Economic theory and asset bubbles

Gadi Barlevy

Introduction and summary

In the late 1990s, there was a dramatic run-up in the stock prices of various publicly traded U.S. firms, particularly those specializing in information technology and its applications. A similar run-up occurred in housing prices in the early 2000s. These episodes have often been cited by pundits as examples of "asset price bubbles." Implicit in this description is the notion that the rapid growth in the prices of the assets in question over such a short period suggests these assets were overvalued and thus at risk of a "correction," or a sharp fall in price. In the case of the stock market, prices did in fact decline in 2000, with some indexes falling quite precipitously, as shown in figure 1 for the NASDAQ-100.[1] The slowdown in the growth of housing prices starting in the summer of 2006 has been offered by some as evidence of the beginning of a similar decline in the housing market.

Some of those who have identified these respective episodes as asset bubbles have also warned of the dangers in letting asset prices rise so rapidly. For example, an editorial in *The Economist* magazine argued that "the risk is not just that asset prices can go swiftly into reverse. As with traditional inflation, surging asset prices also distort price signals and so can cause a misallocation of resources—encouraging too little saving, for example, or too much investment in housing" (Economist Newspaper Limited, 2005). Quite naturally, this has led to calls for policymakers to rein in the prices of assets whose values appear to be inflated before they rise to astounding heights. Yet a number of studies have argued that policies designed to contain asset bubbles might end up doing more harm than good. One problem, highlighted in Cogley (1999), is that it is often difficult to assess whether an asset is overvalued, and so there is a risk of deflating assets that are not overvalued, possibly at great harm. Bernanke and Gertler (1999) raise a different issue, noting that using monetary policy to rein in asset prices may lead to unwanted deviations of inflation from its optimal path. They argue that monetary policy should focus exclusively on inflation, not asset prices.

In this article, I argue that, even absent the practical issues that other researchers have pointed out, policymakers would be wise to proceed with great caution before acting to stem asset bubbles. The reason is that economic theory suggests bubbles can only occur under certain circumstances, and these circumstances may have relevance for whether acting to burst bubbles is desirable. For example, one set of circumstances in which bubbles may emerge is if markets are already distorted in other ways in the absence of a bubble. As is well understood by now, policies that are desirable when markets are efficient may turn out to be undesirable when some markets are already distorted—a result dubbed by Lipsey and Lancaster (1956–57) as the "theory of the second best." Even if policymakers could accurately identify an asset bubble and rein in its price without affecting other economic variables such as inflation, there remains the possibility that intervention by policymakers to rein in asset bubbles can exacerbate the very distortions that allowed the bubble to emerge in the first place.

I first discuss the way most economists would define asset bubbles and distinguish it from the way the term is often invoked in the popular press. I then summarize what economic theory tells us about when bubbles thus defined can occur. Lastly, I review some examples of economic models in which bubbles can occur and discuss the policy implications that emerge from such examples. Although the examples I review

Gadi Barlevy is a senior economist and economic advisor in the Economic Research Department at the Federal Reserve Bank of Chicago. *The author thanks Marco Bassetto and François Velde for their helpful comments.*

are quite specialized and sometimes bear little resemblance to the type of episodes that originally prompted the concern about asset bubbles, they help to convey the reasons why bursting bubbles may not always make society as a whole better off and why, in some cases, such policies may actually make society worse off by exacerbating the distortions that cause the bubble.

What is an asset bubble?

As noted in the introduction, the popular press often uses the term bubble to describe a situation in which the price of an asset has increased significantly in such a short period of time so as to suggest that the price is susceptible to an equally sudden collapse. Academic economists have occasionally invoked this definition as well. For example, Kindleberger (1996, p. 13) defines a bubble as "an upward price movement over an extended range that then implodes." The rise and fall of the NASDAQ-100 Index in figure 1 would certainly qualify as a bubble under this definition. Another episode that fits this description is the U.S. stock market in the 1920s. As illustrated in figure 2, the Dow Jones Industrial Average increased by almost 500 percent over the course of roughly eight years, from a low of 63.9 in August 1921 to a peak of 381.2 in September 1929. Then, in just two days—October 28 and 29, 1929—the Dow lost 24.5 percent of its value. The index continued to tumble over the next three years, reaching a nadir of 41.2 in July 1932 from which it took 22 years to climb back up to its peak level in 1929. Other commonly cited examples of bubbles based on this definition include the rise and fall in the price of tulip bulbs in seventeenth-century Netherlands (an episode often referred to as "Tulipmania") and the rapid rise and fall in the price of shares in the South Sea Company and the Mississippi Company in the early eighteenth century.

However, most economists would find this definition to be problematic. Aside from the fact that it is inherently imprecise—the definition is ambiguous about how much the price of an asset must rise, or how quickly, in order to qualify as a bubble—large price swings sometimes occur naturally and benignly in response to shifts in supply and demand. For example, new fashion accessories sometimes surge in price shortly after they are introduced, but eventually their prices fall. This pattern can be understood as follows. When a new fashion line is introduced, its supply will be limited because producers do not know which particular items will catch on and do not produce much of any given article. As the prices of the most coveted items begin to rise, producers respond by supplying more of them, causing their prices to fall. Eventually, consumers move on to new fashions, and the prices of the original line fall even further. While such prices would look similar to figures 1 and 2, most economists would view such price swings as welcome and desirable rather than as cause for concern. The initial high price of particular fashion lines signals to apparel manufacturers that they should produce more of them, and the subsequent low price signals that they should shift to producing other styles. By contrast, economists do view dramatic increases in the price of assets as a potential cause for concern. This is because the increase in price is sometimes so rapid that it seems unlikely to reflect real changes in the true value of the underlying asset. Instead of signaling a genuine need for more such assets, as in the case of fashion accessories where there is greater demand for one item as opposed to others, a bubble conveys the false impression that an asset is especially valuable when in fact it is not. Economists have attempted to come up with a definition of bubbles that focuses on this aspect and thus distinguishes this phenomenon from scenarios in which the price of a good or asset fluctuates because of cyclical changes in its inherent desirability or worth. The latter phenomena

FIGURE 1

NASDAQ-100, 1995–2005

Note: NASDAQ-100 refers to the National Association of Securities Dealers Automated Quotation system's 100 largest nonfinancial securities, based on market capitalization.
Source: NASDAQ (National Association of Securities Dealers Automated Quotation system).

are sometimes labeled "fads" instead (see, for example, Camerer, 1989).

