , the bubble component, satisfies

$$B_t = \frac{\mathbb{E}(B_{t+1})}{1+r}, \tag{31}$$

where r is the expected return. We emphasize three points relevant to comparing our framework to the rational bubble framework.

First, the rational bubble model does not explain how a bubble gets started in the first place. Under limited liability, the value of B must always be non-negative. But if B is strictly positive in any future state of the world, then, from (31), it must be positive at the current time. Put simply, if a bubble exists, it must always have existed. In our framework, by contrast, bubbles are initiated in a much clearer way, as a consequence of what Kindleberger (1978) calls “displacement”: a sequence of good cash-flow news leads to price increases which, in turn, cause extrapolators to increase their demand for the risky asset.

Second, the rational bubble model has nothing to say about trading volume. In its usual form, agents are assumed to be homogeneous; trading volume is therefore zero.

Third, the rational bubble model does not capture the extrapolative expectations that are often observed during bubbles. In the basic version of this framework, investors expect a constant return on the risky asset at each point in time.

6.2 Disagreement-based models

Building on Harrison and Kreps (1978), Scheinkman and Xiong (2003) present a model in which two risk-neutral investors observe two signals about the fundamental value of a risky asset, but disagree about how useful each signal is. Their disagreement leads to trading volume. With short-sale constraints, their disagreement also leads to overpricing: the price of the risky asset can be higher than the present value of its future cash flows, as perceived by the investor holding the asset. The reason is that the holder of the asset believes that, as more signals and cash-flow news are revealed over time, the other investor will eventually become more optimistic than he is, allowing him to sell the asset on at a gain.

Both in Scheinkman and Xiong (2003) and in our model, the increase in volume during a bubble is due to an increase in disagreement among investors. In Scheinkman and Xiong (2003), this increase in disagreement is exogenous. In our model, disagreement grows *endogenously* over the course of the bubble. As the bubble increases in size, the growth and value signals in equation (10) become very large in absolute magnitude. Extrapolators who, as a consequence of wavering, differ even very slightly in the relative weight they put on the two signals disagree sharply about the expected price change on the risky asset and therefore trade in large quantities. Whereas in Scheinkman and Xiong (2003) an exogeneous increase in disagreement leads to both higher volume and overpricing, in our model, the causation is different: overpricing leads to endogenously higher disagreement and hence higher volume.

Our model differs from disagreement models in other important ways. In our model, many investors hold expectations that depend positively on past returns, consistent with survey evidence on the expectations of actual investors. In Scheinkman and Xiong (2003), the holder of the asset has *constant* expectations about the asset's future return. Our framework also predicts a positive correlation between volume and past returns during bubble episodes, a prediction that we confirm empirically in the next section. Using simulations, we find that, in the model of Scheinkman and Xiong (2003), this correlation is close to zero: the exogeneous process that governs disagreement and hence volume is uncorrelated with the process for fundamentals that is the main determinant of price movements.

7 Empirical analysis

It is beyond the scope of this paper to attempt a full investigation of all the model's predictions. Nonetheless, we present some initial empirical evidence that may be helpful for evaluating two distinctive predictions of the model. One prediction, outlined in Section 4, is that the volume of trading in an overvalued asset is positively related to the asset's return over the previous year. In Section 7.1, we examine this prediction for four different bubble episodes. In Section 7.2, we evaluate another central prediction of our model: as a bubble develops, a larger fraction of its investor base will consist of investors with extrapolator-like characteristics.

7.1 Volume and past returns

We evaluate the prediction that, for an overvalued asset, trading volume is a positive function of its past return for four different bubble episodes: the stock market boom of 1928-1929, the technology bubble of 1998-2000, the 2004-2005 housing bubble, and the 2007-2008 commodity boom.

Stock market boom of 1928-1929

Accounts of the stock market boom of the late 1920s suggest that the bubble began in March 1928 (Allen 1931, Galbraith 1954, White 1990). White (1990) shows that frontier industries, especially utilities, led the stock market boom. Returns in these new-industry stocks far outpaced returns of stocks in older industries such as railroads. Panel A of Figure 7 confirms White's account. It compares the value-weighted cumulative return of public utilities listed in CRSP (SIC codes 4900-4990) with the cumulative return of the broader stock market. Utilities outperformed the broader stock market by more than 80% in the March 1928 - September 1929 period.¹⁷

¹⁷Wigmore (1985, p.42) describes the market for these stocks: "There is no gainsaying the enthusiasm of the financial markets for these public utility holding companies, however. [...] Their trading volume in 1929 exceeded 100% of their outstanding shares. At the high point in the market, their stocks averaged prices 57 times earnings per share, with Electric Bond and Share, which was most prominent because of its size and its relationship with General Electric, selling at 96 times earnings per share." Following the collapse

If we conjecture that utility stocks experienced a bubble in 1928-1929, our model predicts that the volume of trading in these stocks will be positively related to their past return. Panel B of Figure 7 plots the value-weighted monthly turnover of utility stocks over this period alongside their value-weighted 12-month past return; turnover is defined as volume divided by shares outstanding. The panel shows that, after a spike in April 1928, the turnover of utility stocks closely tracks their 12-month past return. For example, the second highest volume month in the series occurs in June 1929, following a 12-month cumulative return of 86%. The correlation between turnover and past returns between January 1927 and December 1932 is 0.67.

The technology bubble of 1998-2000

The explosion of volume during the technology bubble of 1998-2000 is well-known (Ofek and Richardson 2003, Hong and Stein 2007). In Panels C and D of Figure 7, we replicate and extend these findings. Panel C plots value-weighted monthly cumulative returns for the sample of .com stocks used by Ofek and Richardson (2003) and compares them to the cumulative returns of the CRSP value-weighted stock market index; returns for .com stocks are from CRSP. Technology stocks began their climb in December 1997 with a 12% value-weighted return. After a flat month in January 1998, they climbed another 23% in February, another 13% in March, and another 13% in April. Panel D shows that turnover increases steadily as the bubble progresses. Turnover (measured as before, and value-weighted) peaks in April 1999, the same point at which technology stocks reach their highest 12-month return of 429%. Overall, the figure shows that turnover closely tracks the 12-month return, with a time-series correlation of 0.71 between December 1997 and December 2002.

The 2004-2005 housing bubble

The relationship between turnover and past returns also appears during the U.S. housing boom and bust of the mid-2000s. In Panel E of Figure 7, we plot the Case-Shiller 20-City Composite Home Price Index. This index, based on repeat transactions, seeks to measure the value of residential real estate in the 20 largest U.S. metropolitan areas.

of the bubble, utilities were “held up by Franklin Roosevelt and others as spectacles of financial abuse and confusion.”

The Case-Shiller Index rises from a base value of 100 in January 2000 to a peak of 206.61 in April 2006. In Panel F, we show the relationship between 12-month past returns and volume for the U.S. housing market; we use existing home sales by month as a measure of volume.¹⁸ The figure shows that, as for the two stock market bubbles, volume clearly tracks 12-month past returns; the time-series correlation is 0.84.

The 2007-2008 commodity bubble

Whether the run-up in commodity prices in 2007 and 2008 can be easily explained by fundamentals or was instead a bubble is subject to some debate, with some authors suggesting that the “financialization” of derivatives markets instigated demand from institutional investors (Irwin and Sanders 2010, Cheng and Xiong 2014, Hong, de Paula, and Singh 2015). Panel G of Figure 7 shows the run-up in oil prices as reflected in the share price of USO, the largest exchange-traded fund with exposure to oil. USO more than doubled between December 2006 and June 2008. In Panel H of Figure 7, we plot the relationship between the monthly turnover of this ETF and its 12-month past return, both obtained from CRSP. As in our other examples, the turnover of USO closely tracks the past return; the time-series correlation between January 2001 and December 2010 is 0.83.

7.2 Composition of the investor base

In our model, as a bubble develops, extrapolators constitute a larger fraction of its investor base. We can broaden this prediction to say that, as a bubble develops, it draws in new extrapolators whose share demands depend positively on past returns. We here test two predictions of the model using data from the technology bubble of the late 1990s: 1) that new investors are drawn in as the bubble develops, and 2) that a significant fraction of these new investors are extrapolators. We focus on the technology bubble because extensive data on mutual fund ownership is available, even though mutual funds owned less than 25% of the typical .com stock (Griffin et al. 2011).

Our goal is to quantify changes in the ownership of .com stocks as the bubble progresses

¹⁸Existing-home sales are based on closing transactions of single-family homes, townhomes, condominiums, and cooperative homes and are provided by the National Association of Realtors.

in 1998 and 1999. To do so, we again identify .com stocks using the list of securities provided by Ofek and Richardson (2003). This list is then matched to quarterly mutual fund holdings using holdings data from Thomson Spectrum.

We start by analyzing how internet stocks draw in new investors as the bubble progresses. In each quarter t , we record the number of mutual fund owners initiating a position in a stock, in the sense that they held the stock in quarter t but not in quarter $t - 1$, and express this as a ratio to the total number of mutual fund owners in that quarter. We compute the equal-weighted average of this ratio across all technology stocks as the technology bubble progresses. One potential confounding factor in this calculation is the large amount of new issuance during the technology bubble. To avoid mechanically picking up “new” ownership for recently-listed securities, we include a security in our analysis only after it appears in the Thomson Spectrum dataset for at least two consecutive quarters with at least five institutional owners in each of these quarters.

Panel A of Figure 8 shows how the bubble draws in new investors. The figure shows the new-owner-share in each quarter of the bubble. As a benchmark for comparison, the figure also shows the equal-weighted new-owner-share across *non*-technology sector stocks. Figure 8 shows that technology stocks draw in new holders at a greater rate than non-technology stocks. Near the peak of the bubble, in December 1999, over 40% of the typical .com stock owners had not held any shares in the stock in the previous quarter. For example, according to our data, in December 1997, 48 mutual funds reported positions in Yahoo! stock. One year later, 171 mutual funds reported positions, while one year after that, 759 mutual funds reported positions.

