Bubbles and fools

Gadi Barlevy

Introduction and summary

In the wake of the financial crisis of 2007–08 and the Great Recession precipitated by it, a growing chorus has argued that policymakers ought to act more aggressively to rein in asset bubbles—that is, scenarios in which asset prices rise rapidly and then crash. Before the crisis, conventional wisdom among policymakers cautioned against acting on suspected bubbles. As laid out in an influential paper by Bernanke and Gertler (1999), there are two reasons for this. First, while asset prices are increasing, it is difficult to gauge whether these prices are likely to remain high or revert. Second, many of the available tools for reining in asset prices, such as raising nominal short-term interest rates, tend to be blunt instruments that impact economic activity more broadly. Bernanke and Gertler argued that rather than responding to rising asset prices, policymakers should stand ready to address the consequences of a collapse in asset prices if and when it happens.[1] The severity of the Great Recession and the challenge of trying to stimulate economic activity even after lowering short-term rates to zero led many to rethink whether policymakers should wait and see if asset prices collapse and then deal with the aftermath. Of course, just because the stakes are great does not mean policymakers can or should do anything if they are concerned about a possible bubble. To determine whether anything can or should be done, we need to understand when and why bubbles can emerge and what that might mean for policy. This article considers one explanation for bubbles known as the greater-fool theory of bubbles, as well as its implications for policy.

As a first step, let me be clear on what I mean by a *bubble*. Arguably, the distinguishing feature of the various historical episodes that are usually described as bubbles is that asset prices seem to rise too quickly, culminating in an eventual collapse of asset prices and a glut of assets created while prices were rising. This suggests defining a bubble as an asset whose price is somehow too high, which is indeed the way economists typically define the term: A bubble is an asset whose price deviates from its "natural" value. But what is the natural value for an asset? When the cash flows an asset pays out in dividends are drawn from a known probability distribution, a natural benchmark value is the expected present discounted value of the dividends it generates—also known as the fundamental value of the asset. Intuitively, society values assets for the dividends they are expected to yield, so

Gadi Barlevy is a senior economist and research advisor in the Economic Research Department at the Federal Reserve Bank of Chicago.
The author thanks Lisa Barrow, Bob Barsky, Antonio Doblas-Madrid, and François Velde for their helpful comments.

ISSN 0164-0682

the fundamental value represents how much an asset ought to be worth if it is to be supplied efficiently. In practice, though, the distribution of dividends is typically unknown; scenarios that matter for dividends may have few or no historical precedents that can be used to gauge their likelihood. If agents held different beliefs as to the likelihood of these states, it isn't clear what the benchmark value of an asset should be: Whose beliefs should we use to compute the expected present discounted value of cash flows? In what follows, I discuss some ways of extending the definition of a bubble to the case where agents have different beliefs.[2] Indeed, one of the themes of this article is that some models that purport to capture bubbles rely on a particular way of defining the fundamental value of an asset when traders hold different beliefs, though using alternative definitions would imply the asset is in fact properly priced. It is thus unclear whether these models should be viewed as capturing bubbles in the sense that the underlying asset is overvalued. That said, these models unambiguously capture a separate phenomenon, *speculative trading*, by which I mean that agents trade assets not because they expect mutually beneficial gains from trading with others, but because they expect to profit at the expense of others. Speculative trading may well be something policymakers should be concerned about, but the appropriate policy response to it need not be framed in terms of driving asset prices back toward fundamentals, as policy prescriptions for responding to bubbles often are.

It turns out that it is not easy to construct economic models that give rise to bubbles in the sense I have just described. The reason is that people will naturally be reluctant to pay more for an asset than the value of the dividends it generates. Nevertheless, there are settings in which this phenomenon can occur. One explanation is that when asset prices are equal to fundamentals, there will be a shortage of assets relative to the amount agents require for saving or liquidity or to earn a satisfactory return. According to this view, a shortage will lead agents to pile into whatever assets are available. Even if assets trade above their fundamentals, agents might still be willing to buy them given their inherent usefulness.[3]

A different explanation for bubbles is based on risk shifting: If agents can buy risky assets and borrow against them, they would be willing to pay more for assets than their expected payoff, since they can shift their losses on to their creditors by defaulting.[4]

A third explanation for why bubbles can arise—and the one this article focuses on—is known as the greater-fool theory of bubbles. According to this explanation, agents are willing to pay more for an asset than they think it is worth because they anticipate they might be able to sell it to someone else for an even higher price. Such explanations have come to be known as greater-fool theories because they all invariably involve speculative trading in the sense in which I defined it earlier—that is, traders trade assets because they expect to profit at the expense of others (who would be the greater fools) rather than because they expect mutual gains from trading.[5] This feature distinguishes this explanation of bubbles from some of the explanations based on asset scarcity that feature finitely lived agents who buy infinitely lived assets. In the latter case, agents also buy an asset intending to eventually sell it to someone else for a higher price. But in such a case, they do not expect to profit at the expense of those they trade with and would have been willing to hold on to the asset if they could.

Theories of bubbles based on asset shortages or risk shifting are straightforward and fairly well understood. By contrast, greater-fool theories of bubbles raise a host of complications, even though the idea they represent is simple and resonates with many people. For example, Edward Chancellor titled his 1999 book on the history of speculation *Devil Take the Hindmost*, alluding to the fact that whoever is the last to be stuck with the asset ends up losing.[6] Academics and non-academics both refer to investors "riding the bubble" to evoke the way one might ride up an air bubble in a champagne flute, letting go right before the bubble reaches the surface and pops. The problem is that constructing a model where such bubbles arise can be daunting. First and foremost, if we assume traders are rational and understand the underlying environment they face (as is common in most economic models), the greater-fool theory may not hang together: The traders one expects to profit off of would be aware that others are trying to exploit them and might refuse to buy the asset. Economists have figured out ways to get around this problem. But even if we succeed in getting rational agents to trade an asset in the hope of profiting at the expense of other traders, it is not clear whether this asset can be legitimately viewed as a bubble. As noted earlier, if traders hold different views about how valuable the asset is, it is not obvious how to define the asset's fundamental value. Is it the highest value of dividends any trader in the economy expects the asset to generate? As I discuss later on, some models—which I shall call asymmetric information models of bubbles—can be understood as proper models of bubbles. But even in these models, it is not obvious what appropriate policy should be. The remainder of this article explores these issues in more detail.

When greater-fool theories are a fool's errand

A natural starting point for any discussion of greater-fool theories of bubbles is the work of Tirole (1982).[7] He derived conditions under which greater-fool theories can be definitively ruled out. Thus, any successful greater-fool theory of bubbles must violate at least one of the conditions he sets forth. These four conditions are as follows:

1) The number of potential traders is finite.

2) All traders are rational, and this is common knowledge among all traders.

3) Traders start out with common prior beliefs (or "priors," for short) about the economic environment they face.

4) Resources are allocated efficiently prior to any trading taking place.

The first condition requires little explanation or justification. The second condition contains two assumptions. First, traders are assumed to be rational, in the sense that they process information in accordance with the laws of logic and probability and then act to maximize their expected utility. Second, rationality is common knowledge, implying that traders know that other traders are rational and that other traders know they themselves are rational. The third condition holds that all traders begin with the same initial beliefs about the environment they face (for example, the attributes of the assets they trade, of the markets they trade in, and so on). Different traders may subsequently receive different information that leads them to revise their initial beliefs and deviate from what others believe. In other words, the third condition requires only that traders share the same initial understanding of the environment they face before receiving any information, not that they receive the same information or always hold the same views. The fourth condition implies that individuals have no reason to trade assets beyond the attempt to profit at the expense of others, since the initial allocation is efficient and, thus, there is no other reason to trade.

If these conditions are satisfied, attempting to construct a greater-fool theory of bubbles would be a fool's errand: Tirole (1982) proves that these conditions deny the possibility of a bubble altogether. The formal proof is contained in Tirole (1982). Here, I provide a sketch of his argument, which will be useful later for understanding why "greater-fool bubbles" might arise in alternative environments. Consider a trader, whom I will call Alice, who wants to sell her asset to a trader who is not privy to the same information that Alice is and who might therefore believe the asset is more valuable than it truly is. Suppose there were such a trader, whom I will call Bob. Since Bob is rational, he would realize that the only reason Alice wants to sell him the asset is that she received information indicating that the asset is worth less than the price she is offering. Since Bob knows that he and Alice started out with the same beliefs, he realizes that if he saw the same information as Alice did, he would also be convinced that the asset is worth less than the price she is offering. As a result, even without seeing the information Alice has, he knows better than to buy the asset from her. What if Bob had incontrovertible evidence that the asset is worth more than the price at which he can buy it from Alice? In that case, Alice would realize that Bob must have information that she would find compelling, and thus seeing him eager to buy would cause her to refuse to sell. Since Alice is rational and knows that other traders are rational, she would realize that she will not be able to unload an asset for a price above its true value. Given this, she would never agree to buy an asset for more than she believes it is worth. It follows that the asset can never trade above its fundamental value.[8]

In short, developing a greater-fool theory of bubbles requires violating one of the conditions Tirole set forth. As I shall next discuss, the literature has pursued two approaches to modeling greater-fool theories of bubbles, each of which violates at least one of Tirole's conditions.

Fanfare for the uncommon prior

One modification to Tirole's (1982) setup that has attracted a great deal of attention is to dispense with the assumption that traders start out with common prior beliefs about the environment they face. Indeed, work that assumes traders have different prior beliefs—and therefore do not temper their beliefs when they see others taking different trading positions from their own—precedes Tirole's work. Examples include Miller (1977) and Harrison and Kreps (1978), who both frame their results in terms of speculative trading rather than bubbles. That is, both papers are concerned with whether agents trade expecting to profit at the expense of others that hold different beliefs, but neither is explicitly concerned with whether asset prices reflect fundamentals or not.[9] Harrison and Kreps do offer a few brief comments on fundamental valuation in the conclusion of their paper, which I discuss later. But it is only more recently, starting with the work of Scheinkman and Xiong (2003), that models featuring agents with different prior beliefs have come to be associated with bubbles.