In particular, most economists would define a bubble as a situation where an asset's price exceeds the "fundamental" value of the asset.[2] Formally, the fundamental value is the expected value of all dividends the asset yields over its lifetime, properly discounted to reflect the present-day value of dividends paid at future dates. Intuitively, an asset represents a claim to a stream of future payments. From society's point of view, the value of an asset lies in the fact that it provides resources, in the form of dividends, to the various people who will own the asset over time. If the price of the asset exceeds this value, it can be legitimately viewed as overvalued, since it would not be worth it for society as a whole to pay such a price to create an asset on the collective behalf of all future owners of the asset. In many cases, the fundamental value of the asset is straightforward to define. However, there are some circumstances in which it is not possible to properly assign a unique fundamental value to an asset, although it might still be possible to describe an asset as overvalued. I discuss this in more detail in box 1. Note that this definition does not presuppose that there is a run-up in the price of the asset. However, a rapid increase in the price of an asset may be a telling indicator that an asset is overvalued, especially if there are no concurrently large changes in dividends or in news about future dividends. This definition also does not presuppose that the price of the asset must eventually crash. However, the fact that the price of an asset exceeds its true worth suggests there is nothing to justify its high price, and so the price is always liable to collapse. To put it another way, there is always an equilibrium in which the price of the asset suddenly falls to its fundamental value. Thus, economists distinguish an asset bubble not by what actually happens to the asset's price, but by what *could* happen to it.

It is worth noting that for the price of an asset to exceed its fundamental value, astute investors who buy the asset should believe they might be able to sell it at some point in the future. This is because a forward-looking investor would never purchase an asset worth more than its fundamental value with the intent of keeping it forever. If he did, he would be paying more for the asset than it would ever earn him in dividends. This feature will be important for understanding why bubbles can be ruled out in many circumstances; as we shall see, it is only under special circumstances that all investors can credibly believe they will eventually be able to sell the asset to someone else.

FIGURE 2

Dow Jones Industrial Average, 1920–41

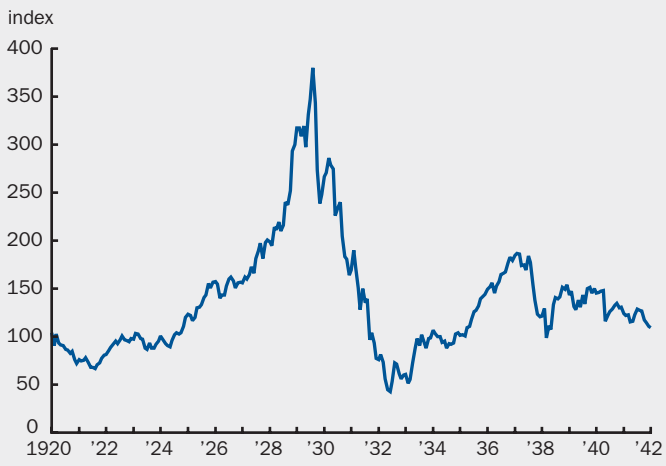

Source: Standard and Poor's.

Although this definition of an asset bubble is simple enough, using it to identify bubbles in practice is difficult. This is because measuring the fundamental value of an asset is hard, especially since it depends on the expectations of dividends that have yet to be realized. Indeed, a coherent case has been made for each of the historical episodes described previously that the asset in question was not truly overvalued.[3] Are there any examples of bubbles that clearly meet this definition? Some economists would argue that money, or cash, should be viewed as a bubble. While not all economists agree with this view, and bubbles associated with money are indeed somewhat special, it is worth examining why money can arguably be viewed as a bubble. Working through the logic of this argument will allow me to explain some of the implications of this particular definition for bubbles and to set the stage for my discussion of why bursting a bubble may not always make society better off.

To focus the discussion, I use a particular model of money, namely, the one proposed by Townsend (1980) and described in box 2. Essentially, Townsend describes money as a worthless asset that serves no purpose in either production or consumption; it is a piece of paper that yields no direct benefits from holding it. As such, it is akin to an asset that pays no dividends, which is also nothing but a piece of paper that entitles its owner to nothing and serves no other purpose. As such, the fundamental value of money

BOX 1

Formal definition of a bubble

In defining the fundamental value of an asset, it will be easiest to begin with the case of an asset that yields a known and fixed stream of dividends. Let d_t denote the dividend income paid out by the asset at date t, where t runs from 0 to infinity. In order to value a dollar payment received in the future, let us further assume there are markets in which any trader could buy and sell bonds that pay off one dollar at any specified date. Let q_t denote the current price of a bond that pays one dollar at date t. As long as all those who trade in the asset could also access the market for these bonds, each would equally value a dollar at date t and q_t dollars today. Hence, the value any trader attaches to the dividend stream from this asset is given by

$$\text{B1)} \quad F = \sum_{t=0}^{\infty} q_t d_t,$$

where F denotes the fundamental value of the asset. An asset bubble is an asset whose price P is not equal to its fundamental value (equation B1), that is, $P \neq F$. Under fairly weak conditions, one can rule out the case where an asset trades for less than its fundamental value, which is why in the main text I define an asset as a bubble if its price exceeds its fundamental value, that is, $P > F$.

Next, consider the case where dividends are uncertain. That is, suppose that at date t, the state of the world can be one of any states in the set Ω_t, which defines the set of all possible outcomes at that date, and let $\omega_t \in \Omega_t$ refer to a particular state of the world at date t, which all traders believe will occur with probability Prob (ω_t). The state of the world determines the value of the dividend at date t, that is, $d_t = d(\omega_t)$. Let $q_t = q(\omega_t)$ denote the value individuals assign today to a dollar they receive at date t in the particular state ω_t. To relate this to the price of a bond as in the case where dividends are known with certainty, note that if there were a market for state-contingent bonds that paid a dollar at date t in a particular state, and this market were available to all traders, then $q(\omega_t)$ would be the price of this bond. The fundamental value traders assign to the asset in this case is the expectation

$$\text{B2)} \quad F = E\left[\sum_{t=0}^{\infty} q(\omega_t) d(\omega_t)\right]$$

$$= \sum_{t=0}^{\infty} \sum_{\omega_t \in \Omega_t} \text{Prob}(\omega_t)\, q(\omega_t)\, d(\omega_t),$$

where the expectation is taken with respect to the distribution over all states of the world in Ω_t. An asset would be considered a bubble if its price $P \neq F$ as defined in equation B2.

Defining the fundamental value of an asset becomes more complicated if we allow households to have different expectations about dividends or different access to financial markets. For example, suppose households hold different views about the likelihood of the different states in Ω_t. In this case, the expectation in equation B2 will differ across traders, and there will not be a unique fundamental value that we can attach to the asset. However, as Harrison and Kreps (1978) and Allen, Morris, and Postlewaite (1993) point out, it is still possible to talk about an asset as unambiguously overvalued if it exceeds the fundamental value *any* trader in the market would assign to it—that is, if there is no trader who would be willing to buy the asset at the going price if he had to hold the asset forever and never sell it. Thus, we would define an asset as a bubble if its price P exceeds the value of F for each trader. In particular, in that case no trader would purchase the stock with the intent of holding it forever; rather, traders would purchase the stock only because they expected to sell it at a later date with some probability.