Are the new owners of bubble stocks extrapolators, as our model predicts? To test this, we examine the other securities held by .com owners and check if these owners behave like extrapolators elsewhere in their portfolios. We sort all stocks—not only .com stocks—into NYSE deciles according to their past 12-month return. We then measure the “growthiness” of a mutual fund’s portfolio as the position-weighted past-return decile of the stocks in that portfolio. A portfolio with a score of 1, for example, contains exclusively stocks that have performed poorly over the past 12 months, while a portfolio with a score of 10 contains exclusively high past return stocks. Panel B of Figure 8 shows the increasing growthiness of

new owners as the bubble progresses. A similar result holds if we classify owners based on their degree of growthiness in December 1997, before the bubble starts.

Overall, the evidence is consistent with the broad prediction of our model that, as the bubble develops, it draws in new investors who resemble the extrapolators in our framework.

8 Conclusion

Although historical accounts of price bubbles typically emphasize extrapolative expectations (Kindleberger 1978, Shiller 2000), recent models of bubbles have moved away from this feature. In this paper, we embrace it. In our model, some investors hold extrapolative expectations, but also waver in their convictions in that they worry more or less about the possible overvaluation of the asset. The model generates occasional price bubbles in asset prices. Such bubbles occur in response to particular patterns of good news, a phenomenon Kindleberger (1978) called displacement. They are characterized by very high trading volume documented in earlier literature, which to a significant extent comes from the trading between the wavering extrapolators. The model generates a sharp new prediction that trading volume is driven by high past returns which distinguishes it from some popular recent models and appears to be consistent with some historical evidence.

Our analysis has left several important issues to future work. First, we have not addressed the controversy of whether bubbles actually exist, and whether investors can tell in the middle of a rapid price increase of an asset that it is actually overpriced. Second, even in the context of our model, we have assumed a very simple and stabilizing form of arbitrage. This specification does not consider the possibility of destabilizing arbitrage, whereby rational investors buy an overpriced asset in the hopes of selling at an even higher price to extrapolators (DeLong et al. 1990, Brunnermeier and Nagel 2004). But we also have not considered other stabilizing forces, such as arbitrage by security issuers themselves through greater issue or asset creation (Galbraith 1954, Glaeser and Nathanson 2015). Perhaps most important, we have simply adopted a standard but ad-hoc formulation of extrapolative beliefs by some investors. The fundamental psychological mechanisms of extrapolation remain to be understood.

9 Appendix

A micro-foundation for the fundamental trader demand in equation (3).

Consider an economy with the time structure and asset structure described at the start of Section 2. There are two types of trader: one type, which makes up a fraction μ^X of the population, has time t demand for shares of the risky asset given by N_t^X ; the other type, which makes up a fraction μ^F of the population with $\mu^F = 1 - \mu^X$, is a fundamental trader who, at time t , chooses his share demand N_t^F by maximizing a utility function with constant absolute risk aversion γ and defined over next period's wealth. In other words, his objective is

$$\max_{N_t^F} \mathbb{E}_t^F \left[-e^{-\gamma(W_t + N_t^F(\tilde{P}_{t+1} - P_t))} \right]. \quad (\text{A1})$$

This trader is boundedly rational, in a way that we make precise in what follows.

To determine his time t demand for the risky asset, the fundamental trader reasons as follows. At the final date, date T , the price of the risky asset P_T must equal the cash flow realized on that date, so that $P_T = D_T$. At time $T - 1$, the fundamental trader's first-order condition implies that his share demand is

$$N_{T-1}^F = \frac{\mathbb{E}_{T-1}^F(\tilde{P}_T) - P_{T-1}}{\gamma \text{Var}_{T-1}^F(\tilde{P}_T - P_{T-1})} = \frac{D_{T-1} - P_{T-1}}{\gamma \sigma_\varepsilon^2}, \quad (\text{A2})$$

where we have used the fact that $\mathbb{E}_{T-1}^F(\tilde{P}_T) = D_{T-1}$ and have also assumed, for simplicity, that the fundamental trader sets the conditional variance of price changes equal to the variance of cash-flow shocks. Market clearing implies

$$\mu^F \left(\frac{D_{T-1} - P_{T-1}}{\gamma \sigma_\varepsilon^2} \right) + \mu^X N_{T-1}^X = Q, \quad (\text{A3})$$

which, in turn, implies

$$P_{T-1} = D_{T-1} - \frac{\gamma \sigma_\varepsilon^2}{\mu^F} (Q - \mu^X N_{T-1}^X). \quad (\text{A4})$$

At time $T - 2$, the fundamental trader's demand is

$$N_{T-2}^F = \frac{\mathbb{E}_{T-2}^F(\tilde{P}_{T-1}) - P_{T-2}}{\gamma \sigma_\varepsilon^2}. \quad (\text{A5})$$

It is here that his bounded rationality comes into play. When computing $\mathbb{E}_{T-2}^F(\tilde{P}_{T-1})$, in other words, when computing the expectation of the quantity in (A4), he must come up with

an estimate of $\mathbb{E}_{T-2}^F(N_{T-1}^X)$. We assume that the fundamental trader does not try to forecast the evolution of the other traders' demand N^X , but instead sets $\mathbb{E}_{T-2}^F(N_{T-1}^X) = Q$; in other words, he assumes that the other traders will simply hold their share of the supply of the risky asset. Under this assumption, $\mathbb{E}_{T-2}^F(\tilde{P}_{T-1}) = D_{T-2} - \gamma\sigma_\varepsilon^2 Q$, so that

$$N_{T-2}^F = \frac{D_{T-2} - \gamma\sigma_\varepsilon^2 Q - P_{T-2}}{\gamma\sigma_\varepsilon^2}. \quad (\text{A6})$$

We assume that the fundamental trader continues in this way, working back from date T to the current time t , and, at each time, forecasting that the future demand from the other traders will simply equal Q . Under these assumptions,

$$N_t^F = \frac{D_t - \gamma\sigma_\varepsilon^2(T-t-1)Q - P_t}{\gamma\sigma_\varepsilon^2}, \quad (\text{A7})$$

which is equation (3).

A micro-foundation for the extrapolation-based demand in equation (5).

Consider an economy with the time structure and asset structure described at the start of Section 2. Now consider a trader who, at time t , maximizes a utility function with constant absolute risk aversion γ and defined over next period's wealth. In other words, his objective is

$$\max_{N_t^X} \mathbb{E}_t^X \left[-e^{-\gamma(W_t + N_t^X(\tilde{P}_{t+1} - P_t))} \right]. \quad (\text{A8})$$

From the first-order condition, optimal demand is

$$N_t^X = \frac{\mathbb{E}_t^X(\tilde{P}_{t+1}) - P_t}{\gamma \text{Var}_t(\tilde{P}_{t+1} - P_t)}. \quad (\text{A9})$$

Suppose that this investor forms beliefs about future price changes by extrapolating past price changes, so that

$$\mathbb{E}_t^X(\tilde{P}_{t+1} - P_t) = (1 - \theta) \sum_{k=1}^{\infty} \theta^{k-1} (P_{t-k} - P_{t-k-1}) \equiv X_t, \quad (\text{A10})$$

which, for an economy that starts at time 0, can be written as

$$\mathbb{E}_t^X(\tilde{P}_{t+1} - P_t) = (1 - \theta) \sum_{k=1}^{t-1} \theta^{k-1} (P_{t-k} - P_{t-k-1}) + \theta^{t-1} X_1. \quad (\text{A11})$$

Suppose also, for simplicity, that he sets the conditional variance of prices equal to the variance of cash-flow shocks, namely σ_ε^2 . His demand function then becomes

$$N_t^X = \frac{1}{\gamma\sigma_\varepsilon^2} \left(\sum_{k=1}^{t-1} \theta^{k-1} (P_{t-k} - P_{t-k-1}) + \theta^{t-1} X_1 \right), \quad (\text{A12})$$

as in (5).

Proof of Proposition 1. From expressions (9) and (10), we see that aggregate demand, $\mu_0 N_t^F + \sum_{i=1}^I \mu_i N_t^{E,i}$, can take an arbitrarily high value if the price P_t is sufficiently low, and a value as low as zero if the price is sufficiently high. Moreover, it is a continuous function of P_t and is strictly decreasing in P_t until it falls to zero. Taken together, these observations imply that there is a unique price P_t at which aggregate demand at time t equals the market supply of Q .

We find the market-clearing price in the following way. As noted in the statement of the proposition, we define \bar{P}_i to be the price at which trader i 's short-sale constraint binds, namely

$$\begin{aligned} \bar{P}_0 &= D_t - \gamma\sigma_\varepsilon^2(T-t-1)Q \\ \bar{P}_i &= D_t - \gamma\sigma_\varepsilon^2(T-t-1)Q + \frac{1-w_{i,t}}{w_{i,t}} X_t, \quad i \in \{1, \dots, I\}. \end{aligned} \quad (\text{A13})$$

We now order these $I+1$ ‘‘cut-off’’ prices, so that

$$\bar{P}_{i(0)} \geq \bar{P}_{i(1)} \geq \dots \geq \bar{P}_{i(I)},$$

where $i(l)$ indexes the trader $i \in \{0, 1, \dots, I\}$ with the $(l+1)$ 'th highest cut-off price. If $N_{\bar{P}_{i(l)}}$ is aggregate demand at price $\bar{P}_{i(l)}$, we have

$$0 = N_{\bar{P}_{i(0)}} \leq N_{\bar{P}_{i(1)}} \leq \dots \leq N_{\bar{P}_{i(I)}}.$$

Finally, let $I(l)$ be the set of traders i who have strictly positive demand at price $\bar{P}_{i(l)}$. Note that $I(0)$ is an empty set, and $I(l)$ is a subset of $I(l+1)$.

We consider two cases. Suppose that $N_{\bar{P}_{i(I)}} < Q$. This indicates that the market-clearing price is below $\bar{P}_{i(I)}$, and that, in equilibrium, all traders in the economy will have strictly positive demand. Aggregate demand at the market-clearing price P_t will therefore equal

$$\sum_{i=0}^I \mu_i \left[w_{i,t} \left(\frac{D_t - \gamma\sigma_\varepsilon^2(T-t-1)Q - P_t}{\gamma\sigma_\varepsilon^2} \right) + (1-w_{i,t}) \frac{X_t}{\gamma\sigma_\varepsilon^2} \right],$$

where $w_{0,t} \equiv 1$, indicating that fundamental traders put 100% weight on the value signal. Setting this aggregate demand equal to the market supply Q leads to the equilibrium price in (11).