Before I discuss whether models featuring heterogeneous prior beliefs among traders can rightly be viewed as giving rise to bubbles, let me reflect on why dispensing with the assumption that traders begin with the same priors about their environment may allow us to avoid Tirole's (1982) conclusion that rules out bubbles. Let me trot out Alice and Bob again. If Alice tries to sell her asset to Bob, then because Bob is rational, he still realizes that there is no reason Alice would want to sell him the asset other than that she believes the asset is worth less than the price she is offering. But since Bob does not start out with the same beliefs as Alice, he will not necessarily be convinced by the evidence Alice sees. Indeed, suppose nobody receives any new information to update their priors. In that case, Bob would know that Alice is trading on the basis of her priors, which he does not agree with. Thus, he might believe the asset is selling for less than its fundamental value even as Alice believes the price is above the fundamental value. Alice and Bob will agree to trade, each believing they are taking advantage of the other. Note that in getting around Tirole's result, I am not abandoning the assumption that agents are rational. Indeed, I invoke rationality throughout my analysis. This is worth pointing out, since models with uncommon priors are sometimes described as models in which agents are irrational, even though they need not be.[10]

To show how dropping the requirement of common priors can lead to scenarios that are suggestive of greater-fool bubbles, consider the following adaptation of the Harrison and Kreps (1978) model. Suppose there is a single asset, available in fixed supply that is normalized to 1. Let d_t denote the dividend this asset yields in period t, which corresponds to a single day. There are two agents, Evelyn and Odelia, who maintain different beliefs about dividends. In particular, Evelyn believes the asset yields one unit of consumption goods in even periods and nothing in odd periods:

$$\text{1)} \qquad \text{Evelyn believes } d_t = \begin{cases} 1 & \text{if } t \text{ is even} \\ 0 & \text{if } t \text{ is odd.} \end{cases}$$

Odelia instead believes that the asset yields one unit of consumption goods in odd periods and nothing in even periods, that is,

$$\text{2)} \qquad \text{Odelia believes } d_t = \begin{cases} 0 & \text{if } t \text{ is even} \\ 1 & \text{if } t \text{ is odd.} \end{cases}$$

I assume that Evelyn and Odelia maintain these beliefs regardless of what dividends are actually paid out. That is, their respective theories about dividends at different periods are logically independent, so even if their theories about dividends at some period are ever proven wrong—and in each period at least one of their two theories about dividends that period will always be proven wrong—neither trader will change or update her expectations about future dividends.[11] I could have replaced the aforementioned beliefs with beliefs that involve nondegenerate probabilities—for example, Evelyn believes that $d_t = 1$ in any even period with probability $1-\varepsilon$ and $d_t = 0$ with probability ε for some small but positive ε. This way, nobody would ever be explicitly proven wrong, since each of their theories would allow both realizations for dividends. But I assume degenerate beliefs to simplify the exposition. The essential feature of my example is that Evelyn is more optimistic about dividends in even periods and Odelia is more optimistic in odd periods.

I assume the asset in question cannot be sold short—that is, a trader can sell any units of the asset she already owns, but she cannot borrow additional units to sell. Evelyn and Odelia take prices as given when they trade.[12] I also assume Evelyn and Odelia have ample endowments each period that allow them to purchase the entire fixed supply of the asset should they wish to do so. Finally, I assume that both have a utility that is linear in the amount of consumption goods they eat, implying both are risk-neutral, and that both discount the future at the same rate β, where $0 < \beta < 1$. Trade takes place in the morning of each period, while the dividend is paid out that evening. Consumption goods are not storable, so a trader who buys the asset in the morning expects to consume any dividends the asset generates that same night.

Consider period 1. On the one hand, because the date is odd, Evelyn believes that the asset will bear no dividend today, but that it will bear a dividend one period from now, three periods from now, five periods from now, and so on. Hence, she would value the present discounted dividends from the asset at

$$\text{3)} \qquad \beta + \beta^3 + \beta^5 + \ldots = \frac{\beta}{1-\beta^2}.$$

On the other hand, Odelia believes the asset will pay a dividend today, two days from now, four days from now, and so on. Hence, she would value the present discounted dividends from the asset at

$$\text{4)} \qquad 1 + \beta^2 + \beta^4 + \ldots = \frac{1}{1-\beta^2}.$$

Since $0 < \beta < 1$, Odelia values the dividends paid by the asset more than Evelyn. In even periods, the two valuations switch, and Evelyn values the asset at $\frac{1}{1-\beta^2}$, while Odelia values it at $\frac{\beta}{1-\beta^2}$.

I will now argue that in equilibrium—when the supply of the asset equals the demand for it—the price of the asset at each date t, denoted p_t, will exceed $\frac{1}{1-\beta^2}$, the most either Evelyn or Odelia thinks the flow of dividends from the asset from that date on is worth. To see this, observe that if the price of the asset were ever below $\frac{1}{1-\beta^2}$ at some period t, there would always be someone (Evelyn if t is even or Odelia if t is odd) who would want to buy as many units of the asset as her endowment would allow, since she thinks she can earn strictly positive profits from buying the asset and holding it indefinitely. Since I assumed that both traders have ample resources but that the supply of the asset is fixed, demand would exceed supply. So the price cannot fall below this level, that is, $p_t \geq \frac{1}{1-\beta^2}$ for all t.

Could the price be equal to $\frac{1}{1-\beta^2}$, the highest valuation any agent assigns to a dividend in each period? Suppose it did. In period 1 (or at any odd date for that matter), Odelia could contemplate the following strategy: Buy the asset, anticipate consuming its dividend that evening, and then sell the asset for a price of at least $\frac{1}{1-\beta^2}$ in the next period. She would discount the proceeds of her sale by β, and so the payoff to this strategy is at least

$$1+\beta\left(\frac{1}{1-\beta^2}\right) = \frac{1+\beta-\beta^2}{1-\beta^2}.$$

But since $0 < \beta < 1$, it follows that

$$\frac{1+\beta\,(1-\beta)}{1-\beta^2} > \frac{1}{1-\beta^2}.$$

Hence, if the price were equal to $\frac{1}{1-\beta^2}$ in period 1, Odelia would expect to earn strictly positive profits from buying the asset. She would therefore want to buy as many units of the asset as her endowment allows, and demand would exceed supply. The fact that the price of the asset can never fall below $\frac{1}{1-\beta^2}$ at any date implies that it must be at least $\frac{1+\beta-\beta^2}{1-\beta^2}$ in period 1. Of course, by the same logic, it must also be at least $\frac{1+\beta-\beta^2}{1-\beta^2}$ in period 2, or else Evelyn could buy the asset, earn the dividend in period 2, and then sell the asset in period 3. And the same argument applies in subsequent periods ($t = 3, 4, 5$, and so on), so the price can never fall below this new level, meaning $p_t \geq \frac{1+\beta-\beta^2}{1-\beta^2}$ for all t. Thus, I have established a new lower bound on prices for all dates—which is higher than the original bound of $\frac{1}{1-\beta^2}$ that I started with.

I can now repeat the argument: Given the price is at least $\frac{1+\beta-\beta^2}{1-\beta^2}$ at all dates, can it ever equal this bound at any date? If that were indeed the price in period 1 (or at any odd date), Odelia could buy the asset in period 1, consume its dividend that evening, and then sell the asset for at least $\frac{1+\beta-\beta^2}{1-\beta^2}$ in the following period. Her payoff in that case would be at least

$$1+\beta\left(\frac{1+\beta-\beta^2}{1-\beta^2}\right) = \frac{1+\beta-\beta^3}{1-\beta^2}.$$

Since $0 < \beta < 1$, it follows that

$$\frac{1+\beta-\beta^3}{1-\beta^2} > \frac{1+\beta-\beta^2}{1-\beta^2}.$$

Hence, Odelia would expect to earn strictly positive profits from this strategy, and so she should buy as many units of the assets as her endowment allows. To ensure supply is equal to demand, the price in period 1 must be at least $\frac{1+\beta-\beta^3}{1-\beta^2}$, the lowest profit Odelia can earn by holding the asset one period and then selling it in the next period. Once again, this argument can be applied to every period, and so I can conclude that $p_t \geq \frac{1+\beta-\beta^3}{1-\beta^2}$ for all t, a bound that is higher than in the previous round.

I can apply this argument repeatedly: Given the price exceeds a new threshold in every period, I can derive a now higher bound on the price for each period. In particular, on the nth iteration, I will be able to conclude that $p_t \geq \frac{1+\beta-\beta^n}{1-\beta^2}$. Since this holds for any n, it follows that

$$5) \qquad p_t \geq \lim_{n\to\infty} \frac{1+\beta-\beta^n}{1-\beta^2} = \frac{1+\beta}{1-\beta^2} = \frac{1}{1-\beta}.$$

The limiting case where $p_t = \frac{1}{1-\beta}$ for all t turns out to be an equilibrium price. To see this, observe that the expected payoff for Odelia from buying the asset in period 1, consuming its dividend that evening, and selling at price $\frac{1}{1-\beta}$ the next day is equal to

$$1+\beta\left(\frac{1}{1-\beta}\right)=\frac{1}{1-\beta}.$$

What Odelia pays for the asset is thus exactly equal to the profit Odelia would earn from buying and selling the asset. One can show that buying and selling after one period is the best Odelia can do—that is, holding the asset for longer and then selling it will be less profitable. Hence, Odelia is just indifferent between buying the asset in period 1 and not buying it at all. Evelyn, by contrast, wants to sell the asset (and would even sell it short if she could), since she believes the asset will yield no dividend that evening. Hence, supply and demand for the asset can be equal. At the equilibrium price path, Evelyn sells all her asset holdings to Odelia in odd periods, and Odelia sells all her asset holdings to Evelyn in even periods. I can appeal to arguments in Harrison and Kreps (1978) to argue that with some additional assumptions, $p_t = \frac{1}{1-\beta}$ for each t is the only possible equilibrium price path for the asset.[13]

To recap, the equilibrium price of the asset exceeds what either Evelyn or Odelia believes the asset can generate in dividends. Some have argued that this implies the asset in my example should be viewed as a bubble. Specifically, they argue that when agents have different beliefs, a bubble should be defined as follows. First, define a fundamental value for each individual as what that individual expects the cash flow from the asset to be or, alternatively, how much each individual would value holding the asset indefinitely and consuming its dividends. Then define an asset to be a bubble if its price exceeds every individual's fundamental value. Note that when the distribution of dividends is known so that all agents have the same beliefs, this definition reverts to the original definition of a bubble for the case where the distribution of dividends is known by all agents. Hence, this definition extends the definition of a bubble for a known distribution for dividends to the case where the distribution of dividends is not known and agents can hold different beliefs.

The equilibrium I have just constructed would thus seem to provide an internally consistent model of a greater-fool bubble. The market clearing price for the asset is higher than anyone in the economy believes dividends are worth. Nevertheless, traders buy the asset at this price, precisely because they expect to sell it later to someone who values the asset even more than they do. The fact that the equilibrium price is constant rather than growing may make this seem like an unusual model of a bubble, since most historical episodes suspected to be bubbles feature rapid price appreciation. But the price is constant because the dividends in my example are constant over time—an assumption I imposed for convenience. The model can be readily modified to allow for dividend growth in a way that would introduce price booms and busts without changing its key features.[14] Still, as I next explain, it is not obvious that this model should be interpreted as a model of a bubble, since alternative ways of extending the definition of fundamental value to the case where traders hold different beliefs do not imply the asset is overvalued.

Is it a bubble?

To illustrate the complications for interpreting the previous example as a bubble, I consider the following related example of an economy with two types of goods—say, apples and bananas. As before, there are two people in the economy—I'll again call them Evelyn and Odelia—each of whom is endowed with an ample amount of apples each period. Evelyn and Odelia have the same beliefs, but now their *preferences* differ. Evelyn enjoys bananas on only even days, when she derives the same pleasure from one apple as she does from one banana. On odd days, Evelyn derives no utility from bananas. Odelia's tastes are the exact opposite: She enjoys bananas on only odd days, deriving the same utility from a banana as from an apple. On even days, Odelia derives no utility from bananas. There is no uncertainty, and both Evelyn and Odelia discount at the same rate $\beta \in (0, 1)$.