Similarly, it might be the case that traders assign different values to dollar amounts at the same state of the world. For example, if there were no markets for contingent bonds, or if only some traders had access to them, traders would typically discount the future at different rates. In this case, the expression in equation B2 would vary across traders because each assigns his own discount rates q_t. Again, there may not be a single fundamental value across all traders, but we can still talk about an asset that is overvalued in the sense that it exceeds the fundamental value *any* trader would assign to it.

Finally, although I have treated q_t and d_t as given, these should in fact be viewed as endogenous sequences that depend on policy. When a policymaker uses policy to change asset prices, she might also change q_t and d_t. This point is discussed, for example, in Weil (1990). Once we acknowledge that the fundamental value of the asset itself can be affected by policy, one has to be careful about what it means to burst a bubble. For example, it might not be possible to burst a bubble without also changing its fundamental value. I ignore this issue in my discussion, although it is certainly a relevant consideration.

BOX 2

The Townsend model of money

The Townsend (1980) model can be described as follows. There are two types of households, labeled *A* and *B*, each of which lives for infinitely many periods. Households of type *i* are endowed with y_t^i units of consumption goods at date *t*. This endowment process is as follows:

$$y_t^A = 0,\ y_t^B = 1 \text{ for } t \text{ even},$$
$$y_t^A = 1,\ y_t^B = 0 \text{ for } t \text{ odd}.$$

That is, the total endowment is always 1, but the two types alternate as to which receives it. Both types of households enjoy consumption according to the utility function

$$\sum_{t=0}^{\infty} \beta^t U(c_t),$$

where $U(\cdot)$ is a concave, increasing, and differentiable function. Assume further that $U'(0) = \infty$, which implies households would want to consume at least some goods in each period.

In each period, both types of household must choose consumption c_t and money holdings for the next period M_{t+1} to maximize their utility. Money holdings are restricted to be positive; that is, money cannot be sold short. Formally, households of type *i* solve

$$\max_{c_t^i, M_{t+1}^i} \sum_{t=0}^{\infty} \beta^t U(c_t^i)$$

subject to the budget constraint

$$\text{B3)} \quad M_{t+1}^i + p_t c_t^i \le M_t^i + p_t y_t^i,$$

where p_t is the price of consumption goods at date *t*, and the constraint that money holdings are nonnegative,

$$\text{B4)} \quad M_{t+1}^i \ge 0.$$

Townsend analyzes various aspects of this model. One of these is the existence of an equilibrium when the price of goods p_t is fixed over time. He shows that there is indeed such an equilibrium. In this equilibrium, a household endowed with one unit of consumption goods will trade a fraction λ of it in exchange for money, and in the period where it is not endowed with consumption goods, it will trade all of the money it holds for goods, where given constant prices will be equal to λ. The value of λ solves the equation

$$U'(1 - \lambda) = \beta U'(\lambda).$$

Since $\beta < 1$ and $U(\cdot)$ is concave, one can show that $\lambda < 1/2$; that is, households consume more when their endowment is high. This equation explains why households expect to earn a capital gain from buying money: The marginal utility of the consumption goods they give up when they buy money is lower than the marginal utility of the consumption goods they purchase when they sell their money holdings, and so they realize a capital gain in utility from this transaction.

To see how money in the model can be viewed as the present discounted sum of quasi dividends, let λ_t^i denote the Lagrange multiplier on the budget constraint (equation B3) and θ_t^i the multiplier on the nonnegativity constraint (equation B4). The first order conditions for households of type *i* imply that

$$\text{B5)} \quad \lambda_t^i = \beta\theta_{t+1}^i + \beta\lambda_{t+1}^i.$$

Repeatedly substituting forward for λ_{t+1} from equation B5 yields

$$\text{B6)} \quad \lambda_t^i = \sum_{j=t+1}^{\infty} \beta^j \theta_j^i.$$

Since λ_t^i is the value the household *i* attaches to a marginal unit of money, the similarity between equation B6 and equation B1 (in box 1) naturally suggests an interpretation in which the value of money to a household is equal to the present discounted value of a stream of "dividend" payments θ_t, where θ_t represents the transaction services the household obtains from money. While this provides an intuitive explanation for the value of money, it is also somewhat misleading. A positive "dividend" θ_t at date *t* reflects the fact that an optimizing household would trade away money at that date if it received it, and thus achieve a utility gain (since it wants to have a negative value of M_t). It does not represent a dividend payment that a household would receive simply from holding on to the asset for a period, in line with what a true dividend payment reflects. Hence, although one can interpret the value households attach to money in the model as the discounted value of certain quasi dividends, this does not contradict the fact that money in Townsend's model can be viewed as a bubble.

should be zero, and the fact that buyers and sellers are willing to trade valuable goods for money implies its price exceeds its zero fundamental value.

The notion that money can be viewed as a bubble may strike some readers as odd. The classic example of a bubble is an asset whose price rises rapidly, encouraging investors to buy it, even though it is overvalued, because they can turn around and sell it at a higher price than they bought it at. At first glance, there does not seem to be an obvious analogue to this in the case of money.[4] For example, in a world with a constant price level over time, a household would give up just as many goods to buy a dollar as it can purchase with that dollar. So in what way is money providing a capital gain to a seller who trades his goods and services for money? Here, Townsend's model reveals that money does in fact accrue a capital gain, but one which is measured in utility terms rather than financial terms. In particular, for reasons inherent to the Townsend model and which are explained in box 2, households value goods more when they trade money for goods than when they trade goods for money. Thus, they buy money with "cheap" goods and turn around and sell it for "valuable" goods, just as a speculative trader might buy an overvalued asset at a lower price than he plans to sell it at. More generally, monetary models tell us that when households hold money, they are trading goods they own but do not like very much for money, and later they trade this money for different goods they like more than the goods they sold. These gains would be reflected in the way a household personally values goods, not in the market price of goods. But these gains are no different in principle than the capital gains one could earn on a stock: Each household buys money with something that gives it low utility, then turns around and sells money for something that gives it high utility.

Another feature of money that might make it seem odd to refer to it as a bubble is that money does play a useful role, specifically by facilitating transactions. This value should presumably be counted as part of the inherent value of money. Indeed, as discussed in box 2, the value a household in Townsend's economy attaches to money is precisely equal to the present discounted value of the transaction services money provides to that household. So why does it make sense to think of money as overvalued? The issue rests largely on the definition of dividends. A dividend for an asset is typically defined to be the benefits an owner of an asset accrues from retaining ownership of the asset, net of any changes in the price of the asset. In Townsend's model, these benefits are exactly zero.[5] It is only when the household sells money for goods that the household becomes better off, which ought to be counted as a capital gain. Of course, if we adopt a more expansive view of dividends according to which money's transaction services would be counted as dividends, we would not choose to define money as a bubble. But the fact that it is possible to view money as a bubble is instructive. In particular, as I discuss later in this article, the reason money plays a useful role in Townsend's model is that it serves to mitigate a distortion that is inherent to Townsend's economy, and it is this very distortion that allows money to exist in the first place. If bubbles occur in environments in which there are already distortions, eliminating the bubble may not always be socially desirable.