We now turn to the other case. Suppose that $N_{\bar{P}_{i(l)}} \leq Q \leq N_{\bar{P}_{i(l+1)}}$. We then know that the market-clearing price is somewhere between $\bar{P}_{i(l+1)}$ and $\bar{P}_{i(l)}$, and that, in equilibrium, only the traders in the set $I(l+1)$, denoted I^* in the statement of the proposition, will have strictly positive demand for the risky asset. Aggregate demand at the market-clearing price P_t will therefore equal

$$\sum_{i \in I(l+1)} \mu_i \left[w_{i,t} \left(\frac{D_t - \gamma \sigma_\varepsilon^2 (T - t - 1) Q - P_t}{\gamma \sigma_\varepsilon^2} \right) + (1 - w_{i,t}) \frac{X_t}{\gamma \sigma_\varepsilon^2} \right].$$

Setting this equal to the market supply of Q , we obtain the equilibrium price in (12). \blacksquare

Proof of Proposition 2. For tractability, we replace the discrete types of extrapolators by a continuum of extrapolators. At time t , each of these extrapolators draws an independent weight $w_{i,t}$ from a bounded and continuous density function $g(w)$, $w \in [w_l, w_h]$, with mean \bar{w} and with $0 < w_l < w_h < 1$. Given these assumptions and the results from Proposition 1, the equilibrium price of the risky asset is

$$P_t = D_t + \alpha_1 \sum_{k=1}^{\infty} \theta^{k-1} (P_{t-k} - P_{t-k-1}) - (T - t - 1) \gamma \sigma_\varepsilon^2 Q - \frac{\gamma \sigma_\varepsilon^2 Q}{\mu_0 + (1 - \mu_0) \bar{w}} \quad (\text{A14})$$

in the first stage of the bubble, where $\alpha_1 \equiv \frac{(1-\theta)(1-\mu_0)(1-\bar{w})}{\mu_0 + (1-\mu_0)\bar{w}}$. In the second stage of the bubble, so long as all the extrapolators are in the market, the equilibrium price is

$$P_t = D_t + \alpha_2 \sum_{k=1}^{\infty} \theta^{k-1} (P_{t-k} - P_{t-k-1}) - (T - t - 1) \gamma \sigma_\varepsilon^2 Q - \frac{\gamma \sigma_\varepsilon^2 Q}{(1 - \mu_0) \bar{w}}, \quad (\text{A15})$$

where $\alpha_2 \equiv \frac{(1-\theta)(1-\bar{w})}{\bar{w}} > \alpha_1$.

From (4), (A14), and (A15), the level of overpricing \mathcal{O}_t , defined as the difference between the price of the risky asset and its fundamental value, is

$$\mathcal{O}_t = \begin{cases} \frac{1 - \bar{w}}{\bar{w}} X_t - \frac{\gamma \sigma_\varepsilon^2 Q}{(1 - \mu_0) \bar{w}} + \gamma \sigma_\varepsilon^2 Q & X_t > \frac{\gamma \sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \\ \frac{(1 - \mu_0)(1 - \bar{w})}{\mu_0 + (1 - \mu_0) \bar{w}} X_t - \frac{\gamma \sigma_\varepsilon^2 Q}{\mu_0 + (1 - \mu_0) \bar{w}} + \gamma \sigma_\varepsilon^2 Q & X_t \leq \frac{\gamma \sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \end{cases}. \quad (\text{A16})$$

Note that \mathcal{O}_t is continuous at the switching point between the first and second stages of the bubble; at this point, $\mathcal{O}_t = \gamma\sigma_\varepsilon^2 Q$. Also note that, when the value of X_t equals its steady-state level of $\gamma\sigma_\varepsilon^2 Q$, the overpricing is zero; in this case, both extrapolators and fundamental traders demand the per-capita supply Q of the risky asset.

From (A14) and (A15) it is apparent that, if the economy stays within stage one or within stage two and if all the extrapolators are in the market, the model has a linear structure: in stage i , a fundamental shock of ε_{t_1} at t_1 and a fundamental shock of ε_{t_2} at t_2 generate, at a later time t , a total overvaluation of $\mathcal{L}_i(t - t_1)\varepsilon_{t_1} + \mathcal{L}_i(t - t_2)\varepsilon_{t_2}$. It is also straightforward to check that $\mathcal{L}_1(\cdot)$ and $\mathcal{L}_2(\cdot)$ can be defined recursively as

$$\mathcal{L}_i(0) = 0, \quad \mathcal{L}_i(1) = \alpha_i, \quad \mathcal{L}_i(l) = (\alpha_i + \theta)\mathcal{L}_i(l - 1) - \alpha_i\mathcal{L}_i(l - 2), \quad l \geq 2 \text{ for } i = 1, 2. \quad (\text{A17})$$

This is a standard difference equation with the following general solution

$$\mathcal{L}_i(j) = A_1(K_1)^j + A_2(K_2)^j, \quad (\text{A18})$$

where K_1 and K_2 are the two roots of the quadratic equation

$$K^2 - (\alpha_i + \theta)K + \alpha_i = 0 \quad (\text{A19})$$

and where A_1 and A_2 can be obtained from the boundary conditions $\mathcal{L}_i(0) = 0$ and $\mathcal{L}_i(1) = \alpha_i$. When $(\alpha_i + \theta)^2 > 4\alpha_i$, (A19) has two real roots; matching (A18) with the two boundary conditions gives (16). When $(\alpha_i + \theta)^2 < 4\alpha_i$, (A19) has two complex roots with non-zero imaginary components; matching (A18) with the two boundary conditions gives (17). When $(\alpha_i + \theta)^2 = 4\alpha_i$, applying L'Hôpital's rule to either (16) or (17) gives (18).

The linear structure implies that, at time t with $l \leq t < j$, the overpricing caused by $\{\varepsilon_i\}_{i=l}^{t-1}$ is $\sum_{m=l}^{t-1} \mathcal{L}_1(t - m)\varepsilon_m$; and that, at time t with $j \leq t \leq N$, the additional overpricing caused by $\{\varepsilon_i\}_{i=j}^{t-1}$ is $\sum_{m=j}^{t-1} \mathcal{L}_2(t - m)\varepsilon_m$.

We now derive \mathcal{O}_t^1 at time $t \geq j$. For $t = j$

$$X_j = (1 - \theta)(P_{j-1} - P_{j-2}) + \theta X_{j-1}. \quad (\text{A20})$$

From (14) we know

$$P_{j-1} - P_{j-2} = \mathcal{O}_{j-1} - \mathcal{O}_{j-2} + \varepsilon_{j-1} + \gamma\sigma_\varepsilon^2 Q \quad (\text{A21})$$

and

$$X_{j-1} = \frac{\mu_0 + (1 - \mu_0)\bar{w}}{(1 - \mu_0)(1 - \bar{w})} \left(\mathcal{O}_{j-1} + \frac{\gamma\sigma_\varepsilon^2 Q}{\mu_0 + (1 - \mu_0)\bar{w}} - \gamma\sigma_\varepsilon^2 Q \right). \quad (\text{A22})$$

Substituting (A21) and (A22) into (A20), and then substituting (A20) back into (A16) gives \mathcal{O}_j^1 in (15). For $j < t \leq N$, similar steps lead to \mathcal{O}_t^1 in (15).

Substituting the price equation (A15) into the extrapolator share demand of extrapolator in (6) shows that, whenever $X_t > w_h \gamma \sigma_\varepsilon^2 Q / [(1 - \mu_0)(w_h - \bar{w})]$, the extrapolator with $w_{i,t} = w_h$ exits the market and hence a price spiral occurs. (A16) shows that this condition is equivalent to $\mathcal{O}_t > \bar{O}$. Applying (A16) at time $j' - 1$ gives $X_{j'-1}$ as a function of $\mathcal{O}_{j'-1}$, and further applying (A20) and (A21) gives (22).

Assume that, at time t , extrapolators with $w_{i,t} \in [w_l, w^*)$ are in the market. Integrating the share demands of these extrapolators in (10) and equating the result to the aggregate per-extrapolator supply of $Q/(1 - \mu_0)$ gives (19) and (20). Setting the share demand of the extrapolator with $w_{i,t} = w^*$ to zero then gives (21). Given that $X_t > w_h \gamma \sigma_\varepsilon^2 Q / [(1 - \mu_0)(w_h - \bar{w})]$, the left-hand side of (21) is smaller than the right-hand side when $w^* = w_h$; however the left-hand side of (21) is greater than the right-hand side when $w^* = w_l$. As a result, there must exist a w^* that solves (21). \blacksquare

Proof of Proposition 3. Substituting the equilibrium asset price in (A14) and (A15) into extrapolator i 's share demand gives

$$N_t^{E,i} = \begin{cases} \max \left(\frac{\bar{w} - w_{i,t}}{\bar{w}\gamma\sigma_\varepsilon^2} X_t + \frac{w_{i,t}Q}{(1 - \mu_0)\bar{w}}, 0 \right) & X_t > \frac{\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \\ \max \left(\frac{\mu_0(1 - w_{i,t}) + (1 - \mu_0)(\bar{w} - w_{i,t})}{(\mu_0 + (1 - \mu_0)\bar{w})\gamma\sigma_\varepsilon^2} X_t + \frac{w_{i,t}Q}{\mu_0 + (1 - \mu_0)\bar{w}}, 0 \right) & X_t \leq \frac{\gamma\sigma_\varepsilon^2 Q}{(1 - \mu_0)(1 - \bar{w})} \end{cases}. \quad (\text{A23})$$

From (A23) we know that, when $-w_l \gamma \sigma_\varepsilon^2 Q / [\mu_0(1 - w_l) + (1 - \mu_0)(\bar{w} - w_l)] \leq X_t < \gamma \sigma_\varepsilon^2 Q / [(1 - \mu_0)(1 - \bar{w})]$, both fundamental traders and extrapolators stay in the market and the change in share demand for extrapolator i , holding the level of the extrapolation signal fixed, is

$$\frac{(w_{i,t+1} - w_{i,t})(\gamma\sigma_\varepsilon^2 Q - X_t)}{(\mu_0 + (1 - \mu_0)\bar{w})\gamma\sigma_\varepsilon^2}. \quad (\text{A24})$$

Taking the absolute value of this quantity and integrating over $w_{i,t+1}$ and $w_{i,t}$ gives the wavering-induced trading volume. A similar derivation covers the case when $\gamma \sigma_\varepsilon^2 Q / [(1 - \mu_0)(1 - \bar{w})] \leq X_t \leq w_h \gamma \sigma_\varepsilon^2 Q / [(w_h - \bar{w})(1 - \mu_0)]$.