Suppose this economy had no bananas initially, and I contemplated introducing a banana tree that bears one banana each day. How much would this tree be worth in terms of apples? Consider first the perspective of an outsider who shares the same discount rate β as that of Evelyn and Odelia. The outsider could sell the tree's yield of one banana each day in exchange for one apple. On even days, he would sell the banana to Evelyn, while on odd days, he would sell it to Odelia. Hence, the present discounted value of the tree for the outsider as measured in apples is just

$$6)\qquad 1+\beta+\beta^2+\beta^3+\ldots=\frac{1}{1-\beta}.$$

The same would be true if I considered the perspective of either Evelyn or Odelia. For example, if I asked Odelia to value the tree, she would reason that on odd days she could consume the banana directly, which she values the same as an apple, while on even days she could sell a banana to Evelyn in exchange for an apple. Thus, she would value the tree as the present discounted value of receiving an apple each day. The same would be true if I asked Evelyn. Since there is no uncertainty, the usual definition of the fundamental value of an asset would imply the banana tree is worth $\frac{1}{1-\beta}$ apples.

Now, what would happen if I precluded people from selling bananas but still let them buy banana trees? That is, I would shut down the market for the dividends generated by the asset, but not the market for the asset. This restriction precludes the arrangements that I used to argue the tree is worth $\frac{1}{1-\beta}$ apples. However, Evelyn and Odelia could still achieve the same allocation as with a market for bananas by trading the banana tree in such a way that ensures the person who values bananas owns the tree when it yields fruit. That is, Evelyn will buy the tree the morning of each even date, consume the banana that evening, and then sell the tree to Odelia the following morning. Since the allocation is the same as before, the value of the tree should be unchanged—that is, it should still be $\frac{1}{1-\beta}$. One can verify that the equilibrium price of the tree would still equal $\frac{1}{1-\beta}$ each period.

Now, suppose I asked Odelia and Evelyn how much they would value the tree if they couldn't sell it and had to consume its bananas themselves. Since Odelia enjoys bananas in only odd periods, if I asked her valuation in period 1, she would say she values owning the tree indefinitely at

$$7)\qquad 1+\beta^2+\beta^4+\ldots=\frac{1}{1-\beta^2}.$$

Evelyn, who enjoys bananas in only even periods, would say that owning the tree in period 1 and consuming its bananas is worth

$$8)\qquad \beta+\beta^3+\beta^5+\ldots=\frac{\beta}{1-\beta^2}.$$

Thus, the asset trades at a price above what Evelyn or Odelia thinks it is worth if either had to consume its yield on her own. Nevertheless, Evelyn or Odelia agrees to buy the asset at this price because each agent intends to sell the tree at a price that exceeds the value of consuming its fruit herself.

The connection between this example and the case with traders holding heterogeneous beliefs should hopefully be clear. The two have the same underlying structure: Each period, there is one person who values the dividend of the asset at 1, while the other values the dividend at 0. In the case where traders hold heterogeneous beliefs, this difference in valuation occurs because one of the traders believes a dividend will be paid out that period and the other doesn't. In the case where traders have heterogeneous preferences, this difference in valuation occurs because one of the traders values the good, while the other doesn't. In both cases, the trader who doesn't value the dividend that accrues that evening sells the asset to the trader who does. The price of the asset is the same in both cases. Given traders with heterogeneous preferences, it seems clear that the asset is trading at its fundamental value, even though it exceeds the value each trader assigns to owning the tree forever and consuming its dividends. Why shouldn't one say the same when agents have different beliefs, rather than different tastes?

Comparing the two examples reveals a shortcoming with defining a bubble as an asset whose price exceeds each individual's fundamental value or, in other words, as an asset whose price exceeds what each individual is willing to pay to consume the asset's dividends indefinitely. This is most readily apparent when there is an explicit market for dividends—for example, when individuals can sell bananas as opposed to just banana trees. In that case, all agents agree that the value of owning the tree indefinitely is $\frac{1}{1-\beta}$, because any agent who owns the asset can sell its dividends to those who value them most. In the case where traders hold heterogeneous beliefs, I implicitly ruled out this possibility by not allowing a market for dividends that was analogous to a market for bananas. Without such a market, agents are forced to trade the asset to achieve the same outcome, making it seem as if trading the asset makes it more valuable. But the same value can be achieved without ever transferring ownership of the asset. Indeed, the same outcome could be achieved by introducing a rental market for the asset. Just as capital equipment can be rented out to others who can keep the cash flow they generate using it, an agent who owns a financial asset can in principle rent it out for a period and let whoever rents the asset accrue its dividends. While rental markets for financial assets may seem odd, they do have historical precedents. Velde (2013) describes rental markets for government lottery bonds in eighteenth-century England. Lottery bonds were structured so that the interest payments on any particular bond were random, and on any given day there was some chance a particular bond would be drawn and receive a prize interest payment. At the time, individuals could rent a lottery bond for a day and earn the associated payout if the bond happened to be drawn that day. These arrangements were known as "horses," and their prices were published regularly.[15]

In my example where traders had heterogeneous beliefs, if I introduced the possibility of renting out the asset, both Evelyn and Odelia would value owning the asset indefinitely at $\frac{1}{1-\beta}$. It therefore seems reasonable to view $\frac{1}{1-\beta}$ as the fundamental value of the asset. That is, in the special case where agents agree on the value of an asset while they hold different beliefs, it would seem natural to define the fundamental value of the asset as this common value—namely, what any agent could earn from buying the asset and holding it indefinitely, but with the option of renting it out in any period. To see further why this is a reasonable definition for the fundamental value of the asset, consider a benevolent social planner who contemplates creating another asset on behalf of agents in this economy. The planner would value the asset in terms of the total surplus that could be created by promising its dividends at different dates to different traders. That is, the planner could collect $\frac{1}{1-\beta^2}+\frac{\beta}{1-\beta^2}=\frac{1}{1-\beta}$ from the two agents to create another asset by promising to give any future dividends that accrue in even periods to Evelyn and any future dividends that accrue in odd periods to Odelia. This suggests $\frac{1}{1-\beta}$ accurately reflects the value to society from creating another asset, which is what the notion of a fundamental value is meant to capture. Harrison and Kreps (1978) offer a similar interpretation, writing in the conclusion to their paper that the equilibrium price they derive "is consistent with the fundamentalist spirit, tempered by a subjectivist view of probability."

Why, then, have some argued for treating the asset as a bubble if its price exceeds how much each agent values holding the asset indefinitely and consuming its dividends? Undeniably, the example in which Evelyn and Odelia hold different beliefs contains features that make it reminiscent of a bubble. For example, Evelyn and Odelia both agree that the asset pays dividends only every other period, although they disagree as to the periods in which these dividends will be paid out. Isn't a price that is equivalent to the asset paying out a dividend every period too high given neither agent believes this to be the case? This characterization of beliefs, however, is misleading. If Evelyn and Odelia disagreed about only when dividends are paid out, Odelia would not expect to sell the asset to Evelyn after consuming its dividends for a price of $\frac{1}{1-\beta}$, since she knows Evelyn would realize she was wrong. Rather, the price of $\frac{1}{1-\beta}$ emerges because Evelyn and Odelia continue to believe dividends will be paid out in different periods regardless of what happened in the past, which is perfectly rational if dividends in different periods are determined through logically independent processes. The price of $\frac{1}{1-\beta}$ can be rationalized using the most optimistic beliefs any trader holds about dividends in each period. In other words, an outsider who could rely on only Evelyn's and Odelia's beliefs would be unable to rule out the possibility that dividends will actually be paid out each period, since for each period he can find a logical theory advanced by either Evelyn or Odelia that implies a dividend will be paid out.

Still, the notion that the price can be supported by always appealing to the most optimistic beliefs about dividends may seem suspect. Isn't it implausible that it is always the most optimistic traders who are correct? Depending on how agents form their beliefs, it may indeed be implausible to always rely on the most optimistic agents to determine the fundamental price of the asset. For example, Scheinkman and Xiong (2003) assume agents receive signals about dividends but attribute too much precision to their signals. This implies that the traders who are the most optimistic tend to also be *excessively* optimistic. But this does not mean that the reason assets are overvalued is because individuals have different beliefs. Even when traders hold the same beliefs, they might still be overconfident about the signals they receive. The reason Scheinkman and Xiong are correct to call the asset in their model a bubble is because they drop Tirole's second condition, which holds that traders are rational and process information correctly, rather than his third condition, which holds that they have common prior beliefs. Without any information about how traders form their beliefs, there is no reason to dismiss the beliefs of the most optimistic agents any more than those of other agents. Relying on the most optimistic beliefs corresponds to the usual notion of maximum willingness to pay that economists routinely rely on to determine how resources should be allocated.

Finally, models in which agents hold heterogeneous beliefs, as in the example I've constructed, imply traders who buy and sell the asset expect to profit at the expense of others they think value the asset incorrectly. This feature makes my example a good model of a greater-fool theory, but not necessarily a model of a bubble. In other words, this feature makes the example a good model of speculative trading as opposed to a good model of an asset that is overvalued. Indeed, what is striking about the example is that even though traders disagree about dividends, they can agree on what the asset is worth. In particular, both traders in my example view the asset as worth $\frac{1}{1-\beta}$ if given the option to rent it out, and would view accepting any price for the asset below this one as a bad trade. The fact that traders believe others are fools does not necessarily imply that they must think the asset is overvalued.

Disagreement on valuation

In my example in which Evelyn and Odelia had different beliefs, both valued holding the asset indefinitely equally provided they could rent out the asset. This equality in valuation is due to a particular feature of this example—namely, that traders agree about the distribution of the most optimistic valuation for dividends in every period. This feature can arise in other environments. For example, Scheinkman (2014) presents a model in which beliefs are independent across time. Specifically, Scheinkman assumes two types of traders. Type A traders believe the dividend in each period is equally likely to be 0 or 1. Type B traders, independent of their beliefs in other periods, will with probability $1 - 2q$ share the same beliefs that type A traders hold; but with probability q, type B traders believe the dividend that period will be 1, and with probability q, these same type B traders believe the dividend will be 0. In this case, type A and type B agents still agree about the expected value from holding the asset indefinitely given they both have the option to rent it out.

More generally, though, traders might disagree about the distribution of the most optimistic beliefs in future periods. In that case, they will disagree about the value of holding the asset indefinitely. Indeed, this is true in both the Harrison and Kreps (1978) model and the Scheinkman and Xiong (2003) model. To illustrate this possibility, I consider the following example. Suppose there is a single asset that pays one dividend in period 4, which can assume one of four values, that is,

$$9)\quad d_4 \in \{0, 1, 2, 3\}.$$

To motivate this example, suppose d_4 represents the profits of an agricultural company that plants three trees, each of which can yield a harvest of at most 1. If a tree bears fruit, it will do so in period 4. However, it will be possible to tell whether some trees will bear fruit before period 4. In particular, whether the first tree will bear fruit is revealed in period 2, whether the second tree will bear fruit is revealed in period 3, and whether the third tree will bear fruit is revealed in period 4, at the time of the harvest.[16]

There are two traders, Alice and Bob, who can trade shares in the agricultural company as news about the trees is revealed. Neither discount consumption over time. Alice's beliefs can be summarized as follows:

1) Unless given evidence to the contrary, Alice believes that with probability 0.9, all three trees will bear fruit and that with probability 0.1, none of the trees will bear fruit.