For the remainder of this article, I use the term bubble to refer to an asset whose price cannot be justified by the value of the dividends that society expects to earn from this asset class collectively. This definition says nothing about the rises and falls in asset prices that characterized the most prominent historical episodes often identified as bubbles, but it does capture the oft-cited notion that during these episodes asset prices seemed to have become temporarily unhinged from their true worth.

When can asset bubbles arise?

The idea that asset prices can assume arbitrary values rather than reflect their underlying worth is rather disturbing. Not surprisingly, economists have sought to determine whether this possibility can be ruled out. There are in fact various conditions that can be used to rule out the existence of bubbles on theoretical grounds. This section reviews the results for when bubbles can and cannot occur, with the purpose of eventually exploring what the conditions necessary for bubbles to emerge imply for the desirability of setting out to keep the prices of bubbles in check.

A natural starting point for a discussion of when bubbles can emerge is a paper by Milgrom and Stokey (1982). They consider the distinct but related question of whether it is possible for agents to engage in "speculative trading" in one-shot markets. To understand their result, consider an example of a stock that can be traded today and then never again. Two investors, Alice and Bob, participate in the market, and initially Bob owns all of the stock. Suppose that after doing research on the company, Alice decides that the company is likely to be profitable and pay out high dividends. At the same time, Bob decides after doing his own independent research that the company is likely to be unprofitable and pay out low dividends. Milgrom and Stokey asked whether Bob and Alice would agree to trade the asset, with each speculating

that they will be benefiting at the other's expense. This question does not bear directly on the existence of asset price bubbles, which crucially depends on the possibility of resale. Nevertheless, this question is related to the possibility of bubbles. In particular, investors who purchase an overvalued asset are also speculating that they will be able to sell the asset to another trader in the future. As such, they must believe they will be able to sell the asset for more than the dividends it will pay future owners of the asset. This is similar to the way Bob expects to sell Alice an asset for more than he believes it is worth. Not surprisingly, the conditions that rule out speculative trade between Alice and Bob in this example can also rule out the existence of asset bubbles.

Milgrom and Stokey argue that in their example, speculative trade will not be possible under certain circumstances. This is because the fact that the other party is willing to engage in trade will convince both Alice and Bob that they may have failed to take into account some information about the asset that the other side has access to. This would dissuade both sides from completing the trade, even if they never actually see the information of the other party. The formal argument that rules out speculative trade requires several assumptions. First, all traders must share the same initial beliefs about the company before they engage in research.[6] Second, before trade occurs, the stock is assumed to have been allocated efficiently. That is, there would have been no reason for Alice to buy the stock from Bob back when both of them shared the same initial beliefs about dividends. This is important, since without this assumption Bob would not be able to tell whether Alice was willing to buy the stock because she has favorable information about dividends or because she simply has a greater preference for owning the asset. Finally, all traders are assumed to be rational and profit maximizing, and this is commonly known by all traders. That is, not only are Alice and Bob rational, but they know each other to be rational, they know that the other trader considers the counterpart to be rational, and so on. If all of these conditions are satisfied, Alice and Bob will not trade.

Around the time Milgrom and Stokey published their result on speculative trade, Tirole (1982) separately showed that a similar set of assumptions can rule out the possibility of asset bubbles. As in Milgrom and Stokey, Tirole (1982) assumed that traders start out with a common set of beliefs, that it is common knowledge that all traders are rational, and that resources are allocated efficiently prior to trading. In addition, Tirole (1982) assumed there are only finitely many traders who can trade in the asset, although each can trade it infinitely many times. Under these assumptions, Tirole showed that asset bubbles could not occur. The intuition is as follows. The assumption that the initial allocation of resources is efficient implies there is no mutually advantageous trade between the agents. As a result, any investor who agrees to buy an overvalued asset must believe that doing so benefits him at the expense of remaining traders taken as a whole. In particular, he expects to gain at their expense by selling them a stream of dividends for more than it is worth. When there are only finitely many traders, these beliefs cannot all simultaneously be true: If everybody who buys the overvalued asset expects to benefit at the expense of all other traders, there won't be any traders left from whom to actually benefit. The argument is analogous to why the children in Garrison Keillor's fictional town of Lake Wobegon cannot all be above average as its residents claim, since deviations from the average have to sum up to zero. Likewise, not all traders can expect to gain at the expense of all remaining traders taken together, since the expected gains of all traders together have to sum up to zero.

For an asset bubble to occur, then, at least one of Tirole's (1982) assumptions must be contradicted: The potential number of traders who trade in the asset is infinite, traders start out with different prior beliefs or they believe other traders are irrational, or there must be some inefficiency in the economy prior to the initiation of trade.[7] I now consider each of these scenarios in turn.

Infinitely many traders

If there are infinitely many potential traders, then the argument for why bubbles could not exist sketched earlier fails. Suppose each investor buys the asset, sells it to someone else, and then leaves the market. The first person believes he benefits at the expense of all the traders that come after him, since they pay him more than he gives them in dividends. The second trader similarly believes that he will benefit at the expense of all traders who come after him, and so on. Since we never reach a last trader, we never reach a contradiction. To borrow an analogy from Tirole (1982), investing in a bubble amounts to a game of "hot potato" in which each person tries to pass on an overvalued asset to somebody else so as not to get stuck with a stream of dividends worth less than what he paid for it. With only finitely many people, the last person in line will refuse to buy the asset. But then the person before him will refuse to buy the asset, knowing he will be stuck with it. By the same logic, no investor would ever agree to invest in an overvalued asset. But with infinitely many traders, we never run out of

hands to which to pass the hot potato. Various examples of such bubbles are formally demonstrated in Tirole (1985). However, the fact that a bubble can arise when there are infinitely many people who can trade the asset was demonstrated much earlier, most notably in early work on monetary economics such as Samuelson (1958).

Requiring that there be literally infinitely many traders seems like an odd condition to rely on for the existence of bubbles. Indeed, Santos and Woodford (1997) have argued that bubbles that hinge on infinitely many traders are both exceptional and fragile, in the sense that they typically will not survive slight modifications in the underlying environment. Another potential problem with relying on infinitely many traders to sustain a bubble is that since an asset must potentially pass through infinitely many hands, the economy must grow at least as fast as the bubble to ensure that there are enough resources to keep trading the asset. Without this assumption, only finitely many traders will be able to buy the asset, and then Tirole's (1982) original argument for the non-existence of bubbles would apply. Abel et al. (1989) propose a way to test whether the growth rate of the economy is large enough to accommodate an asset bubble, and find that this condition has never been satisfied in the U.S., including the period of the rise and crash of stock prices throughout the 1920s depicted in figure 2.

Differences in initial beliefs and the possibility of irrational traders

If all traders begin with the same initial beliefs about the fundamental value of the asset but then conduct their own independent research to learn more about the true fundamental value, the logic of Milgrom and Stokey (1982) rules out the possibility that two traders starting with the same initial beliefs would come to believe that they can each benefit at the expense of the other. But if traders have different initial beliefs or else believe that other traders are irrational, a trader might think that as long as there are other traders with incorrect beliefs, he can profit at their expense because they mistakenly fail to realize they are being taken in when they buy the asset.