When $X_t > w_h \gamma \sigma_\varepsilon^2 Q / [(w_h - \bar{w})(1 - \mu_0)]$, extrapolators with a sufficiently high level of w stay out of the market but may re-enter in the next period. For those extrapolators who stay in for both periods, replace \bar{w} , $1 - \mu_0$, and $\bar{\Delta}_0$ by $\bar{w}(X_t)$, $(1 - \mu_0)\eta(X_t)$, and $\bar{\Delta}(X_t)$, respectively. For those extrapolators who are in at time t but out at time $t + 1$, their change in share demand is

$$\frac{X_t}{\gamma \sigma_\varepsilon^2} - \frac{w_{i,t}((1 - \mu_0)\eta(X_t)X_t - \gamma \sigma_\varepsilon^2 Q)}{(1 - \mu_0)\eta(X_t)\gamma \sigma_\varepsilon^2 \bar{w}(X_t)} \geq 0 \quad (\text{A25})$$

for $w_{i,t} \leq w_\eta(X_t)$. Integrating (A25) over $w_{i,t}$ from w_l to $w_\eta(X_t)$ and then further integrating it over $w_{i,t+1}$ from $w_\eta(X_t)$ to w_h gives $((1 - \eta(X_t))Q)/(1 - \mu_0)$. The trading volume generated by extrapolators who are out at time t but in at time $t + 1$ can be computed in a similar way; it also equals $((1 - \eta(X_t))Q)/(1 - \mu_0)$. ■

10 References

- Allen, F. (1931), *Only Yesterday: An Informal History of the 1920's*. Harper and Brothers, New York.
- Bagehot, W. (1873), *Lombard Street: A Description of the Money Market*. Wiley, New York.
- Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2015), “X-CAPM: An Extrapolative Capital Asset Pricing Model,” *Journal of Financial Economics*, 115, 1–24.
- Barberis, N., and A. Shleifer (2003), “Style Investing,” *Journal of Financial Economics* 68, 161–199.
- Blanchard, O., and M. Watson (1982), “Bubbles, Rational Expectations, and Financial Markets,” NBER Working Paper 945.
- Brunnermeier, M., and S. Nagel (2004), “Hedge Funds and the Technology Bubble,” *Journal of Finance* 59, 2013–2040.
- Brunnermeier, M., and M. Oehmke (2013), “Bubbles, Financial Crises, and Systemic Risk,” in George Constantinides, Milton Harris, and Rene Stulz, eds., *Handbook of the Economics of Finance*, Volume 2, North Holland.
- Carlos, A., L. Neal, and K. Wandschneider (2006), “Dissecting the Anatomy of Exchange Alley: The Dealings of Stockjobbers during and after the South Sea Bubble,” Working paper, University of Illinois.
- Case, K., R. Shiller, and A. Thompson (2012), “What Have They Been Thinking? Home-buyer Behavior in Hot and Cold Markets,” Brookings Papers on Economic Activity.
- Cheng, I.-H., and W. Xiong (2014), “The Financialization of Commodity Markets,” *Annual Review of Financial Economics*, forthcoming.
- Cutler, D., J. Poterba, and L. Summers (1990), “Speculative Dynamics and the Role of Feedback Traders,” *American Economic Review Papers and Proceedings* 80, 63–68.
- De Long, B., A. Shleifer, L. Summers, and R. Waldmann (1990), “Positive Feedback Investment Strategies and Destabilizing Rational Speculation,” *Journal of Finance* 45, 379–395.
- Fama, E. (2014), “Two Pillars of Asset Pricing,” *American Economic Review* 104, 1467–1485.

- Galbraith, J. (1954), *The Great Crash: 1929*. Houghton Mifflin Co., Boston, MA.
- Glaeser, E., and C. Nathanson (2015), “An Extrapolative Model of House Price Dynamics,” NBER Working Paper 21037.
- Greenwood, R., and A. Shleifer (2014), “Expectations of Returns and Expected Returns,” *Review of Financial Studies* 27, 714–746.
- Griffin, J., J. Harris, T. Shu, and S. Topaloglu (2011), “Who Drove and Burst the Tech Bubble?,” *Journal of Finance* 66, 1251–1290.
- Griffin, J., F. Nardari, and R. Stulz (2007), “Do Investors Trade More when Stocks Have Performed Well? Evidence from 46 Countries,” *Review of Financial Studies* 20, 905–951.
- Harrison, M., and D. Kreps (1978), “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations,” *Quarterly Journal of Economics* 89, 323–336.
- Haruvy, E., Y. Lahav, and C. Noussair (2007), “Traders’ Expectations in Asset Markets: Experimental Evidence,” *American Economic Review* 97, 1901–1920.
- Hong, H., A. de Paula, and V. Singh (2015), “Hoard Behavior and Commodity Bubbles,” NBER Working Paper 20974.
- Hong, H., and D. Sraer (2013), “Quiet Bubbles,” *Journal of Financial Economics* 110, 596–606.
- Hong, H., and J. Stein (1999), “A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets,” *Journal of Finance* 54, 2143–2184.
- Hong, H., and J. Stein (2007), “Disagreement and the Stock Market,” *Journal of Economic Perspectives* 21, 109–128.
- Irwin, S., and D. Sanders (2010), “The Impact of Index and Swap Funds on Commodity Futures Markets: Preliminary Results,” OECD Food, Agriculture and Fisheries Working Paper.
- Jin, L. (2015), “A Speculative Asset Pricing Model of Financial Instability,” Working paper, California Institute of Technology.
- Kindleberger, C. (1978), *Manias, Panics, and Crashes: A History of Financial Crises*. Basic Books.

- Mackay, C. (1841), *Extraordinary Popular Delusions and the Madness of Crowds*. Harmony Books, New York.
- Ofek, E., and M. Richardson (2003), “DotCom Mania: The Rise and Fall of Internet Stock Prices,” *Journal of Finance* 58, 1113–1138.
- Pastor, L., and P. Veronesi (2006), “Was there a Nasdaq Bubble in the 1990s?” *Journal of Financial Economics* 81, 61–100.
- Scheinkman, J., and W. Xiong (2003), “Overconfidence and Speculative Bubbles,” *Journal of Political Economy* 111, 1183–1219.
- Shiller, R. (2000), *Irrational Exuberance*. Princeton University Press, New Jersey.
- Smith, V., G. Suchanek, and A. Williams (1988), “Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets,” *Econometrica* 56, 1119–1151.
- Statman, M., S. Thorley, and K. Vorkink (2006), “Investor Overconfidence and Trading Volume,” *Review of Financial Studies* 19, 1531–1565.
- Tirole, J. (1985), “Asset Bubbles and Overlapping Generations,” *Econometrica* 53, 1071–1100.
- White, E. (1990), “The Stock Market Boom and Crash of 1929 Revisited,” *Journal of Economic Perspectives* 4, 67–83.
- Wigmore, B. (1985), *The Crash and Its Aftermath: A History of Securities Markets in the United States, 1923-1933*. Greenwood Press, Westport CT.
- Xiong, W. (2013), “Bubbles, Crises, and Heterogeneous Beliefs,” in Jean-Pierre Fouque and Joe Langsam, eds., *Handbook for Systemic Risk*, Cambridge University Press.
- Xiong, W., and J. Yu (2011), “The Chinese Warrants Bubble,” *American Economic Review* 101, 2723–2753.