2) If just one of the first two trees bears fruit, Alice believes that the third tree will bear fruit.

Bob's beliefs can be summarized as follows:

1) Unless given evidence to the contrary, Bob believes that only the second tree will bear fruit.

2) If neither of the first two trees bears fruit, Bob believes the third tree will bear fruit.

3) If the first tree bears fruit, Bob believes no other trees will bear fruit.

4) If both of the first two trees bear fruit, Bob believes that the third tree will not bear fruit.

These conditions fully describe what Alice and Bob believe depending on what they know at the beginning of each period. Figure 1 shows the same information

FIGURE 1

Alice's and Bob's beliefs given the number of good trees in each period

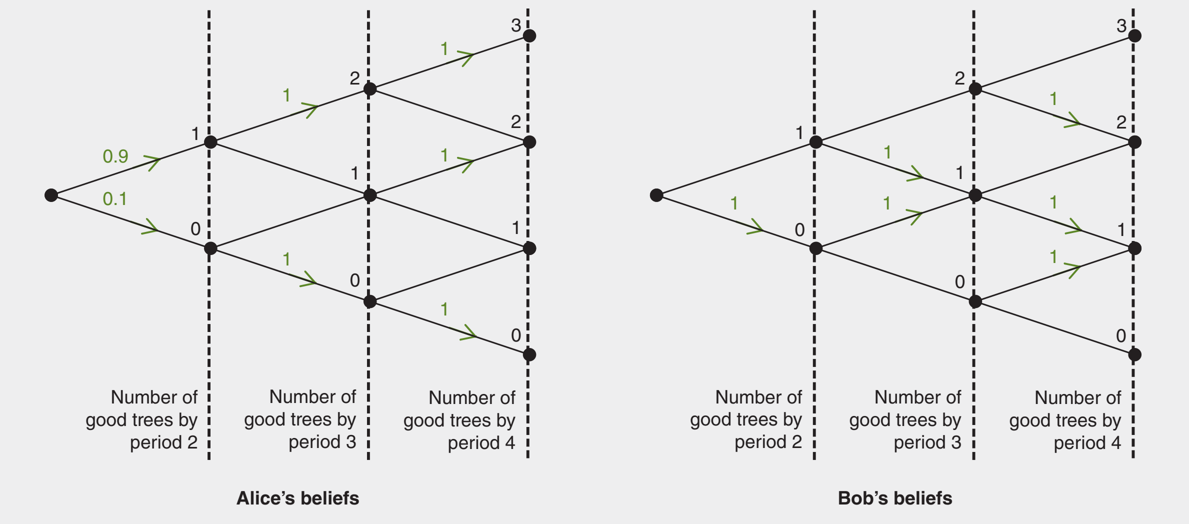

graphically, with the numbers in green indicating the probability each trader assigns to what will happen at each node in the information tree.

Given these beliefs, consider how Alice and Bob value owning the agricultural company in period 1, before the status of any tree is revealed. First, consider the value of consuming the dividend—that is, ignoring the possibility of renting out the asset while maintaining ownership of it. In that case, Alice values dividends at the following units of consumption:

10) $0.9 \times 3 + 0.1 \times 0 = 2.7.$

Bob instead values dividends at 1 unit, fully expecting only the second tree to bear fruit.

However, I argued earlier that the value of owning an asset indefinitely should incorporate the value of renting out the asset to those who have the most optimistic beliefs. In this example, the asset makes a single dividend payment in period 4. Whoever owns the asset in period 4 can thus rent it out before the status of the last tree is revealed. In this case, Alice understands that if the first two trees do not yield any fruit, Bob will still believe that third tree will yield fruit, and she will be able to rent her share to him for 1 unit. Hence, she values the asset at

11) $0.9 \times 3 + 0.1 \times 1 = 2.8.$

Bob instead expects that only the second tree will bear fruit, at which point Alice will still expect the third tree to bear fruit. Hence, he can count on renting the asset to Alice rather than consuming the fruit it yields, and so he would value owning the asset but being able to rent it out to the highest bidder at 2 units of consumption.

Alice and Bob now disagree as to the value of owning the asset indefinitely, even when given the option to rent it out. How then should the fundamental value of the asset be defined in this case? One possibility is to define the fundamental value of the asset as the most any agent would pay at any given date for the right to own the asset indefinitely but still rent it out. In this case, the value in period 1 would be 2.8, the amount Alice thinks the asset is worth. I next show that the equilibrium price of the asset in period 1 will exceed this value, so according to this definition the asset should be viewed as a bubble.

Since Alice values the asset more than Bob in period 1, she will outbid him and own the asset at that point in time. If the first tree turns out not to bear any fruit, though, she could sell it to Bob. Recall that in this case Bob will believe that the second tree will bear fruit and that he will then be able to rent the asset to Alice for 1 unit, so Bob would value the asset at 2 units. By selling the asset to Bob if the first tree does not bear any fruit, Alice would guarantee herself an expected payoff of

12) $0.9 \times 3 + 0.1 \times 2 = 2.9.$

Hence, if the price of the asset was only 2.8, Alice would want to buy infinitely many units of the asset. The only way to ensure she demands finitely many units of the asset is if the price is 2.9. In this case, the asset is more valuable to Alice precisely because she can transfer the asset to Bob, something that cannot be replicated by simply renting the asset to him. Essentially, if the first tree does not bear any fruit, Alice and Bob disagree in period 2 about what Alice will believe in period 4; Bob expects she will think the asset is worth 2, while Alice expects she will believe the asset is worth 1. As a result, in period 2, Bob thinks that owning the asset and later renting it out is more valuable than Alice does. The only way for Alice to profit from Bob's beliefs is by selling him the right to rent out the asset in the future.

Should this example be viewed as a bubble? The price exceeds what any trader believes holding the asset indefinitely is worth (even with the option to rent out the asset). However, once again, care must be taken in how the fundamental value of the asset is defined. Suppose that agents have different initial beliefs and that tomorrow all agents might receive news that reveals the asset's dividends will likely be higher. If the news comes, since traders have different beliefs, not all of them will process this information in the same way. Some traders may not update their beliefs. However, these traders should still recognize that the asset would be more valuable now for society as a whole, and incorporate the impact of the news on the beliefs of others when assessing the fundamental value of the asset. Just as an agent would naturally take into account that tomorrow he might receive news that affects his beliefs about dividends when valuing the asset today, he should also take into account that tomorrow others might receive news that affects their beliefs, even if it doesn't affect his own beliefs, when making his valuation today. For the example in figure 1, to determine the value of the asset, it is essential to know what Alice will believe about the dividend in period 4 given she might be the one who consumes this dividend. But agents disagree about what Alice will believe then. Without any additional information on how agents form their beliefs, there is no reason to suppose Alice knows what her beliefs will be better than others. By the same logic that I described in the previous examples, it would seem natural to rely on the most optimistic beliefs about what Alice will believe, rather than on Alice's own beliefs about what she will believe, to determine the value of the asset.

This logic suggests that the definition of a bubble in environments where agents disagree about future beliefs should not be based on using the same individual's beliefs at all dates—for the same reason as in my earlier example where Evelyn and Odelia held different beliefs. In finite-horizon settings, one could determine which agent holds the most optimistic beliefs at each possible state of the world on the final date in which the asset yields a dividend. Working backward, one could evaluate at each prior node the agent who holds the most optimistic beliefs when that agent's future beliefs are substituted with the beliefs of the trader who was already determined to be the most optimistic at future nodes. Because this is notationally cumbersome, I omit the formal details. This approach is faithful to Harrison and Kreps's (1978) observation that the equilibrium price in a model where traders hold different priors has a fundamentalist spirit, tempered by subjective beliefs. Moreover, when agents have the same beliefs, this construction would revert to the usual definition of fundamental value when the distribution of dividends is known. The fact that agents disagree on the value of holding the asset indefinitely, even with the option to rent it out, doesn't automatically imply that the asset should be considered a bubble.

What does this all mean for policy?

So far, I have argued that models in which traders have different prior beliefs exhibit speculative trading, but there are good reasons not to view them necessarily as models of bubbles in which the underlying asset is overvalued. However, the relevant question is arguably not whether traders holding heterogeneous beliefs give rise to a bubble *per se*, but whether the possibility of traders holding heterogeneous beliefs can somehow justify policy intervention. That is, irrespective of whether these models give rise to a bubble or not, do they imply that policymakers should discourage agents from speculative trading, meaning trading with the aim of profiting at the expense of others rather than to achieve mutual gains? After all, the definition of a bubble is sufficiently difficult to apply in practice that the more relevant question may be whether policymakers should act to curb speculation. Unfortunately, the social welfare analysis of models where agents hold heterogeneous beliefs turns out to also be ambiguous.

As a starting point, consider the case where the asset is in fixed supply. This avoids the question of whether speculation results in too high of a price that would encourage agents to create too many units of the asset. Instead, the relevant policy question is whether people should be allowed to trade on the basis of different opinions. The similarity between trade among agents with heterogeneous beliefs and trade among agents with heterogeneous preferences in my earlier examples suggests traders should be allowed to enter into such trades. In both of my examples, Evelyn and Odelia wanted to trade. In one case, Evelyn preferred owning the asset to her endowment of consumption goods and Odelia preferred the opposite, while in the other case

Evelyn preferred an apple to a banana and Odelia preferred the opposite. It is true that in the former case, the agents' willingness to trade is premised on mutually inconsistent beliefs. However, when agents have inconsistent beliefs, they are aware that their beliefs are incompatible and still wish to trade. Absent additional information on why they disagree, why should they be denied their mutual desire to trade?

Several economists have argued that policymakers should deny people the right to trade on the basis of heterogeneous beliefs. Most of these offer variations on a similar point: When two agents trade on the basis of heterogeneous preferences, like trading an apple for a banana, they will not regret their trade ex post. By contrast, when two agents trade on the basis of heterogeneous beliefs, at least one of them is bound to be proven wrong and regret making the trade. For example, in discussing trades based on differences in beliefs, Stiglitz (1989, p. 106) argues that "impeding trade is (Pareto) inefficient when viewed from the perspective of their ex ante expectations"[17] and "impeding trade may actually improve social welfare when viewed from the perspective of their ex-post realizations." He likens this to a parent who forces his child to study in a way the child will appreciate later. Mongin (2005) discusses the notion of *spurious unanimity* in which all individuals agree to take the same action but only because they believe the action will result in different outcomes that they value differently. He offers an example in which a majority of people in a country support building a bridge to a neighboring country—some of them because they believe it will lead to a massive inflow of people who will revitalize the local economy and others because they believe unwanted newcomers will largely stay away, while locals will be able to travel outside. It is easy to construct examples in which, regardless of which hypothesis proves to be correct, a majority of people will oppose building the bridge knowing this hypothesis to be true. Trade between agents with heterogeneous beliefs has a similar flavor: The two parties do not receive mutual gains from trade but are engaged in a zero-sum game where each one is expecting to benefit at the expense of the other. Since both cannot be correct, one of them is bound to regret having entered into the trade.