Various papers have formally demonstrated how these conditions can lead to the existence of bubbles. One such paper is Harrison and Kreps (1978). They assume that traders begin with different initial beliefs and then stick to their beliefs regardless of what others believe and regardless of what they observe about dividends. One of their results shows that an asset can trade for more than what any one trader in the economy believes the fundamental value should be, even the trader with the most optimistic view of dividends. Traders are willing to engage in trade in an asset they believe is overvalued because they all believe (at least some of them mistakenly) that by trading the asset they can profit at the expense of others. Another paper that follows this approach is De Long et al. (1990). They assume that both rational and irrational traders participate in asset markets. Unlike in Harrison and Kreps (1978), there is no restriction that irrational traders must believe their strategies to be profitable. Rational traders would have an incentive to take advantage of arbitrage opportunities when assets are mispriced, but because they are assumed to be risk averse and to have short time horizons, they will not drive the price of an overvalued asset down to its fundamental value, and bubbles can exist.

There are several empirical findings that support the notion that bubbles are associated with the possibility of at least some traders committing systematic mistakes in their trades. On the one hand, many of the historical episodes of bubbles cited previously seem to coincide with the entry of unsophisticated traders into their respective markets. This view is summarized in White (1990), who writes:

> During the tulipmania, Garber notes that the middle classes and even monied workers began to speculate in the market for tulips, which previously had been the province of specialists. In their enthusiasm to participate, new investors seem incredibly incautious. ... In 1873, new German investors played a prominent role. They busily acquired new, untested foreign securities for their portfolios. In the 1920s the shift in business financing from short-term commercial bank loans to bonds and stocks meant that instead of commercial banks who had considerable experience in evaluating firms, the general investing public became the chief creditors of corporations. Americans who had never owned stocks before were now buying. Given the increased difficulty of evaluating fundamentals and the general optimism from the decade of prosperity, it is not surprising that prices were pushed above fundamentals. Similarly in the 1980s, new financial instruments drew in new investors from both at home and abroad. (p. 238)

Further support for this view of why bubbles can occur comes from controlled experiments in which test subjects (typically undergraduate students) were allowed to trade assets in simulated markets. The best known example is Smith, Suchanek, and Williams (1988). They set up an experiment in which subjects were allowed to trade assets whose dividends were

explicitly described to all participants prior to the experiment. Subjects were paid according to the wealth their trading strategies netted in each experiment, providing them with an incentive to trade optimally. Since all traders were given the same information by design, there should not have been a bubble. Nevertheless, assets in their experiment traded above their fundamental value, first rising above the fundamental value and then crashing back down to it as the asset reached the end of its life. This suggests either that the test subjects failed to correctly compute the true value of the asset, or else they believed other subjects would compute the value incorrectly and so one could trade profitably at the expense of such subjects. Subsequent work by Lei, Noussair, and Plott (2001) finds some support for the former explanation, since bubbles occurred even when participants were told they could not resell an asset after they bought it. Since precluding resale precludes the possibility of exploiting other traders, a bubble could occur in such a market if traders fail to calculate the fundamental value of the asset. That said, Smith, Suchanek, and Williams (1988) find that once participants have some experience trading in these experimental markets, bubbles no longer emerge. Following up on this finding, Dufwenberg, Lindqvist, and Moore (2005) argue that bubbles fail to emerge in laboratory settings where at least one-third of the traders in the experiment are experienced. This lends support to the view that uninitiated traders do in fact make mistakes. What is less clear from the experiments is whether experienced traders are willing to speculate on buying overvalued assets when they believe there are many uninitiated traders active in the market to exploit.

Thinking about how to deal with bubbles that arise because some traders make systematic mistakes is tricky. First, if traders disagree over the true value of an asset, it is not obvious what policymakers should believe about the asset. For example, in Harrison and Kreps' (1978) framework, all traders are firmly convinced that their beliefs are correct. Clearly, at least some of them must be wrong. So how confident should policymakers be that their own beliefs are correct? One compromise might be to wring out at least some excesses by driving down the price of the asset to the most optimistic assessment any trader has. However, it is not clear that this in itself is desirable. Moreover, driving down the price of the asset may be infeasible or counterproductive. One of the points in Harrison and Kreps (1978) is that, given the conflicting beliefs of the traders in their model, the price of the asset has to exceed the value all traders assign to that asset for the market to clear. Simply driving down the current price of the asset might encourage some investors to buy more of the asset and bid its price back up. In addition, as long as traders maintain different beliefs about the underlying dividends, they will always have an incentive to engage in speculative trade. Even if policymakers were committed to keeping the price of a particular asset contained indefinitely, traders with different beliefs would still want to engage in speculative trades, possibly in other markets.

To the extent that bubbles are caused by naive traders' failure to assess the fundamental value of an asset correctly, as some of the historical evidence might suggest, policymakers might be able to rein in bubbles by educating these traders about the true value of the asset. There may be some wisdom to doing this. However, in a certain sense this intervention amounts to redistributing wealth from one party to another, namely, from traders who incurred the cost of acquiring information about the asset to uninformed traders who enter the market. On its own, this would be a poor justification for public policy, since policy should aim to benefit society at large and not some group at the expense of others. Indeed, even where government policies are aimed at curtailing abuse against certain parties, they are usually motivated by a desire to benefit all of society. For example, regulation to limit collusion by producers is viewed as serving the public interest because collusion involves inefficiently low levels of production, not because it favors consumers over producers. There may be similar reasons why it is inefficient to let informed traders exploit uninformed traders. But a compelling justification for bursting bubbles would require an explicit argument for why this should make society as a whole better off.

Inefficiency

The final scenario in which bubbles can occur is if resources are allocated inefficiently in the absence of a bubble. Consider once again the Townsend (1980) model described in box 2. The economy consists of two types of households, each of which lives for infinitely many periods. In each period, households receive some endowment of consumption goods. Different types of households will receive different endowments. In even periods, households of type 1 receive all of the consumption goods available in the economy and households of type 2 receive none. In odd periods, households of type 2 receive all of the goods available and households of type 1 receive none. Thus, the fortunes of any given household fluctuate over time. Both types of households are assumed to prefer consuming the same amount on average every period to consuming more in some periods and less in others. But by

assumption, goods cannot be stored, there are no other assets in the economy, and households are unable to commit to long-term contracts with one another. This model satisfies all of the assumptions in Tirole's (1982) model except for the requirement that the initial allocation is efficient. In particular, all households would prefer a situation in which they could trade claims on each other's endowments so as to achieve a flat consumption profile to being forced to consume whatever their endowment is in each period.