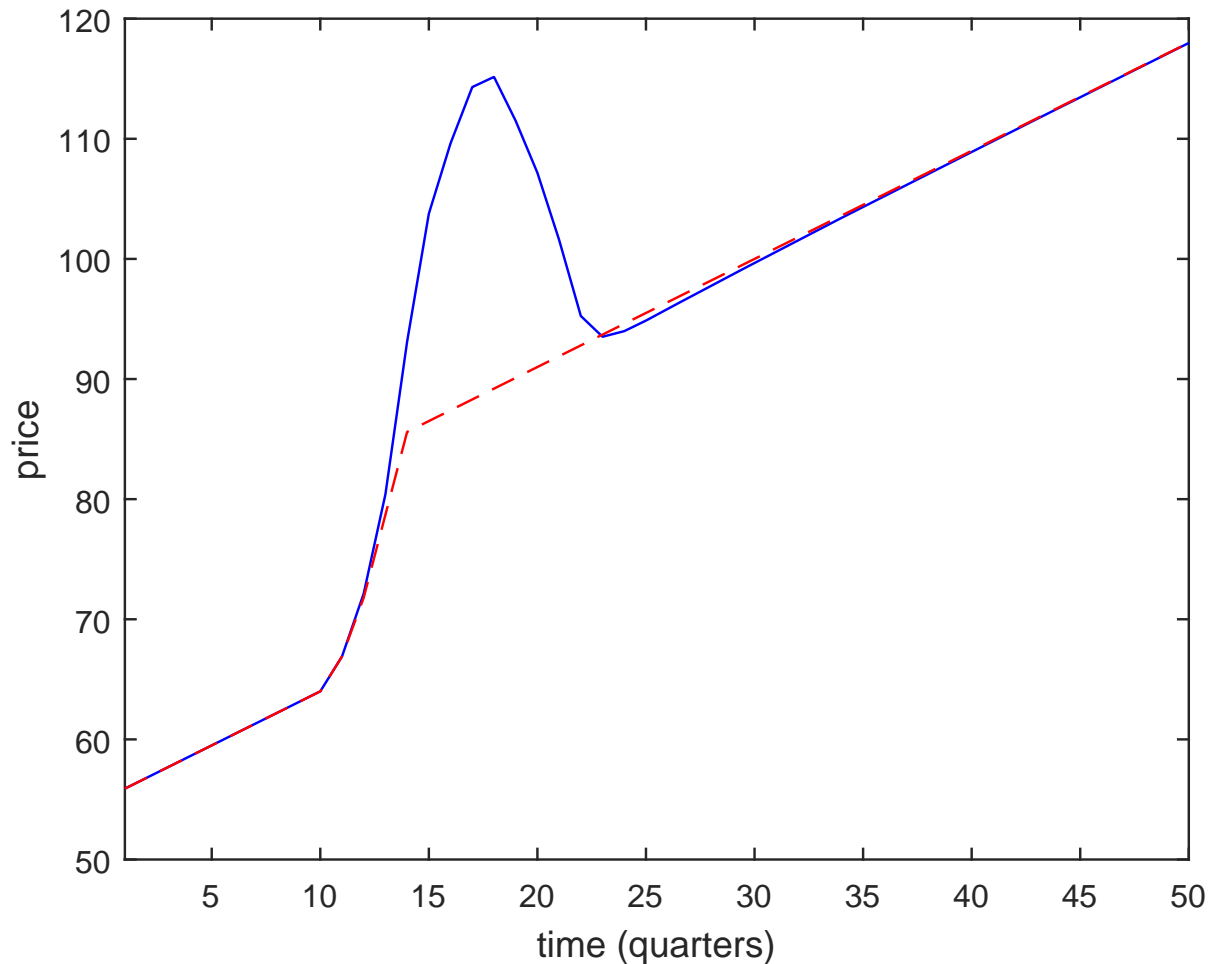


Figure 1: Prices in a bubble. The solid line plots the price of the risky asset for the following sequence of 50 cash-flow shocks: 10 shocks of zero, followed by shocks of 2, 4, 6, 6, followed by 36 shocks of zero. 30% of the investors are fundamental traders; the remainder are extrapolators with an extrapolation parameter θ of 0.9 and who also put a base weight of $\bar{w}_i = 0.1$ on a value signal. The dashed line plots the fundamental value of the asset for the same cash-flow sequence. The other parameters are $D_0 = 100$, $\sigma_\varepsilon = 3$, $Q = 1$, $\gamma = 0.1$, $\sigma_u = 0.3$, and $I = 50$.

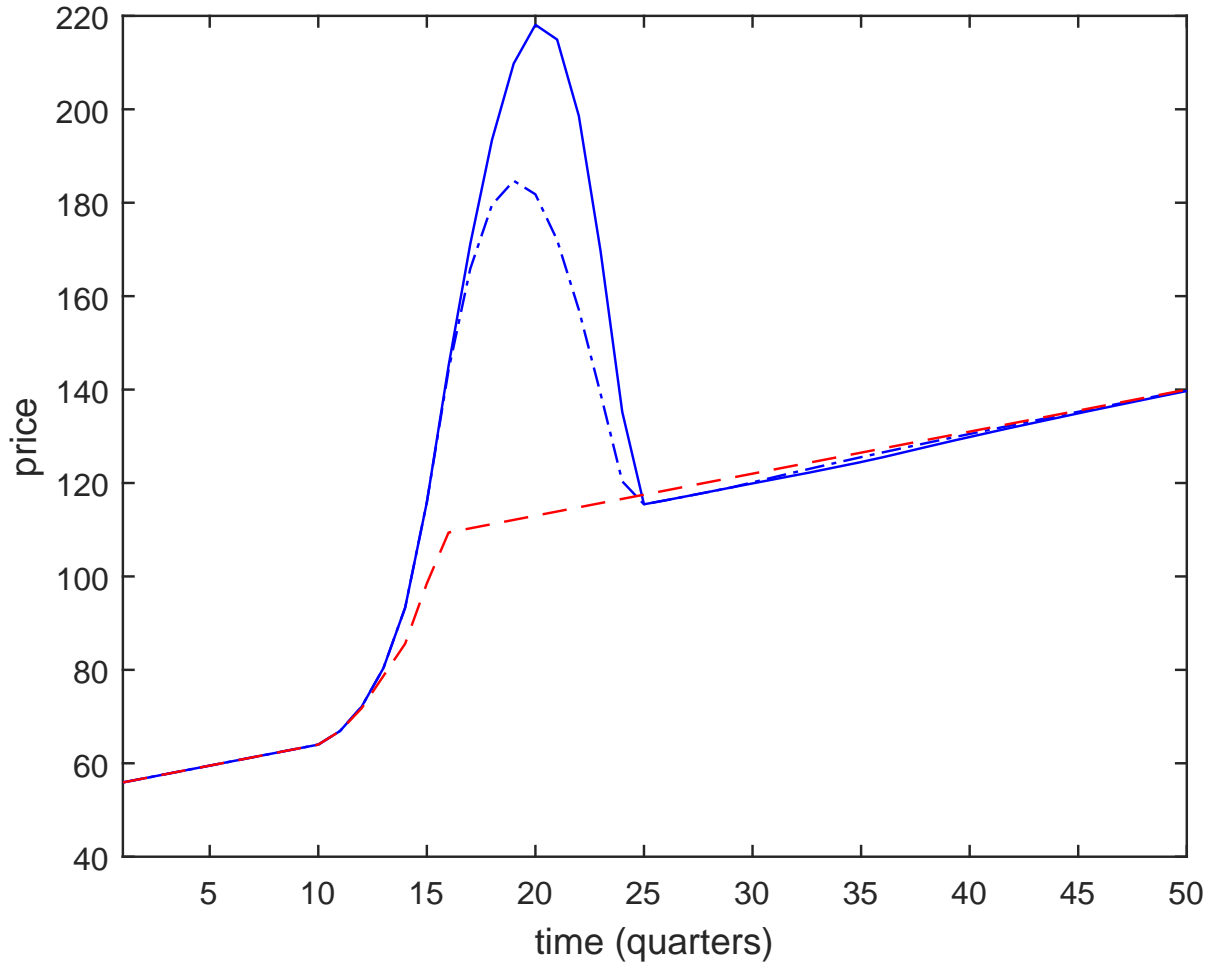


Figure 2: Price spiral. The solid line plots the price of the risky asset for the following sequence of 50 cash-flow shocks: 10 shocks of zero, followed by shocks of 2, 4, 6, 6, 12, 10, followed by 34 shocks of zero. 30% of the investors are fundamental traders; the remainder are extrapolators with an extrapolation parameter θ of 0.9 and who also put a base weight of $\bar{w}_i = 0.1$ on a value signal. The dashed line plots the fundamental value of the asset for the same cash-flow sequence. The dash-dot line plots the price in an economy where the extrapolators are homogeneous, placing the same, invariant weight of $\bar{w}_i = 0.1$ on the value signal. The other parameters are $D_0 = 100$, $\sigma_\varepsilon = 3$, $Q = 1$, $\gamma = 0.1$, $\sigma_u = 0.3$, and $I = 50$.

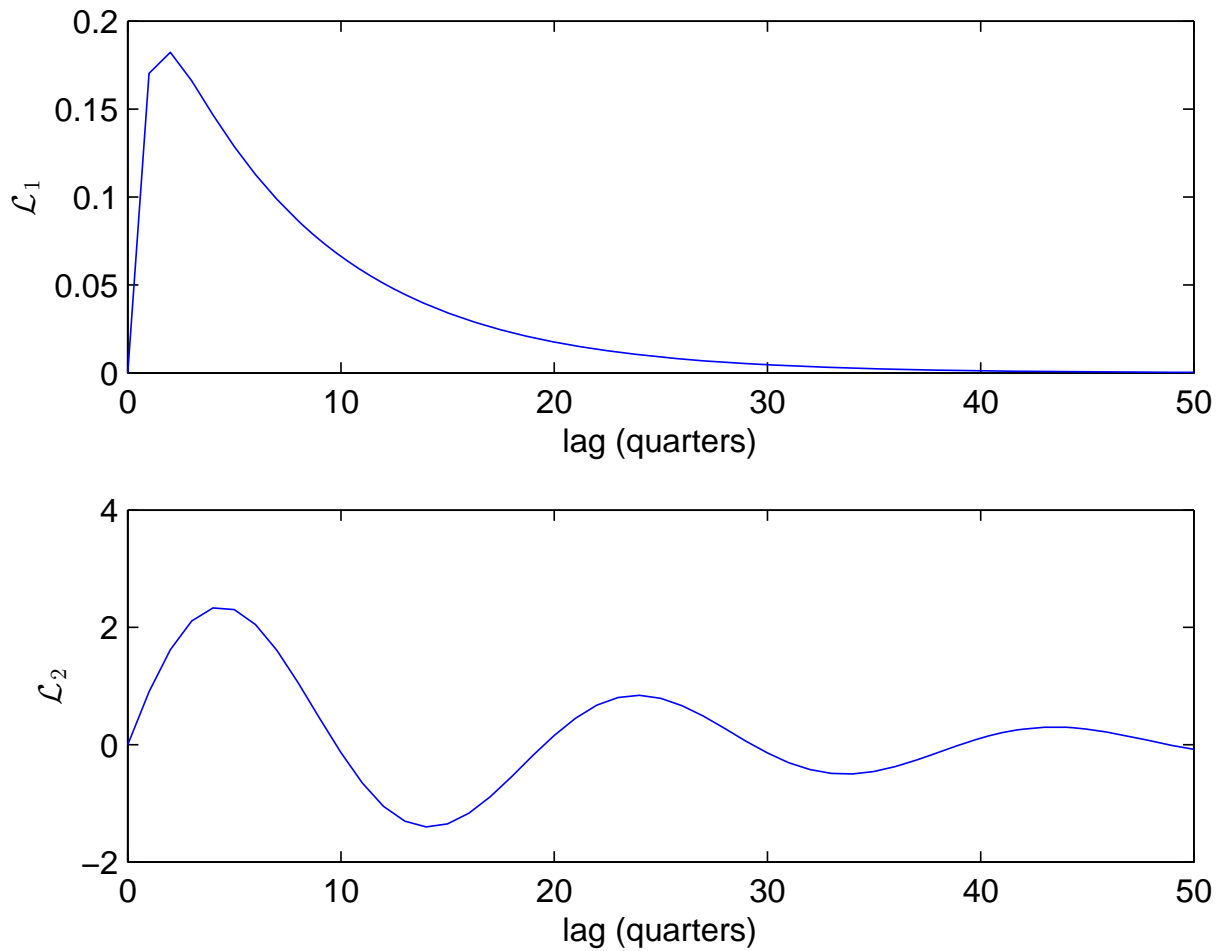


Figure 3: Impulse response for overvaluation. The top panel plots $\{\mathcal{L}_1(j)\}_{j \geq 0}$ for lag $j = 0, 1, \dots, 50$. The bottom panel plots $\{\mathcal{L}_2(j)\}_{j \geq 0}$ for lag $j = 0, 1, \dots, 50$. The parameters are $\theta = 0.9$, $\mu_0 = 0.3$, and $\bar{w} = 0.1$.