A related but distinct argument is laid out in Blume et al. (2014) and Brunnermeier, Simsek, and Xiong (2014). They consider the case where agents hold different beliefs but where trade is costly because it introduces consumption volatility.[18] To appreciate their argument, suppose individuals are risk averse. Start with the case where all individuals have the same fixed endowment, so that they don't need to face any consumption risk. However, if they enter a bet based on their different prior beliefs, they will each be convinced they are correct and expect to gain from the bet. Of course, only one of them will win the bet. Thus, they expose themselves to unnecessary consumption volatility. If the individuals were told in advance, before they ever form beliefs, that they will be exposed to a risk that their levels of consumption will be volatile and perfectly negatively correlated with one another, they would prefer to enter into an ex-ante insurance arrangement to guard against this risk. Here, the risk involves the two of them forming different priors that can't both be correct. Insuring against this possibility is equivalent to agreeing in advance not to trade on the basis of heterogeneous beliefs. Thus, trade on the basis of heterogeneous beliefs in this case is not a zero-sum game, but a negative-sum game, and the parties may be better off ex ante if they could commit not to enter into such trades.

Guided by these observations, Brunnermeier, Simsek, and Xiong (2014) propose a notion of *belief-neutral Pareto improvement*:[19] According to their definition, allocation A is said to improve on allocation B if, given a particular set *S* of probability distributions, it can be verified that no agents are worse off under allocation A and some are strictly better off than under allocation B when all agents' expected utilities are evaluated at each of the probability distributions in the set *S* instead of by what individuals actually believe. In particular, everyone must be no worse off and some strictly better off under allocation A when expected utility is computed using each of the agent's beliefs, as well as any mixture of the beliefs of the different agents.[20] According to this criterion, for risk-averse agents with equal endowments but different beliefs, betting with each other is dominated by a policy that precludes them from betting with each other. This suggests policymakers might want to disallow such trades.

Brunnermeier, Simsek, and Xiong (2014) go on to argue that when the supply of the asset is variable rather than fixed, there can be an additional social cost from allowing agents to trade: If the price of the asset exceeds its fundamental value, too many resources will be allocated to creating this asset. This argument of course presumes that agents with heterogeneous beliefs who are allowed to trade lead to bubbles. As I discussed earlier, it is not obvious whether the asset should be viewed as a bubble if all that is known is that agents differ in beliefs. Brunnermeier, Simsek, and Xiong (2014) recognize this, and argue that their analysis only applies when beliefs are distorted. In this case, the notion that speculation encourages an oversupply of bubble assets does not rely on the fact that people have different beliefs; rather, that notion relies on some people having distorted beliefs.

Although these arguments for how to treat trade based on agents' heterogeneous beliefs have their merits, it is safe to say that their implications for policy remain controversial. Stiglitz's (1989) analogy to paternalism is imperfect, since parental intervention is typically defended on the grounds that children are unable to reason or comprehend the consequences of their actions. Far fewer would argue that parents should continue to intervene in their children's decisions when their children are adults. If the agents who hold different beliefs are rational in the sense of reasoning based on logic and probability, a planner who argues they shouldn't be allowed to trade because their beliefs are incompatible would not be telling them anything they don't already realize. The fact that they are nevertheless willing to trade substantially weakens the case for intervention. As for the argument about ex-post regret, if beliefs correspond to nondegenerate probability distributions about events that are rarely replicated, individuals may never learn whether they were correct or not. Traders may simply chalk up their losses to bad luck, in the same way that a risk-averse agent will understand the fact that a calamity didn't happen does not mean it was a mistake to buy insurance. Finally, even if some agents come to regret entering into trades, those whose beliefs were correct will not regret entering into the same trades. Protecting those who will be proven wrong from trading does not amount to making everyone at least as well off as when they are allowed to trade, so the usual Pareto improvement argument for policy does not apply. Gilboa, Samuelson, and Schmeidler (2014) argue for a compromise of sorts by introducing the notion of *no-betting Pareto improvement*, a refinement on the usual notion of Pareto improvement. Under their notion, an allocation is viewed as superior not only if all agents prefer it to an alternative allocation, but also if there exist some common beliefs—which may be different from the beliefs that any agent holds—such that if all agents maintain these beliefs, all agents are no worse off and some are strictly better off than under the alternative. By this logic, allowing agents to trade on the basis of their heterogeneous beliefs will not be viewed as a Pareto improvement, but preventing them from trading will not be viewed as a Pareto improvement either.

Perhaps the best case for preventing agents with different beliefs from trading is the scenario emphasized in Blume et al. (2014) and Brunnermeier, Simsek, and Xiong (2014) in which trade is a negative-sum game. In this case, society may be better off with restrictions that prevent such trading from taking place before knowing what beliefs any agent might have. This notion may be in line with the emerging view on bubbles in the wake of the Great Recession: The apparent bubble in housing might have left some better off (for example, homeowners and developers who sold houses in the years leading up to the recession and traders such as those profiled in Michael Lewis's book *The Big Short* who managed to short housing[21]) and some worse off (for example, those who bought housing or invested in mortgages just before the recession hit); however, on the whole, society was worse off because of misallocated resources (for example, excess housing and workers whose skills were specific to housing-related activities) that might have depressed subsequent economic activity. Posner and Weyl (2013) are the most forceful in making this case. But there are two important caveats that make this policy prescription difficult to implement in practice. First, the extent to which the investment in an asset (whether it be housing in the mid-2000s or dot-com ventures in the late 1990s or railroads in the 1800s or tulips in the Netherlands in the seventeenth century) is excessive ex ante, before we know how things turn out, hinges not on agents holding different beliefs but on them holding beliefs we know to be distorted. Many would balk at the notion that policymakers can judge when agents hold distorted beliefs and whether agents' beliefs are correct. Referring to models featuring agents with heterogeneous beliefs as models of bubbles can be misleading in that regard, since evidence that people hold different beliefs does not prove that their beliefs are distorted. And yet, distorted beliefs, rather than heterogeneous beliefs, are what imply asset prices are too high. Second, since agents are eager to trade, there is strong incentive for agents to claim they are trading because of fundamental reasons rather than because of differences in their beliefs. Indeed, the response to financial reform in the wake of the financial crisis suggests market participants have actively sought to evade restrictions on when they can trade. Cochrane (2014) makes a similar argument.

An alternative approach: Asymmetric information

I now turn to the other approach for modeling greater-fool theories of bubbles. For lack of a consensus term, I will refer to these as asymmetric information models. This is because a key feature of these theories is that agents receive private information other agents may not be privy to. In particular, they may receive information that all agents would agree establishes that the asset is overvalued. However, since agents are unsure what other agents know, they might still buy the asset in the hope of selling it to a less informed agent. Thus, agents engage in speculative trading because they hope other agents are not privy to the same information as they have, rather than because they think other agents disagree with them. Loosely speaking, they do not expect to profit off of those with whom they trade because their counterparties hold different views, but because their counterparties are less informed. Of course, even agents who begin with different priors may receive asymmetric information. Indeed, an unfortunate source of confusion is that some of the papers on bubbles that feature asymmetric information assume that traders have different prior beliefs, obscuring the differences between these two approaches.

In discussing these asymmetric information models, I find it once again natural to begin with the analysis in Tirole (1982). Recall that his setup allowed individuals to obtain heterogeneous information. To allow for the possibility of bubbles, then, one of Tirole's four conditions must be violated. Theories based on asymmetric information essentially drop the requirement that resources be allocated efficiently before any trades take place (that is, Tirole's fourth condition). If this condition is dropped, then if some trader named Carol offers to sell an asset to some other trader named Ted at what seems to him like a good price, he will not be able to conclude whether the offer has been made because they can both gain from trade or because Carol received information that the asset is worth less than Ted believes it to be. For example, Carol may have immediate liquidity needs and is willing to sell the asset at a price that Ted thinks is a good value. Or Carol may have different hedging needs than Ted, and so both of them will be better off trading, since Carol can then go and purchase another asset that better suits her needs. Ted will of course still be cautious, knowing Carol might have received private information that the asset is not as valuable as he believes and might now be taking advantage of him. But because there is some possibility of gains from the trade, he need not refuse to trade altogether. Note that, as I have essentially already shown, one reason agents may want to trade is that they have different prior beliefs. Hence, one way to violate Tirole's fourth condition (which holds that there is no reason to trade because resources are already allocated efficiently) is to relax his third condition (which requires that agents have common priors). The first papers to construct asymmetric information models of bubbles did in fact just that, since it is relatively easy to analyze models where agents trade because they have different priors. But these models differ in important respects from the models that rely on different priors that I discussed earlier.

The main difference is that with asymmetric information, one can generate bubbles, rather than just speculative trading, when all traders are rational. Recall that speculative trading implies traders expect to profit at the expense of others. This requires them to have different beliefs from others—or else those whom they expect to profit from would refuse to trade. For speculative trading to be sustained, it does not matter whether agents hold different beliefs because they started with different priors or because they receive different information. But the exact reason why agents hold different beliefs does matter for whether one can view an asset as being overvalued. Researchers initially working on asymmetric information in asset markets ignored the question of bubbles, focusing only on the possibility of speculation. For example, Grossman and Stiglitz (1980) were interested in whether agents with common initial beliefs would ever engage in speculative trading. As Grossman (1976) observed, rational agents will try to infer what information others observed given the price at which assets trade, and the market-clearing price can reveal enough information so that all agents are equally informed. In this case, there is no scope for speculation. Grossman and Stiglitz (1980) instead introduced into these models "noise traders"—traders whose trades are motivated by some consideration other than profit maximization, such as liquidity. As their name suggests, a key feature of noise traders is that their impact is random, which makes the price of the asset a noisy signal among all the aggregate information agents receive.[22] But the presence of these traders can also be understood as a way of getting around the Tirole (1982) results that deny the possibility of speculation. Grossman and Stiglitz (1980) did not go on to show that this structure can also get around Tirole's (1982) results that deny the possibility of a bubble; this was established only in subsequent work, which involved an explicit dynamic setting that was missing from previous work on speculative trading.

Bubbles and asymmetric information

The first researchers to show that asymmetric information models can give rise to asset bubbles were Allen, Morris, and Postlewaite (1993). Their analysis was subsequently sharpened and refined by Conlon (2004), who developed a different setup and showed that some of the features of their model were not essential for their results. Although these models are too involved to reproduce here in detail, the basic insight from these papers can be understood as follows.