Townsend (1980) used this model to provide a theory for the existence of money. But as I alluded to earlier, we can reinterpret his results as an example of an asset price bubble. In particular, suppose we introduce an asset that pays no dividends into this economy. Although the fundamental value of this asset is zero, Townsend's results imply that it will be possible for this asset to trade at a positive price. In particular, households that are rich (that is, those that receive an endowment today) will trade some of their endowment in exchange for the asset, while households that are poor (that is, those that do not receive an endowment today) will trade any of the asset they own for goods. In this case, the asset would trade back and forth between the two types of households at a positive value. The asset plays the role of money in that it allows households to sell excess goods today in the hope of trading the asset for goods tomorrow. To understand why Tirole's (1982) non-existence result breaks down, note that since the initial allocation is inefficient, an investor who considers buying an overvalued asset no longer has to believe that to gain from trade he must benefit at the expense of another trader. When a household that is rich buys the asset and then sells it again when it is poor, it is not exploiting the rich household it sells the asset to; rather, the rich household wants to buy the asset because it wishes to hold it now and trade it in the future. By contrast, Tirole's (1982) argument ruled out asset bubbles because all traders had to simultaneously believe they were taking advantage of other traders.

One of the features of Townsend's model is that the asset bubble can be traded infinitely many times. This feature is essential in his model, since otherwise, households would refuse to purchase an overvalued asset in the last period, knowing that there would be no opportunity to sell it again. But then in the previous round of trading no investor would agree to purchase the asset if it were overvalued, knowing he would be unable to sell it at a profit, and so on. The infinite nature of the model might seem essential for a bubble to exist. However, this would be an incorrect conclusion, as demonstrated by Allen, Morris, and Postlewaite (1993). They provide an example of a bubble in an environment where the initial allocation of resources is inefficient, but there are only finitely many opportunities to trade as well as finitely many traders. Their example hinges on the assumption that information about the asset is not common knowledge. In particular, all traders in their example know the asset is worth less than its price, but traders are not sure whether other traders know this to be true. Thus, an investor believes that there might be another trader who would be willing to buy the asset from him, even if that trader expects to hold the asset to maturity. The inefficient allocation of resources is crucial to their example as well, for precisely the same reason: It allows traders to believe that the asset might ultimately be sold to someone who will not be made worse off from buying it.

The possibility that bubbles occur because of some underlying economic inefficiency is especially important for evaluating the desirability of asset price bubbles. If a bubble exists because of some underlying inefficiency in the economy, we would need to account for how the bubble interacts with this underlying inefficiency to assess whether policy intervention is desirable. In particular, we would need to determine whether eliminating the bubble ameliorates or exacerbates this same inefficiency.

Examples where bursting a bubble is distinctly undesirable

One point worth noting is that if bubbles arise because of an underlying inefficiency elsewhere in the economy, it might be possible that bubbles serve to mitigate the inefficiency that allows them to exist in the first place. In that case, bursting the bubble may very well exacerbate the inefficiency and make society worse off. We have already seen one example of this in Townsend's model. The key feature of that model, as well as of most other monetary models, is that there are mutually advantageous trades that for various reasons people fail to enter.[8] In Townsend's model, the problem is that agents cannot commit to trade over time. In particular, rich households would like to trade some of the goods they are endowed with today for goods they can consume when they are poor tomorrow. But this necessitates trading over time: In exchange for goods today, a household has to promise to provide goods tomorrow. But if it is not possible to track down a person in the future or to rely on courts to enforce such contracts, these trades will not be possible. The virtue of money, or any repeatedly traded asset for that matter, is that it can overcome such difficulties. As long as people speculate that there will

be a future market in which they can trade money (or an asset) for goods, there is no need to track down the original person with whom they traded their goods. As far as each individual is concerned, he supplies goods in exchange for an asset that he expects to sell later. There is no need to enforce contracts or engage in bilateral agreements with people who may be hard to track down later. The bubble works to alleviate the underlying inefficiency by encouraging the very trades that would not take place without it. Bursting the bubble would then limit trade between agents, and households would be made worse off because they would have to face more volatile consumption than if the bubble was allowed to persist.

One of the features of the bubble in Townsend's economy is that it never bursts. But in a world where a bubble is likely to burst eventually, the relevant choice for a policymaker would not be whether to let a bubble persist, but rather whether to hasten its collapse. It turns out that allowing for the possibility that a bubble collapses on its own would not change any of the implications in a model like Townsend's, a point emphasized in Weil (1987). Weil explored a monetary model in which there is a constant probability each period that money will lose all of its value. In that model, an asset bubble will play a useful role in fostering trades that would not occur otherwise for as long as it survives, and hastening its collapse would still be detrimental to the households in that economy.

Although money is an unconventional example of a bubble, the lesson inherent from monetary models of the type described previously is that bubbles may encourage mutually beneficial trades that would not have taken place otherwise. To see why these issues may have relevance beyond just monetary models, consider the following example. One of the distinguishing features of most housing markets is that they tend to be distorted in various ways by government regulations. One common distortion is due to the way governments assess home values to determine the amount of property taxes homeowners must pay. In an economy where housing prices tend to appreciate, say, for demographic reasons as demand for housing increases, the assessed values of homes that have not sold recently often lag behind. This creates an incentive for homeowners to stay put in order to keep their tax payments low; if they purchase another home, the value will be assessed more accurately and their tax bill will rise. But there may be reasons to encourage turnover in homeownership. For example, suppose that certain neighborhoods in the city are inherently suited to families with young children. If these areas do not have well-developed rental markets, the dampening effect on turnover would mean that too few of the houses in these areas would be purchased by families with young children for whom these areas are ideal. Anything that serves to encourage greater turnover in homeownership despite these tax distortions thus has the potential to improve the allocation of resources within the city.

A bubble in the housing market is one example of something that can serve to increase turnover in homeownership. In particular, a bubble implies that the price of a house exceeds its fundamental value. For owners to be willing to purchase such a house, they must anticipate selling the house in the future. Thus, a bubble in the housing market may end up generating turnover that would not have occurred otherwise. Of course, a bubble in the housing market could be distorting in and of itself, for example, by encouraging excessive building in particular neighborhoods. What the example reveals is that if there is some inherent distortion that precludes agents from trading resources efficiently, there may be certain negative consequences to bursting a bubble.

Bubbles and resource diversion

Even if bubbles serve to facilitate trades that would not occur otherwise, one of the commonly cited reasons for why it would be desirable to burst them is that they divert resources from other productive uses. The idea is that since agents respond to prices when they undertake economic decisions, a price that deviates from the true value of an asset can distort the incentives of agents and lead to suboptimal allocation of resources. I now consider this possibility. Once again, the main lesson from the economic literature on bubbles is that it is important to take into account the conditions that allow a bubble to emerge in the first place when determining the likely benefits from bursting the bubble.