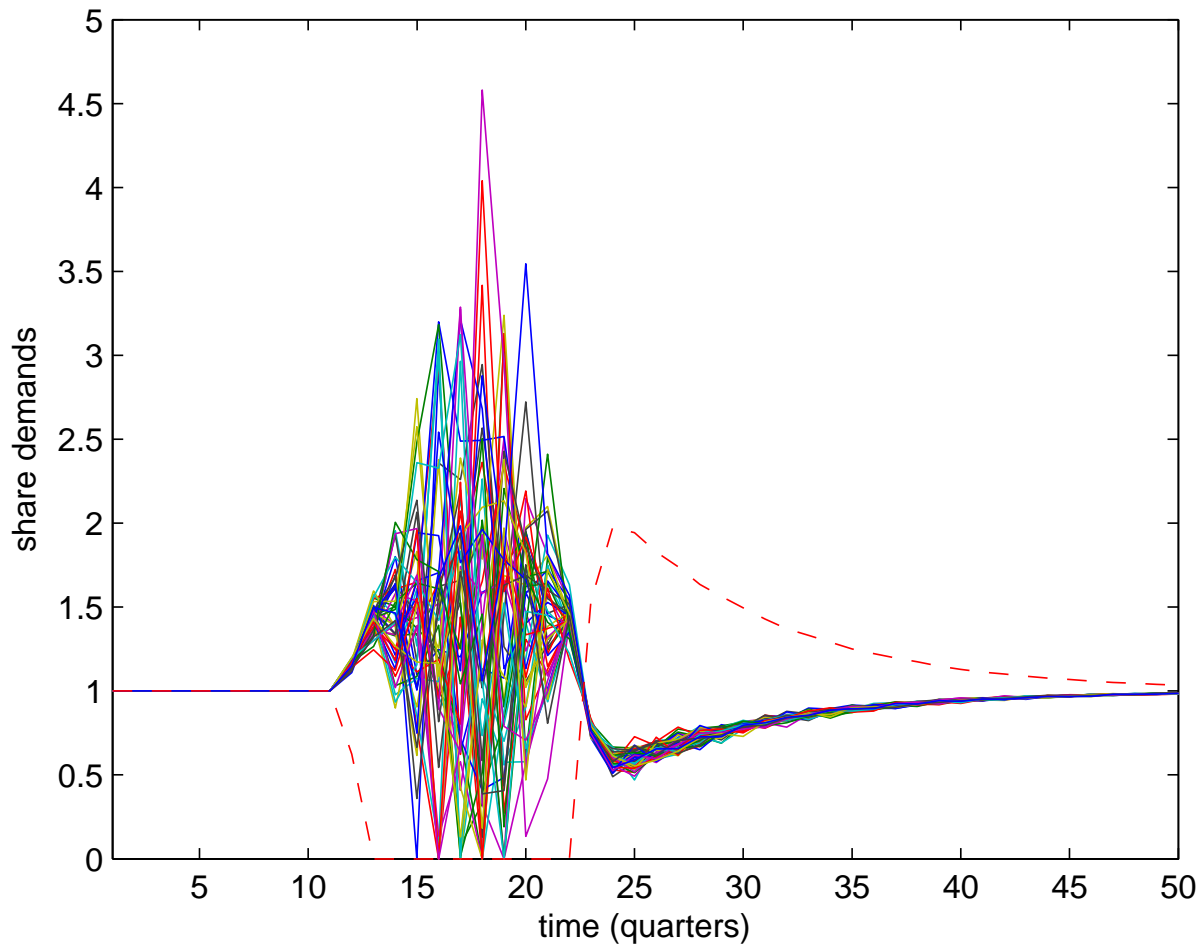


Figure 4: Share demands in a bubble. The solid lines plot the risky asset share demands of extrapolators for the following sequence of 50 cash-flow shocks: 10 shocks of zero, followed by shocks of 2, 4, 6, 6, followed by 36 shocks of zero. The dashed line plots the share demand of the fundamental traders for the same cash-flow sequence. 30% of the investors are fundamental traders; the remainder are extrapolators with an extrapolation parameter θ of 0.9 and who also put a base weight of $\bar{w}_i = 0.1$ on a value signal. The other parameters are $D_0 = 100$, $\sigma_\varepsilon = 3$, $Q = 1$, $\gamma = 0.1$, $\sigma_u = 0.3$, and $I = 50$.

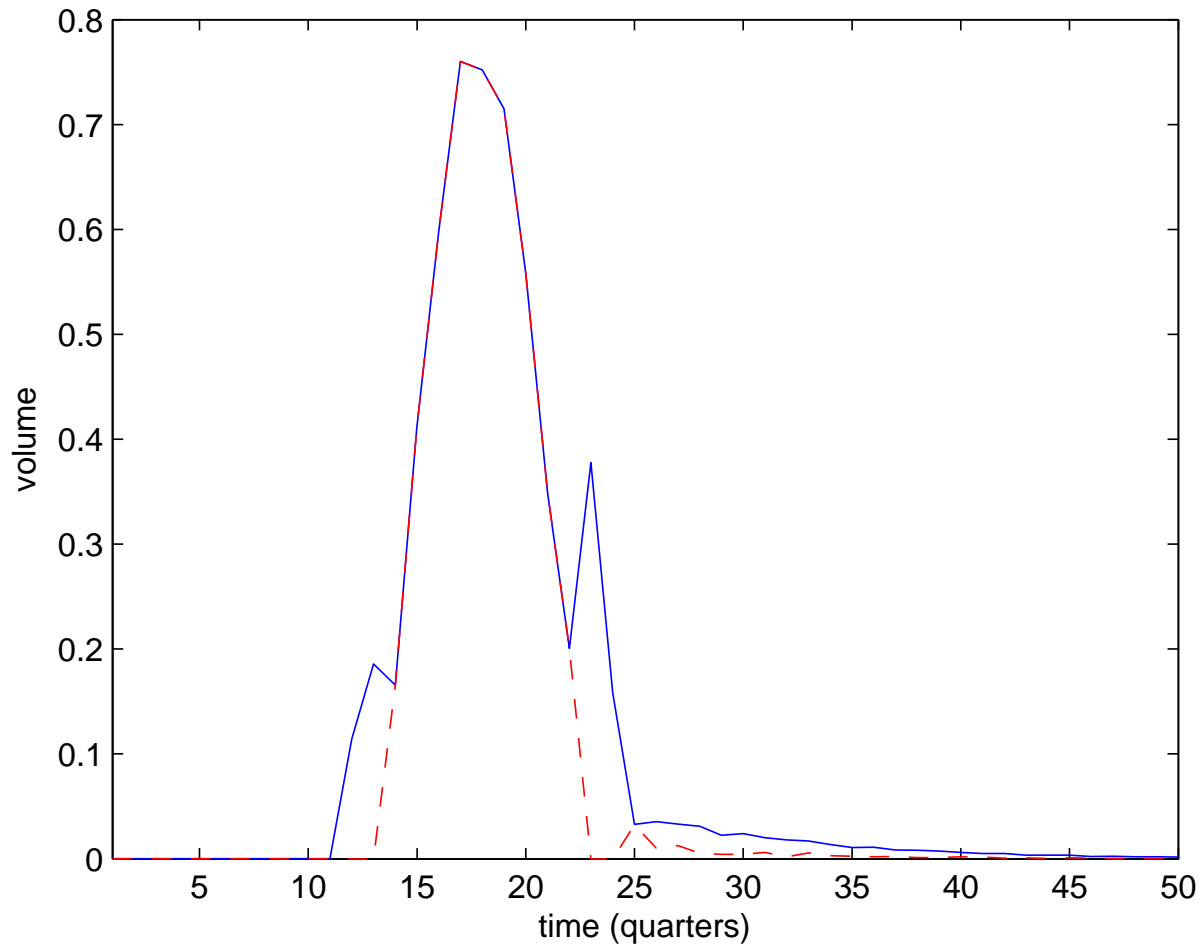


Figure 5: Volume in a bubble. The solid line plots the total trading volume in the risky asset for the following sequence of 50 cash-flow shocks: 10 shocks of zero, followed by shocks of 2, 4, 6, 6, followed by 36 shocks of zero. The dashed line plots the trading volume between the extrapolators for the same cash-flow sequence. 30% of the investors are fundamental traders; the remainder are extrapolators with an extrapolation parameter θ of 0.9 and who also put a base weight of $\bar{w}_i = 0.1$ on a value signal. The other parameters are $D_0 = 100$, $\sigma_\varepsilon = 3$, $Q = 1$, $\gamma = 0.1$, $\sigma_u = 0.3$, and $I = 50$.