Suppose that in some state of the world there was information that the price of the asset exceeded the true discounted value of its dividends. For example, Allen, Morris, and Postlewaite (1993) and Conlon (2004) consider a situation in which there is information that the asset pays no dividend, so its fundamental value is zero, and yet the equilibrium price of the asset is positive. Suppose that in these states of the world, every trader receives the information that indicates the dividend is zero.[23] Although each agent knows this information, none of the agents know what other information the other traders have. In particular, consider a setting where after observing the information that dividends are zero, each trader believes two scenarios are possible:

1) All traders know the dividend is zero, so at any positive price, the asset is overvalued.

2) There are some traders who still believe that at a positive price, the asset is worth buying.

That is, the situation under consideration is one in which reality corresponds to the first case, but no agent is sure whether the truth is the first or the second case. Given this uncertainty, a trader (again, call her Carol) can take a gamble and buy the asset at a positive price with the aim of selling it after one period. From Carol's perspective, if it turns out the truth is the first case, she will incur a loss, since she will be unable to sell the asset given all traders know the asset is worthless. If it turns out the truth is the second case, she will be able to sell the asset at a positive price to some other trader (again, call him Ted). This trade will only be profitable if the price rises between when Carol buys it and when she sells it to Ted; the aforementioned papers design an environment in which the equilibrium price rises over time by the requisite amount.

The reason a trader would be willing to buy the asset in the second case is precisely because of the possibility of gains from trade related to dropping Tirole's fourth condition, which holds that there is no reason for agents to trade. That is, in the state of the world where Carol knows the asset is worthless but Ted does not, Ted understands that Carol may be selling him an overvalued asset. But Ted cannot distinguish that state from other states of the world in which Carol would offer to sell him an asset at the exact same price but in which there are mutual gains from trade—for example, because she has a need for liquidity and would be willing to sell the asset for even less than the expected value of its dividends. In short, Carol is willing to buy the asset in period 1 at a price she knows exceeds its fundamental value because she hopes to sell it to Ted for an even higher price in period 2, when he isn't sure if it is overvalued or not. Hence, a bubble can emerge in equilibrium. That is, the price of the asset can be positive even when all traders are aware that the asset is worthless, so long as traders don't know that everyone else realizes the asset is worthless. The possibility of asymmetric information is crucial for why a bubble can arise.

In the preceding paragraphs of this section, I have described a coherent example in which the emergence of a bubble is a logical possibility. However, the bubble in the Allen, Morris, and Postlewaite (1993) model bears no resemblance to the historical episodes people usually have in mind when they talk about bubbles. First, when all agents know the asset is overvalued, its price collapses after one period: As soon as a Carol tries to look for a Ted to sell the asset to, she will immediately learn she cannot find one, and the overvaluation will disappear. Second, the bubble asset in their model never actually changes hands. By contrast, the historical episodes that many have taken to be examples of bubbles involved high trade volumes and periods of prolonged asset price appreciation before prices collapsed, allowing traders to "ride the bubble," or hold on to the asset and let its price appreciate before selling it. Conlon (2004) modifies the model in a way that allows the bubble to be sustained beyond one period and the asset to be traded back and forth between two agents. Essentially, traders keep gambling on the exact date at which it will become common knowledge that the asset is worthless. However, the bubble he constructs remains fragile, in the sense that small perturbations to beliefs or payoffs will lead the bubble to disappear. More recent work has sought to construct robust asymmetric information models of bubbles that persist for several periods. These are sometimes known as models of riding bubbles, since they feature agents who hold assets while they appreciate and then sell them. I discuss some of them next.

Riding an asymmetric information bubble

Abreu and Brunnermeier (2003) were among the first to try to model the phenomenon of riding a bubble. In their model, agents are sequentially informed that an asset is overvalued from some randomly chosen date t_0. However, no agent observes t_0. Thus, each agent learns that the asset is overvalued, but not how many others know the asset is overvalued or how long they have known this. Abreu and Brunnermeier assume the price of the asset rises over time, just as it does in Allen, Morris, and Postlewaite (1993) and Conlon (2004). Hence, if a trader is among the first to learn the asset is overvalued and the first to sell, he will make a profit. If he is among the last to know and among the last to sell, he will be unable to find a buyer by the time he acts. Abreu and Brunnermeier (2003) show that under additional assumptions, the optimal strategy for a trader is to wait a fixed period of time from when he learns the asset is overvalued and then sell. Depending on the pace at which the asset price grows, the rate at which agents discount, and the distribution of t_0, a trader may wait to sell for longer than it takes all agents to learn the asset is overvalued. Thus, there can be a situation where every agent knows that the asset is overvalued, yet the asset continues to trade at a price that exceeds its fundamental value, just as in Allen, Morris, and Postlewaite (1993) and in Conlon (2004).

Unfortunately, the analysis in Abreu and Brunnermeier (2003) shows only that agents will engage in speculative trading if a bubble exists. But their work does not prove that a bubble can in fact exist. However, subsequent work by Doblas-Madrid (2012) shows that it is possible to construct an internally consistent model of a bubble that exhibits many of the features of the Abreu and Brunnermeier model. Doblas-Madrid's analysis offers several insights. First, contrary to some of the suggestions in Abreu and Brunnermeier (2003), he shows that it is not necessary to assume that some agents hold exotic beliefs to sustain a bubble. This should not be surprising given my earlier observation that a bubble may arise even if agents are rational, as long as there is some reason for them to trade. Indeed, Doblas-Madrid assumes that in every period there will be some traders who require immediate liquidity, so there can be gains from trade between agents with a pressing need for liquidity and those willing to hold the asset. For this explanation to hang together, he needs to impose a limit on how many units of the asset buyers can absorb. This allows the asset price to remain below the present discounted value of earnings, so those who buy the asset are strictly better off. In particular, Doblas-Madrid assumes traders cannot borrow, so their demand for the asset is constrained by their income.[24] Another issue Doblas-Madrid explores is under what condition the bubble will persist even after the first cohort of traders sell their asset holdings, so that the downward pressure on prices the first cohort exert when they sell their assets doesn't tip off other agents that some traders have started to sell their assets. In particular, Doblas-Madrid shows that some source of randomness is necessary so that prices can fall even when the first traders to learn the asset is overvalued sell without alerting other agents. Thus, sustaining trade in an overvalued asset requires more uncertainty than assumed in the Abreu and Brunnermeier (2003) setup. This feature is certainly plausible (the real world is complicated and features many sources of uncertainty), but it suggests that asymmetric information models of bubbles in which agents trade the asset are likely to be fairly complicated. Finally, Doblas-Madrid (2012) shows that to sustain a bubble, his model requires certain restrictions on the way agents can trade. Intuitively, a greater-fool theory can only work if each trader expects he might profit from selling the asset to a greater fool. The problem is that when agents sell, they reveal to everyone that the asset is overvalued. If this revelation scares away buyers, it will be impossible to profit from selling assets. Doblas-Madrid gets around this by assuming agents must submit their orders before they know the price of the asset. In finance parlance, this means agents are only allowed to place *market orders*, which dictate how much to buy or sell at the market price, but they cannot place *limit orders*, which restrict the range at which a trade will be executed (for example, an order that says to only buy an asset if its price is below some cutoff). Whether it is possible to sustain bubbles when traders are unrestricted in the orders they can place remains an open question.[25]

Back to policy

To recap, economists have been able to construct models of bubbles based on asymmetric information in which the price of an asset exceeds what one can objectively argue the asset is worth. Recall that the models that feature agents with different prior beliefs that I discussed earlier could also give rise to scenarios that can be described as bubbles, but only if traders have distorted beliefs so that the most optimistic beliefs tend to be wrong. In that case, policy intervention is justified only if policymakers know that agents' beliefs are erroneous. By contrast, in asymmetric information models of bubbles, all agents know the asset is overvalued, so their beliefs are not erroneous. Instead, it is only because traders are uncertain as to what others know that they are willing to buy assets and gamble that they can sell them to others who are less informed than they are. The question is whether letting agents gamble this way is undesirable—and, more generally, whether the fact that asset prices can exceed the fundamental value of the asset is socially costly. Unfortunately, little work has been done to analyze these issues in models of asymmetric information.

An important exception is Conlon (2015), who studied the role of policy in an asymmetric information model of bubbles along the lines of his earlier paper (Conlon, 2004). Specifically, he assumed the policymaker also receives information that the asset is trading at a price above its fundamental value, and can announce this information publicly. If the policymaker were to make such an announcement, he would eliminate the prospect of exploiting less informed traders. To be sure, this is not the policy response that advocates of more forceful action against potential bubbles have in mind. They typically argue that a central bank should raise interest rates to head off possible bubbles. That said, Conlon's thought experiment is still informative, since it reveals the social welfare consequences of deflating a bubble when it can be achieved costlessly.[26]

To understand Conlon's (2015) results, it will be helpful to return to the key intuition behind bubbles in asymmetric information environments: Traders are willing to buy an asset they know to be overvalued because they are taking a gamble. Either they will be able to sell it at an even higher price to another less informed trader, or else they will find out that no other trader is willing to buy the asset and they will incur a loss. The reason a trader may be able to sell the asset to a greater fool is that the buyer believes he may be entering into a mutually beneficial trade. Thus, an inherent feature of greater-fool bubbles based on asymmetric information is that when agents trade, sometimes it is because assets are overvalued and sometimes it is because there are mutual gains from trade. In other words, Carol can profit at Ted's expense only because there are other situations in which both Carol and Ted gain from trading and Ted doesn't know which state they are in while Carol does. If a policymaker were to reveal that the asset is overvalued, this information would affect the price in both scenarios. In particular, it would lead to a reduction in the price when the asset is overvalued, and it would lead to an increase in the price when there are mutual gains from trade. The first part is straightforward: By telling everyone the asset is overvalued, the policymaker prevents those who know the asset is overvalued (for example, Carol) from passing it off to less informed traders (for example, Ted), and so the price of the asset will not exceed the fundamental value. As for the second part, in the state of the world where there are mutual gains from trade between Carol and Ted, one should note that when Ted buys the asset he remains nervous that Carol might be taking advantage of him. If this concern were mitigated, he would be willing to pay more for the asset, and the price would be higher.

When the asset is available in fixed supply, announcing a bubble will generally have an ambiguous effect on social welfare. Given that the price of the asset rises in some states of the world and falls in others, those who sell the asset will be better off in some states but worse off in others. We can abstract from these considerations by assuming that the gains and losses exactly cancel each other out. In this case, a commitment by a fully informed policymaker to announce whenever she knows the asset is overvalued will have no effect on welfare when the asset is in fixed supply. But a commitment by the policymaker to reveal when she knows there is a bubble could still improve welfare, even when these two effects exactly offset each other, if the asset were in variable supply. This statement holds true because the way the price of the asset changes in different states of the world leads to fewer units of the asset being created when there are no gains from trading it and more units of the asset being created when there are gains from trading it. This is reminiscent of the welfare results in the case of traders with uncommon priors: When an asset is in fixed supply, the case for preventing agents from trading is ambiguous, but when an asset is in variable supply, there can be welfare gains from reducing the cost of resource misallocation due to mispricing.