The earliest example of a model in which bubbles can affect the allocation of resources is Diamond (1965). His model builds on the earlier work in monetary theory by Samuelson (1958). Instead of an endowment economy, Diamond considers a production economy in which it is possible to think about how the presence of an overvalued asset affects the allocation of resources for production. Diamond assumes households have two options when saving: They can hold capital, which can then be used in production, or they can hold an intrinsically worthless asset that pays no dividends. Since this asset trades at a positive price in the model, it can be viewed as a bubble. A bubble can emerge in this framework because Diamond assumes there are infinitely many households that can trade in the asset,

which is one of the conditions Tirole (1982) ruled out when deriving his non-existence result. However, even with infinitely many traders, it turns out that a bubble can only emerge if the economy grows at least as fast as the bubble, since households need to be able to afford to purchase the asset as it increases in value. In the model, the bubble grows at the rate of interest. Hence, a necessary condition for a bubble to exist is that the economy grows faster than the rate of interest. But it is well known that this condition indicates that resources are allocated inefficiently, since such a low interest rate implies there must be too much capital accumulation. Thus, once again, a bubble exists when the underlying economy exhibits some inefficiency, although the inefficiency is not the essential element that gives rise to the bubble.

Since households can hold their wealth either in the bubble asset or capital, the emergence of a bubble will necessarily crowd out capital accumulation. Thus, the emergence of the bubble will divert resources from other activities. But in Diamond's economy, a bubble can only exist when capital accumulation is already excessive. That is, bubbles will only divert resources away from capital in his economy if there is too much capital accumulation to begin with. Hence, the conditions that are necessary for a bubble to exist imply that resources would have been allocated inefficiently without the bubble, and so reallocation is actually socially beneficial. Tirole (1987) strengthens this intuition by showing that in Diamond's environment, the presence of an "asymptotic" bubble—that is, a bubble whose value grows at the same rate as the rest of the economy—necessarily implies the economy is consumption-efficient. Thus, the presence of a bubble in this model should be seen as evidence of an efficiently operating economy, not as something undesirable. Work by Abel et al. (1989) mentioned earlier raises doubts that the U.S. economy exhibits excessive capital accumulation. The scenario envisioned in Diamond's model may therefore be of little relevance in practice. However, it is a useful illustration of the more general point that for a bubble to emerge in the first place may require the presence of distortions elsewhere in the economy, and so we need to think carefully about how any diversion of resources due to bubbles interacts with these distortions.

Although in Diamond's model a bubble can only exist when there is excessive capital accumulation, the more important feature of this model that allows a bubble to emerge is the assumption of infinitely many traders. Indeed, subsequent work by Saint-Paul (1992), Grossman and Yanagawa (1993), and King and Ferguson (1993) shows that one can modify the Diamond model so that a bubble exists when capital accumulation is not excessive. All three do so by introducing an externality in capital accumulation. That is, they assume that the productivity of a firm depends on the economy-wide average level of capital. This assumption relies on a broad interpretation of capital that includes not only physical capital but also organization and knowledge capital. The more firms in the economy accumulate useful knowledge, the more likely remaining firms are to use this accumulated knowledge base to improve their own productivity. Because of this externality, the level of capital in the economy is inefficiently low rather than excessive. In this case, the emergence of a bubble on an intrinsically useless asset would divert resources from already scarce capital. This suggests an environment in which it might be desirable to burst bubbles after all. But as pointed out by Grossman and Yanagawa (1993), once a bubble emerges in their model, even though bursting the bubble will allow for more capital accumulation, it will necessarily make some households worse off. The reasoning is as follows. In their model, households buy assets when they are young to save for retirement, which they then sell when they are old to finance consumption. Bursting a bubble will therefore erode the retirement savings of those older households that happen to hold the asset at the time the bubble is burst. In order to make sure no household is left worse off after bursting the bubble, we would need to compensate these older households for their loss out of current resources. But the resources we would need to compensate older households are the same resources that would have been released for capital accumulation. Essentially, society can benefit from bursting the bubble only if the resources that younger households would have transferred to older households by buying the older households' overvalued assets are reallocated to capital accumulation.

Interestingly, subsequent work by Oliver (2000) argues that even when we allow for externalities that imply there is too little capital accumulation, there may be bubbles that serve to mitigate the underaccumulation of capital rather than exacerbate it. His key insight is that the bubble has to be tied to the right asset. While it is true that a bubble on an asset that substitutes away from capital would exacerbate the underprovision of capital, a bubble associated with capital would encourage greater capital accumulation by increasing the return to holding capital. Proceeding to burst such a bubble would make society as a whole worse off, not better off.

When would bursting a bubble be theoretically desirable?

So far, I have presented several theoretical models of bubbles in which bursting a bubble is distinctly undesirable, since doing so either exacerbates some underlying inefficiency or makes only some households better off while making some others worse off. Are there situations in which bursting an asset bubble would make all households better off, assuming the bubble could be accurately identified and deflated without affecting other variables? In Grossman and Yanagawa's (1993) model, the answer is no: The only way to make households better off by eliminating the bubble is to take resources from some traders and reallocate them in a more efficient way. But a key feature of this model is that by the time the gains from a more efficient use of resources are realized, those whose resources were seized are already gone and cannot be compensated. One can potentially make a case that bursting bubbles is desirable, assuming one can deal with the practical issues of implementing such a policy, if it were somehow possible to compensate those who are hurt from the fall in asset prices.

One possibility is to compensate the original owners of the asset with some delay after all. In Grossman and Yanagawa's model, the problem is that by the time the benefits of bursting the bubble are realized, the original owners—older households in the model that are already retired—are dead. Although this stark result is due to the stylized nature of the overlapping generations model Grossman and Yanagawa use, it is true empirically that older households tend to own a large fraction of total wealth; therefore, there is genuine concern that bursting a bubble may not benefit many of the households that suffer a capital loss when the bubble is burst. One way to provide redress to these households is to compensate their heirs. This idea is not so far-fetched. Households save into retirement in part because they intend to leave some of their wealth to their children. Thus, it could very well be that bursting the bubble and wiping out the savings of older households affect them most by denying them the opportunity to leave a large inheritance for their children. But if younger households are better off once resources are freed up for capital accumulation, they might be better off on the whole despite having a smaller inheritance.

Although appealing to intergenerational linkages this way seems reasonable, one has to be careful in pursuing this argument. The reason is that while taking into account the possibility that households care about their future generations allows us to compensate those who are hurt when asset prices fall, it may also make bubbles less likely to occur in the first place. A key feature of the Grossman and Yanagawa model is that each household has a finite time horizon, which does play a role in allowing the bubble to emerge. Once we allow traders to care about future generations, their time horizon would effectively become infinite. If each generation cares about the well-being of its own children, households would effectively behave in the best interests of their family dynasties. But as Santos and Woodford (1997) point out, as long as even some households with a nonnegligible fraction of the economy's total wealth have an infinite horizon, bubbles might no longer be possible. Problems arise not because individuals have an infinite horizon per se, but because they have an infinite horizon and a positive share of total wealth. As Weil (1989) demonstrated, bubbles may emerge when agents are infinitely lived if there is enough entry of new dynasties that are not linked to pre-existing families, so that in the limit the fraction of aggregate wealth of each dynasty becomes negligible.