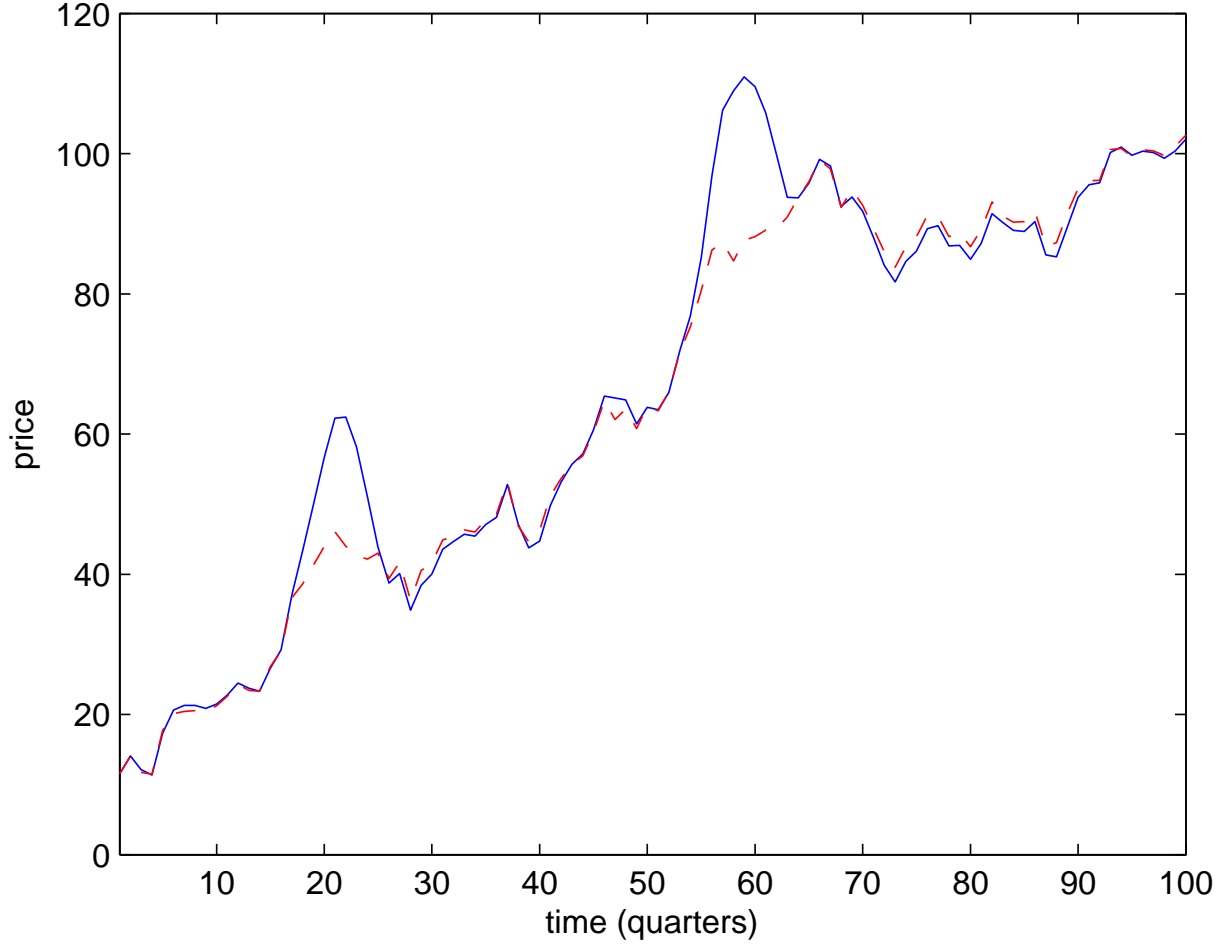


Figure 6: Price and fundamental value over time. The solid line plots the price of the risky asset for a 100-period cash-flow path that is simulated from the cash-flow process. 30% of the investors are fundamental traders; the remainder are extrapolators with an extrapolation parameter θ of 0.9 and who also put a base weight of $\bar{w}_i = 0.1$ on a value signal. The dashed line plots the fundamental value of the asset for the same cash-flow sequence. The other parameters are $D_0 = 100$, $\sigma_\varepsilon = 3$, $Q = 1$, $\gamma = 0.1$, $\sigma_u = 0.3$, and $I = 50$.

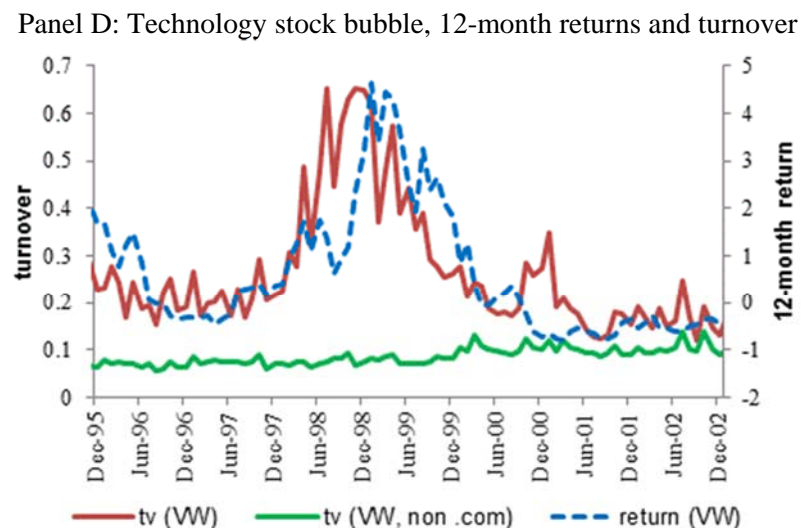
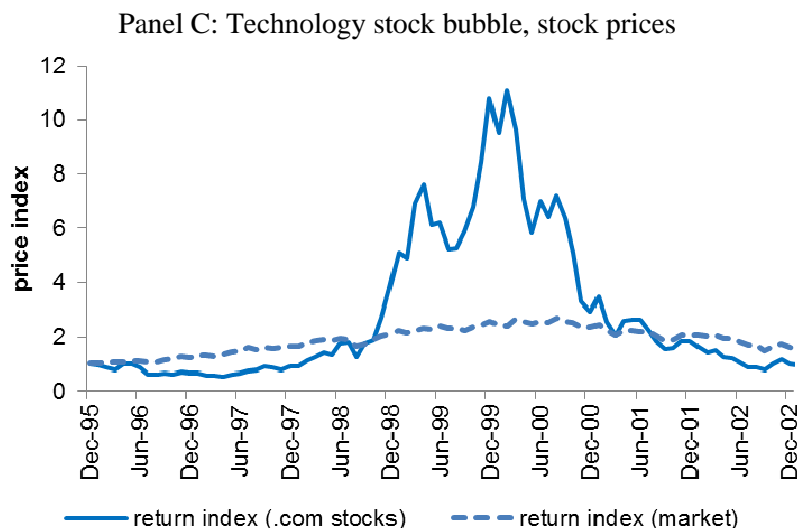
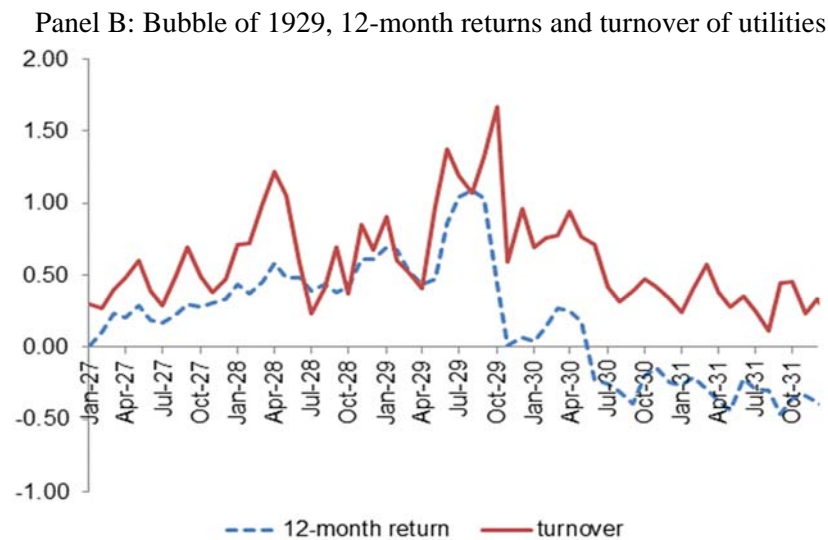
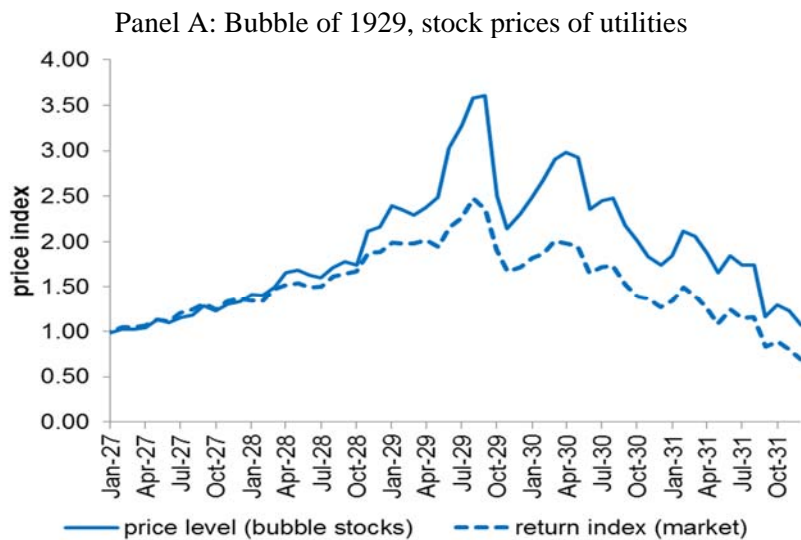


Figure 7: Prices, returns, and volume during bubble episodes. Prices, past 12-month returns, and value-weighted turnover during four bubbles episodes: utility stocks in 1929; technology stocks 1995-2002; house prices as measured by the Case-Shiller 20 city index, and oil 2007-2009 as proxied by the price of the USO ETF. All data are monthly and value-weighted across stocks. For housing, turnover is measured as the number of existing home sales.