Conlon (2015) goes on to show that the case for policy intervention against a bubble crucially hinges on the policymaker being able to identify a bubble whenever it arises. His argument is different from the more conventional logic that allowing policymakers to act against bubbles can be costly if they mistakenly act thinking an asset might be a bubble when in fact it is not.[27] Conlon shows that even if policymakers are conservative and only react when they are certain there is a bubble, policy intervention may make agents worse off. This can happen because if policymakers deflate bubbles in some states of the world but not others, the bubbles that remain in other states can be worse than the ones that policymakers actually lean against. Consequently, the resulting misallocation of resources from policy intervention can be exacerbated. The case for intervention may therefore rest on a policymaker being perfectly informed about bubbles, since responding either too aggressively or too timidly may undercut the case for intervention. That is, while asymmetric information models of bubbles suggest intervention can be helpful, they also highlight the difficulty of justifying intervention in practice.

I conclude my discussion of policy implications with one final observation. Arguably, the primary reason policymakers cite for why they are concerned about bubbles is distinct from those I discussed in this article. The justifications for intervention I have discussed so far involve preventing a glut of assets in cases where the assets are overvalued. But the case for policy intervention has tended to focus on the dire consequences of the bubble bursting rather that resource misallocation while asset prices are too high. Here, it is worth pointing out that the models of bubbles based on asymmetric information I have described imply that if a bubble arises, it will eventually burst. This is because a bubble corresponds to a scenario in which all agents believe the asset is worth less than its price, yet they are willing to buy it because they are unsure whether other traders are aware of this. Eventually, uncertainty about what other traders know is resolved, at which point the price of the asset collapses. This is not true for other theories of bubbles. For example, in models where bubbles arise because of asset shortages, bubbles can in principle persist indefinitely. In models where bubbles arise because of risk shifting, agents who buy the asset are gambling on a risky asset that sometimes pays off. If that happens, the price of the asset will rise further rather than collapse. The fact that greater-fool theories of bubbles based on asymmetric information imply that bubbles necessarily burst makes them of natural interest for further study, especially to determine whether merely avoiding an eventual asset price collapse can justify policy intervention.

Conclusion

This article described the literature on greater-fool theories of bubbles, that is, theories in which agents are willing to buy assets they know to be overvalued because they believe they can profit from selling the assets to others. The idea behind this theory is intuitive and seems to capture aspects of what often happens during real episodes that are suspected to be bubbles. This theory can also capture the unsustainable nature of a bubble that makes asset bubbles a concern for policymakers. And yet it turns out to be a surprisingly difficult theory to model and analyze.

What specific lessons should be taken away from this discussion? In this article, I highlight two distinctions that are important to keep in mind to better sort through the various results in the existing literature. The first is the distinction between speculation and bubbles. Speculative trading concerns why agents trade—namely, to profit at the expense of others as opposed to intending to find mutually beneficial gains. Asset bubbles concern features of an equilibrium price—namely, whether the price faithfully represents what the asset is fundamentally worth. The fact that agents engage in speculative trading does not necessarily imply that the asset must be a bubble. In line with this, some models that try to capture greater-fool theories of bubbles are really models of the greater-fool theory of trading rather than models of bubbles per se.

The second distinction that this article highlights is one between models based on uncommon priors and those based on asymmetric information. Any greater-fool theory requires that traders hold different beliefs. But it matters whether these different beliefs arise because traders start out with distinct priors or because they receive different information. This difference is reminiscent of the line from Shakespeare's *As You Like It*: "The fool doth think he is wise, but the wise man knows himself to be a fool."[28] In models that feature only uncommon priors, agents are willing to trade because they are convinced their beliefs are correct, even though they cannot all be right—they are like the fool who thinks he and only he is wise. In models that feature asymmetric information, traders are aware that those they are trading with receive private information that may allow them to exploit the other traders around them. While agents are willing to trade, they are also cautious about being exploited—they are wise to the fact that they might be the greater fool, although ultimately they are willing to trade. The two types of models are thus quite different. Without additional restrictions on how beliefs are formed, models based on uncommon priors arguably do not generate asset bubbles, while models based on asymmetric information can. For the same reason, the policy implications of the two types of models are not identical.

Although some of the work that assumes different priors does explicitly refer to asset bubbles, the term *bubble* can be appropriate if beliefs are assumed to be somehow distorted. The general insight from these models is that asset prices are determined by the most optimistic traders, in the same way that the price of a good is determined by those who most prefer that good. But without a theory of how beliefs are formed, there is nothing that tells us that the most optimistic person must be wrong or that the average belief is correct. The heterogeneity of agents' beliefs on its own is thus not a good basis for talking about bubbles or for arguing that policymakers should drive asset prices to their correct values. Still, models with uncommon priors are easy to work with, and so it may be appealing to use them together with some restriction on how beliefs are formed that implies excessive optimism. Asymmetric information models can generate bubbles without such restrictions, but they rely on a lot of structure to make sure information isn't somehow revealed through prices or actions that traders take. Which framework is better depends on the particular application one is interested in.

Finally, with regard to policy implications, my discussion highlights various difficulties in using greater-fool theories of bubbles to justify action against potential bubbles. Although these theories can provide some justifications for why policymakers should intervene, these rationales come with many caveats. For example, policymakers may have to know that traders have incorrect beliefs, even though policymakers would not necessarily be any better at forecasting future dividends than members of the private sector. Other justifications for intervention require policymakers to be perfectly attuned to when bubbles arise—a condition that seems implausible in practice. In fact, greater-fool theories of bubbles naturally suggest the opposite, that is, that detecting bubbles is likely to be difficult. Recall that in asymmetric information models, bubbles can arise only because there is the possibility of mutual gains from trade. Thus, there may be plausible reasons for why agents trade assets beyond trying to benefit at the expense of others. Finally, the social welfare implications that emerge most clearly in these models do not seem to capture the main issue policymakers are concerned with in regard to bubbles. For example, those who argue for a more forceful policy response to potential bubbles typically expect this response to come from central banks. This reflects a view that bubbles are fueled by loose credit conditions, as well as the idea that the collapse of a bubble causes the most harm when assets were purchased on leverage and a collapse in their price would trigger a subsequent round of defaults. Yet in most models of the greater-fool theory of bubbles, credit plays only a minor role or is missing altogether. As I discuss in Barlevy (2012), risk-shifting theories of bubbles seem particularly well suited for exploring these issues. However, introducing credit into models of the greater-fool theory of bubbles, which some have attempted to do, may help tackle these issues as well.

NOTES

[1]Note that this logic concerns only how policymakers should respond to evidence of a possible bubble. In principle, though, policy intervention might prevent bubbles from arising in the first place. Indeed, some have argued that policies such as restricting how much agents can borrow against an asset or taxing transactions to make trading less profitable may prevent bubbles. These policies are also more targeted than the interest rate rules Bernanke and Gertler (1999) considered.

[2]Allen, Morris, and Postlewaite (1993) provide a clear discussion of why defining a bubble when agents have different beliefs can be difficult. Rather than attempt to provide a general definition for a bubble, they argue for constructing specific circumstances in which there is enough structure to argue that the price of the asset deviates from its fundamental value.

[3]There is an extensive literature on bubbles of this type. The classic reference on bubbles that arise when agents need assets to serve as a store of value is Tirole (1985), who builds on the original work of Samuelson (1958). Caballero and Krishnamurthy (2006) and Farhi and Tirole (2012) discuss the case where agents need assets to serve a liquidity role.

[4]For more on the risk-shifting theory of bubbles, see Allen and Gorton (1993), Allen and Gale (2000), and Barlevy (2014).

[5]It is not clear where the term *greater fool* originated, but it seems to have been first used by market practitioners. For example, a discussion of potential broker-dealer misconduct in the Securities and Exchange Commission's annual report for the fiscal year ending June 30, 1963, contains the following description: "What has been colloquially referred to as the 'bigger fool' theory ... is simply the assurance that regardless of whether the price paid for a security is fair and/or reflective of the intrinsic value of the security or even reflective of a rational public evaluation of the security, the security is still a good buy because a 'bigger fool' will always come along to take it off the customer's hands at a higher price" (Securities and Exchange Commission, 1964, p. 74).

[6]Chancellor (1999).

[7]Milgrom and Stokey (1982) independently established similar results to those in Tirole (1982), although they framed their findings in terms of speculative trading rather than bubbles. I therefore refer to Tirole's work in my discussion.

[8]The condition that the number of traders be finite, which I have ignored in my discussion, also plays an important role in ruling out the possibility of a bubble. Even if each trader understands that the trader he buys from is profiting at his expense, he might be willing to buy the asset if he thinks there is another trader at whose expense he can profit. If one never runs out of traders to exploit, it may be possible to sustain such trading chains. An important step in Tirole's (1982) argument is to show that this is not possible when the number of potential traders is finite.

[9]Harrison and Kreps (1978) define speculation differently than I do, using the term to refer to a situation in which traders assign positive value to the right to resell an asset. The problem with their definition is that it implies that finitely lived agents who buy infinitely lived assets are speculators, even when they do not expect to profit at the expense of younger cohorts that buy assets from them. This distinction was irrelevant for Harrison and Kreps, who assumed infinitely lived agents for their model's environment, but their definition may not generalize well to other environments.

[10]Morris (1995) discusses the common prior assumption and its connection to rationality. As he notes, some have argued that since at least one agent must be wrong whenever two agents hold different beliefs, that agent must not be rational. But as Morris notes, rationality restricts only how agents update their priors, not what their priors can be.

[11]For example, suppose dividends reflect the profits of a multiproduct company that sold different products at each date. The fact that a theory about how much profit the firm would earn selling apples in Australia in period 1 was wrong may not lead us to revise our theories about how much profit the firm would earn selling bicycles in Burundi in period 2.

[12]The usual motivation for assuming agents are price takers is that it can always be assumed there are many identical replicas of Evelyn and Odelia in the market, in which case the actions of any one agent have no influence on the price of the asset.

[13]Harrison and Kreps (1978) show that $p_t = \frac{1}{1-\beta} + c\beta^{-t}$ for $c > 0$

can also be an equilibrium price path. If Evelyn's and Odelia's endowments do not grow, at some point the one who values the asset more could not afford the asset, yet the other party would want to sell all of her holdings. Hence, in this case, such a path cannot be an equilibrium. Harrison and Kreps refer to $c\beta^{-t}$ as a bubble, although they use this term in the sense of an explosive solution of a difference equation rather than the way I use the term. Still, their terminology

suggests $\frac{1}{1-\beta}$ behaves somewhat like a fundamental value—a

theme they pick up on in their paper.

[14]In particular, one can introduce a mechanism similar to the one in Zeira (1999). Suppose that both agents believe that positive dividends grow at a constant rate until some random date T, where the distribution of date T is known to both parties and has unbounded support. At date T, both traders Evelyn and Odelia agree that dividends will cease growing thereafter, even if they disagree on when they will be paid out. As long as dividends continue to grow, asset prices will rise faster than the risk-free interest rate. Evelyn and Odelia will therefore sell the asset at a higher price than they paid to buy it. Moreover, at date T the price will crash, so the model admits both a boom and a crash. Note that speculative trading would continue beyond date T even when asset prices stop growing unless one assumes differences in beliefs also disappear at date T.