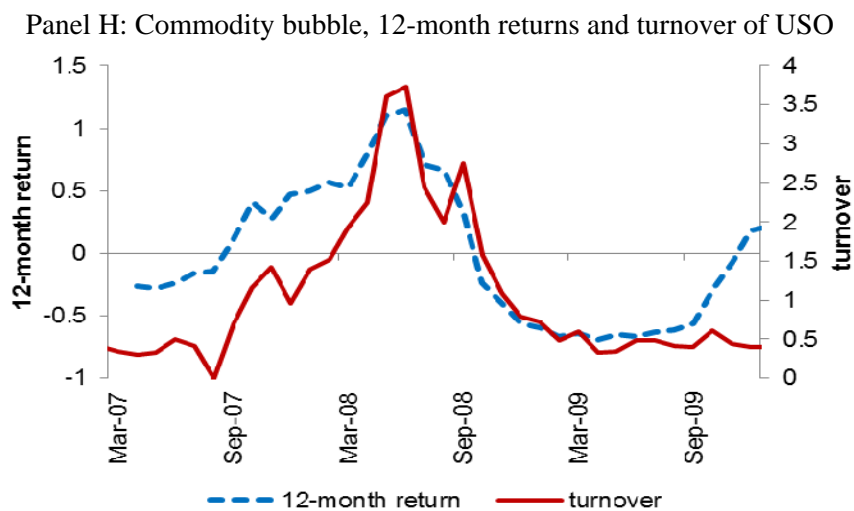
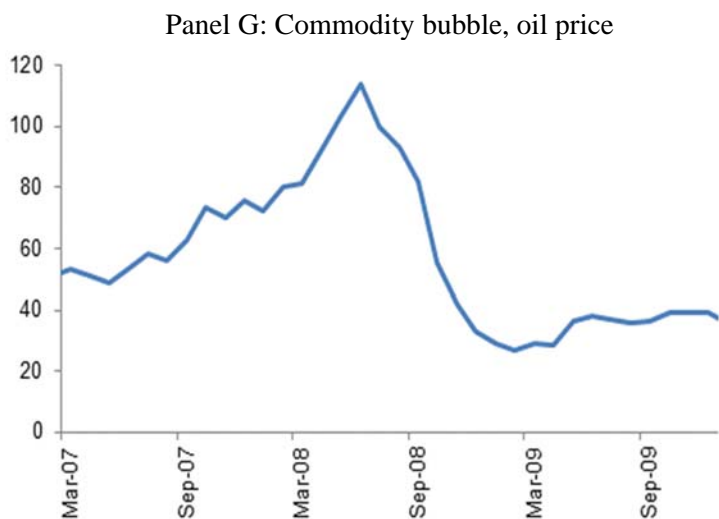
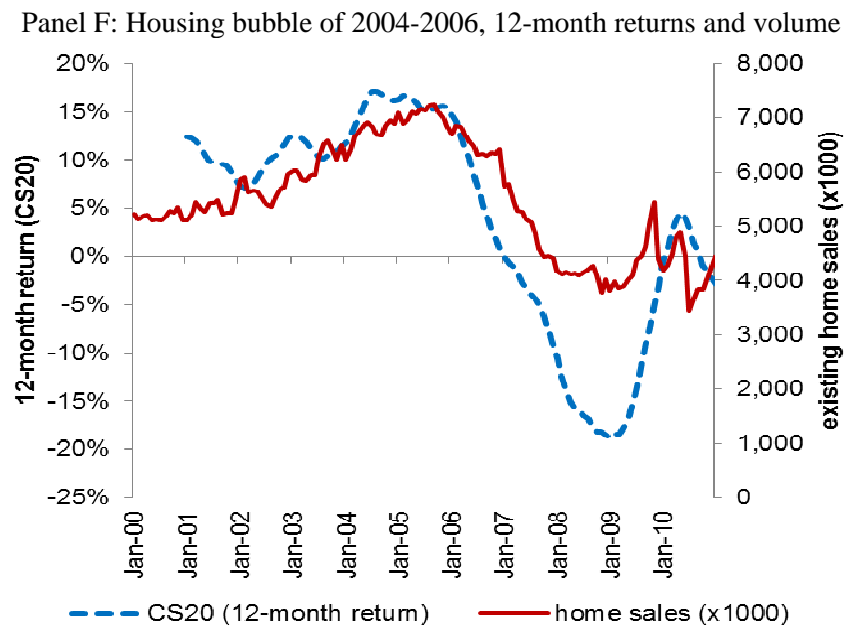
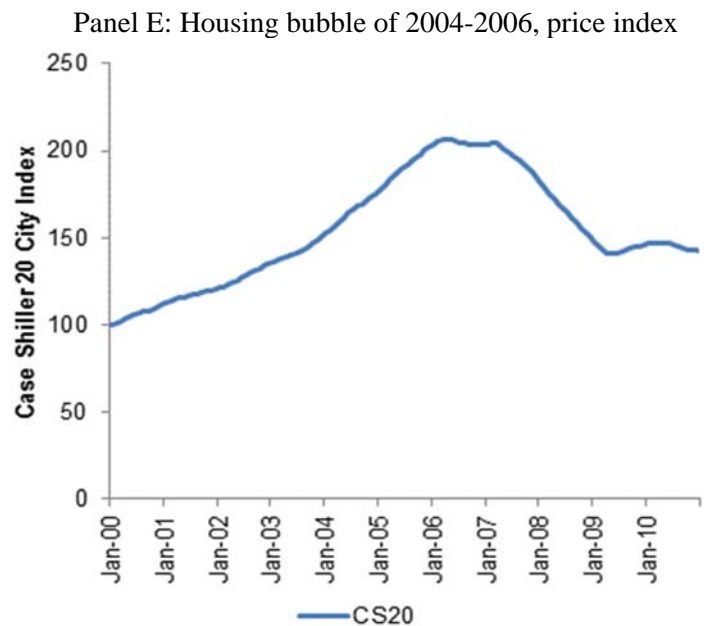
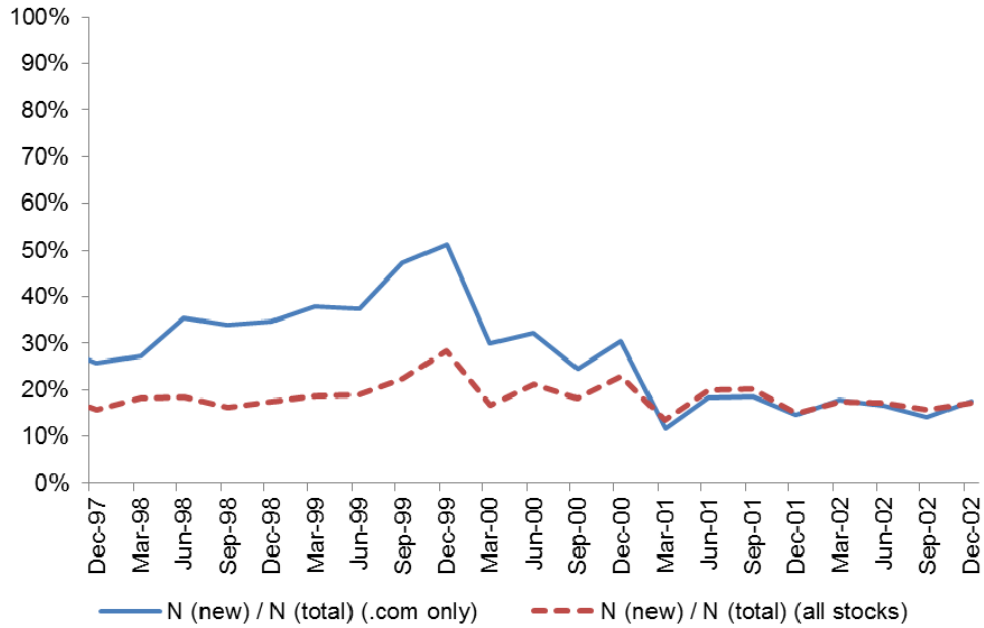


Figure 7: Prices, returns, and volume during bubble episodes [continued].

Panel A. Bubble stocks attract new investors



Panel B. "Growthiness" of owners

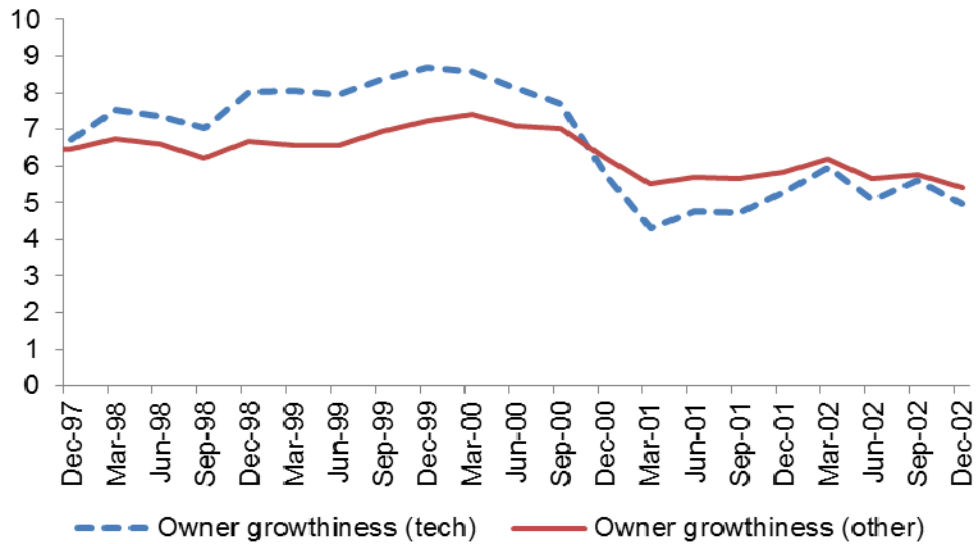


Figure 8: Changing ownership composition during the bubble. The sample includes all .com stocks classified as such by Ofek and Richardson (2003). Stocks are only considered following two quarters in which there were five or more mutual fund owners. Panel A shows the ratio of new owners in quarter t to total owners, averaged across .com stocks and all other stocks. Panel B shows the growthiness of new owners, where growthiness of a portfolio is defined as the position-weighted past return decile.