[15]A modern-day example is repurchase agreements (repos), under which a security is sold with the promise that the seller will buy it back. This can be viewed as effectively renting the asset, although legally repos do transfer ownership of the asset.

[16]Formally, the dividend corresponds to a binomial tree. Such a process is often used in models with heterogeneous beliefs among agents because of its tractability; see, for instance, the example in section III.C of Brunnermeier, Simsek, and Xiong (2014).

[17]Pareto efficiency is a standard criterion for evaluating policies in economics. Stiglitz (1989, p. 113, note 3) offers the usual definition for this term: "An economy is Pareto efficient if no one can be made better off without making someone else worse off."

[18]Kreps (2012) offers another example where agents hold different beliefs and trade is costly. Kreps (2012, p. 193) describes a bet between two economists, Joe Stiglitz and Bob Wilson, over the contents of a pillow. Each is willing to bet a small sum of money that his belief is right. However, to prove which one is correct, Stiglitz and Wilson must destroy the pillow and purchase a new one. The destroyed pillow is the social cost associated with trade.

[19]In line with the definition of Pareto efficiency described in note 17, a Pareto improvement is a change that harms no one and helps at least one person.

[20]For example, suppose there are two individuals—one who believes there is a probability of 0.7 that a tossed coin will come up heads rather than tails and the other who believes there is a probability of 0.3 it will come up heads. Achieving belief-neutral Pareto improvement requires evaluating both of their utilities with the same probability p for every p between 0.3 and 0.7. If one allocation makes both agents at least as well off as another allocation at each probability, that allocation is said to represent a belief-neutral Pareto improvement.

[21]Lewis (2010).

[22]Formally, the presence of noise traders prevents the price from being an invertible function of aggregate information.

[23]Allen, Morris, and Postlewaite (1993) refer to this case as a strong bubble. They contrast this with the case where the asset trades at a positive price but not all agents know the dividend is zero.

[24]Doblas-Madrid and Lansing (2014) consider a variation of the Doblas-Madrid (2012) model in which agents can borrow. As long as there are constraints on how much agents can borrow, it will still be possible for agents to profit from buying the asset from those with liquidity needs. The rate at which the price of the asset grows will now be tied to the rate at which credit is growing.

[25]Given my earlier remark on Zeira (1999), it is worth noting that the Doblas-Madrid (2012) model achieves a boom and bust using a similar structure to that of Zeira's model. Recall that Zeira assumes dividends grow until some random date T. At date $T + 1$, all agents learn that dividends will no longer grow and the price crashes. In Doblas-Madrid's setup, there is a date t_0 at which the fundamental value ceases to grow, but not all agents learn this at date t_0 because they receive information sequentially. The crash that would ordinarily happen when fundamentals cease growing is thus delayed. To avoid having agents learn the fundamentals from dividends, Doblas-Madrid assumes dividends are paid out far off in the future. But conceptually the setup is similar to Zeira's.

[26]Conlon's (2015) notion of relying on announcements also has a historical precedent. In December 1996, Alan Greenspan in his capacity as Federal Reserve Chairman publicly expressed his concern about "irrational exuberance" in financial markets, which some have interpreted as an attempt to affect the beliefs of market participants (Greenspan, 1996).

[27]See Cogley (1999) for an example of this type of argument.

[28]See http://www.shakespeare-online.com/plays/asu_5_1.html (act 5, scene 1, lines 30–31).

REFERENCES

Abreu, Dilip, and Markus K. Brunnermeier, 2003, "Bubbles and crashes," *Econometrica*, Vol. 71, No. 1, January, pp. 173–204.

Allen, Franklin, and Douglas Gale, 2000, "Bubbles and crises," *Economic Journal*, Vol. 110, No. 460, January, pp. 236–255.

Allen, Franklin, and Gary Gorton, 1993, "Churning bubbles," *Review of Economic Studies*, Vol. 60, No. 4, October, pp. 813–836.

Allen, Franklin, Stephen Morris, and Andrew Postlewaite, 1993, "Finite bubbles with short sale constraints and asymmetric information," *Journal of Economic Theory*, Vol. 61, No. 2, December, pp. 206–229.

Barlevy, Gadi, 2014, "A leverage-based model of speculative bubbles," *Journal of Economic Theory*, Vol. 153, September, pp. 459–505.

__________, 2012, "Rethinking theoretical models of bubbles: Reflections inspired by the financial crisis and Allen and Gorton's paper 'Churning bubbles,'" in *New Perspectives on Asset Price Bubbles: Theory, Evidence, and Policy*, Douglas D. Evanoff, George G. Kaufman, and A. G. Malliaris (eds.), New York: Oxford University Press, pp. 41–61.

Bernanke, Ben, and Mark Gertler, 1999, "Monetary policy and asset price volatility," *Economic Review*, Federal Reserve Bank of Kansas City, Fourth Quarter, pp. 17–51.

Blume, Lawrence E., Timothy Cogley, David A. Easley, Thomas J. Sargent, and Viktor Tsyrennikov, 2014, "The case for incomplete markets," New York University and Cornell University, mimeo, February 27, available at https://easley.economics.cornell.edu/BECST-20140226.pdf.

Brunnermeier, Markus K., Alp Simsek, and Wei Xiong, 2014, "A welfare criterion for models with distorted beliefs," *Quarterly Journal of Economics*, Vol. 129, No. 4, November, pp. 1753–1797.

Caballero, Ricardo J., and Arvind Krishnamurthy, 2006, "Bubbles and capital flow volatility: Causes and risk management," *Journal of Monetary Economics*, Vol. 53, No. 1, January, pp. 35–53.

Chancellor, Edward, 1999, *Devil Take the Hindmost: A History of Financial Speculation*, New York: Farrar, Straus and Giroux.

Cochrane, John H., 2014, "Challenges for cost-benefit analysis of financial regulation," *Journal of Legal Studies*, Vol. 43, No. S2, June, pp. S63–S105.

Cogley, Timothy, 1999, "Should the Fed take deliberate steps to deflate asset price bubbles?," *Economic Review*, Federal Reserve Bank of San Francisco, No. 1, pp. 42–52.

Conlon, John R., 2015, "Should central banks burst bubbles? Some microeconomic issues," *Economic Journal*, Vol. 125, No. 582, February, pp. 141–161.

__________, 2004, "Simple finite horizon bubbles robust to higher order knowledge," *Econometrica*, Vol. 72, No. 3, May, pp. 927–936.

Doblas-Madrid, Antonio, 2012, "A robust model of bubbles with multidimensional uncertainty," *Econometrica*, Vol. 80, No. 5, September, pp. 1845–1893.

Doblas-Madrid, Antonio, and Kevin Lansing, 2014, "Credit-fuelled bubbles," Michigan State University and Federal Reserve Bank of San Francisco, mimeo, May, available at http://www.bostonfed.org/Macro-Finance-2014/Doblas-Madrid-Lansing.pdf.

Farhi, Emmanuel, and Jean Tirole, 2012, "Bubbly liquidity," *Review of Economic Studies*, Vol. 79, No. 2, April, pp. 678–706.

Gilboa, Itzhak, Larry Samuelson, and David Schmeidler, 2014, "No-betting-Pareto dominance," *Econometrica*, Vol. 82, No. 4, July, pp. 1405–1442.

Greenspan, Alan, 1996, "The challenge of central banking in a democratic society," remarks by Federal Reserve Chairman at Annual Dinner and Francis Boyer Lecture of the American Enterprise Institute for Public Policy Research, Washington, DC, December 5, available at http://www.federalreserve.gov/boarddocs/speeches/1996/19961205.htm.

Grossman, Sanford, 1976, "On the efficiency of competitive stock markets where trades have diverse information," *Journal of Finance*, Vol. 31, No. 2, May, pp. 573–585.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, "On the impossibility of informationally efficient markets," *American Economic Review*, Vol. 70, No. 3, June, pp. 393–408.

Harrison, J. Michael, and David M. Kreps, 1978, "Speculative investor behavior in a stock market with heterogeneous expectations," *Quarterly Journal of Economics*, Vol. 92, No. 2, May, pp. 323–336.

Kreps, David M., 2012, *Microeconomic Foundations I: Choice and Competitive Markets*, Princeton, NJ: Princeton University Press.

Lewis, Michael, 2010, *The Big Short: Inside the Doomsday Machine*, New York: Norton.

Milgrom, Paul, and Nancy Stokey, 1982, "Information, trade and common knowledge," *Journal of Economic Theory*, Vol. 26, No. 1, February, pp. 17–27.

Miller, Edward M., 1977, "Risk, uncertainty, and divergence of opinion," *Journal of Finance*, Vol. 32, No. 4, September, pp. 1151–1168.

Mongin, Philippe, 2005, "Spurious unanimity and the Pareto principle," London School of Economics, Centre for Philosophy of Natural and Social Science (CPNSS), LSE Choice Group, working paper, Vol. 1, No. 5, available at http://eprints.lse.ac.uk/23942/.

Morris, Stephen, 1995, "The common prior assumption in economic theory," *Economics and Philosophy*, Vol. 11, No. 2, October, pp. 227–253.

Posner, Eric A., and E. Glen Weyl, 2013, "An FDA for financial innovation: Applying the insurable interest doctrine to twenty-first-century financial markets," *Northwestern University Law Review*, Vol. 107, No. 3, Spring, pp. 1307–1358.

Samuelson, Paul A., 1958, "An exact consumption-loan model of interest with or without the social contrivance of money," *Journal of Political Economy*, Vol. 66, No. 6, December, pp. 467–482.

Scheinkman, José A., 2014, *Speculation, Trading, and Bubbles*, New York: Columbia University Press.

Scheinkman, José A., and Wei Xiong, 2003, "Overconfidence and speculative bubbles," *Journal of Political Economy*, Vol. 111, No. 6, December, pp. 1183–1219.

Securities and Exchange Commission, 1964, *29th Annual Report of the Securities and Exchange Commission: For the Fiscal Year Ended June 30, 1963*, Washington, DC: U.S. Government Printing Office.

Stiglitz, Joseph E., 1989, "Using tax policy to curb speculative short-term trading," *Journal of Financial Services Research,* Vol. 3, Nos. 2–3, December, pp. 101–115.

Tirole, Jean, 1985, "Asset bubbles and overlapping generations," *Econometrica*, Vol. 53, No. 6, November, pp. 1499–1528.

__________, 1982, "On the possibility of speculation under rational expectations," *Econometrica*, Vol. 50, No. 5, September, pp. 1163–1181.

Velde, François R., 2013, "Lottery loans in the eighteenth century," Federal Reserve Bank of Chicago, mimeo.

Zeira, Joseph, 1999, "Informational overshooting, booms, and crashes," *Journal of Monetary Economics*, Vol. 43, No. 1, February, pp. 237–257.