An alternative argument for why bursting a bubble might make society as a whole better off is if bursting the bubble yields immediate benefits as opposed to only in the distant future, and these benefits can in principle be used to compensate those who take a capital loss when the bubble bursts. The key feature of the Grossman and Yanagawa model is that bursting the bubble releases resources for capital accumulation, but capital is only productive with some delay. If bursting the bubble freed up resources that could be more immediately put to use, it might be possible to compensate the owners of the bubble asset without much delay. Suppose we modify the Grossman and Yanagawa model so that maintaining the bubble uses up current resources—that is, buying a bubble requires using up some labor and capital resources. For example, the excess trading that occurs when an asset is overvalued might require labor resources (real estate agents or stock brokers) that could alternatively be used for production. If bursting the bubble releases these resources for production fairly quickly, there may be enough new resources to compensate the original owners of the asset who lose out when the value of their assets collapses.

As a final note, much of the preceding discussion is based on the notion that bursting bubbles could be theoretically desirable because a bubble diverts resources from more productive uses. Another potential argument for bursting bubbles is that it might be better to contain asset prices while a bubble is still small than to wait until it becomes large. Intuitively, the

larger the bubble, the larger the price collapse and spillover effects on output and consumption would be when the bubble bursts. In other words, bubbles might contribute to making the economy more volatile than necessary. This is an intriguing idea, although there is little formal work that can help to evaluate this particular argument.

Conclusion

In this article, I have argued that in addition to the practical difficulties associated with identifying asset bubbles and using monetary policy to affect asset prices, there are theoretical reasons why policymakers should proceed with caution if they intend to combat asset bubbles.

In many ways, economic theory views bubbles as oddities, since under ideal conditions they would never occur. More precisely, if traders are rational and markets operate efficiently, as is typically assumed in various economic models, the existence of bubbles can be ruled out explicitly. Some economists, most notably Garber (1990), have argued on this basis that most of the commonly cited historical examples of asset bubbles are not really bubbles as economists define them. Of course, if we can be assured that bubbles are unlikely to ever occur, there is little reason for policymakers to worry about the prospect of reining them in if they were to occur. Yet, as the work reviewed in this article suggests, bubbles can occur once we depart from ideal conditions, and so there is some merit to contemplating whether and how policymakers should respond to a bubble. At the same time, the conditions that allow bubbles to exist may have important implications for the desirability of such policies. Equating asset prices with their fundamental values—a feature that is desirable in classical models—may no longer be possible or desirable once we depart from such ideal conditions. Thus, one has to be careful not to rely too much on the intuition that it is always desirable to have the price of an asset reflect its true worth—an idea that implicitly assumes perfectly functioning markets—to guide policy. Simply put, the emergence of a bubble may signal that the economy already suffers from certain structural problems, and then one has to be careful to distinguish whether bursting the bubble can mitigate or exacerbate these problems.

NOTES

[1]The NASDAQ-100 is an index that corresponds to the value of the National Association of Securities Dealers Automated Quotation system's 100 largest nonfinancial securities, based on market capitalization.

[2]More generally, a bubble is a situation in which the price of the asset deviates from its fundamental value, be it because the price exceeds the fundamental value or falls short of it. The latter case is sometimes referred to as a negative bubble. In this article, I focus on positive bubbles, which more closely correspond to the asset bubbles that figure most prominently in policy discussions.

[3]For example, McGrattan and Prescott (2004) argue that the stock market prior to the 1929 crash was if anything undervalued and that the subsequent crash could have been caused by monetary or regulatory policy. Garber (1989) argues that rare tulip varieties in seventeenth-century Netherlands were highly prized, justifying the high price of bulbs that could produce these exceptional varieties. Garber (1990) goes on to argue that the Mississippi and South Sea episodes might not correspond to bubbles either. For a more comprehensive discussion of the Mississippi Company, see Velde (2003), who, among other things, discusses whether the stock was overvalued and should rightly be viewed as a bubble.

[4]Note that in a period of general price deflation, a dollar costs less in terms of goods than it will purchase in the future. In that case, holding money would in fact earn a capital gain. But the argument that money can be viewed as a bubble is independent of any overall trends in the general price level.

[5]In particular, households in Townsend's (1980) model cannot lend their money holdings to others who need money and charge them some interest rate. In practice, this option may be available to households; for example, most depository institutions offer some interest on checking accounts. Thus, money may earn some dividends rather than strictly zero, as is so starkly assumed in the model.

[6]Formally, this is known as the common priors assumption. For a discussion on the merits and validity of this assumption, see Morris (1995).

[7]In addition to violating at least one of Tirole's (1982) assumptions, a necessary condition for bubbles to exist is that traders are short-sales constrained. Since a bubble is worth more than the present discounted value of the dividends it yields, agents can make arbitrage profits by selling short shares of the bubble and then using the proceeds they earn to pay out dividends, pocketing the leftover as profit. If people were free to sell assets short, the threat of infinite supply would ensure that the price of an asset could never exceed its fundamental value, thus precluding the possibility of a bubble.

[8]An excellent collection of papers on the theory of money is Kareken and Wallace (1980). A follow-up compendium of more recent advances was published in a special May 2005 issue of the *International Economic Review*.

REFERENCES

Abel, Andrew, Greg Mankiw, Lawrence Summers, and Richard Zeckhauser, 1989, "Assessing dynamic efficiency: Theory and evidence," *Review of Economic Studies*, Vol. 56, No. 1, January, pp. 1–19.

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

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.

Camerer, Colin, 1989, "Bubbles and fads in asset prices," *Journal of Economic Surveys*, Vol. 3, No. 1, pp. 3–41.

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.

De Long, J. Bradford, Andrei Shleifer, Lawrence Summers, and Robert Waldmann, 1990, "Noise trader risk in financial markets," *Journal of Political Economy*, Vol. 98, No. 4, August, pp. 703–738.

Diamond, Peter, 1965, "National debt in a neoclassical growth model," *American Economic Review*, Vol. 55, No. 5, Part 1, December, pp. 1126–1150.

Dufwenberg, Martin, Tobias Lindqvist, and Evan Moore, 2005, "Bubbles and experience: An experiment," *American Economic Review*, Vol. 95, No. 5, December, pp. 1731–1737.

Economist Newspaper Limited, 2005, "Monetary conundrums," *Economist*, Vol. 374, No. 8415, February 26–March 4, pp. 12–14.

Garber, Peter, 1990, "Famous first bubbles," *Journal of Economic Perspectives*, Vol. 4, No. 2, Spring, pp. 35–54.

_____, 1989. "Tulipmania," *Journal of Political Economy*, Vol. 97, No. 3, June, pp. 535–560.

Grossman, Gene, and Noriyuki Yanagawa, 1993, "Asset bubbles and endogenous growth," *Journal of Monetary Economics*, Vol. 31, No. 1, February, pp. 3–19.

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

Kareken, John, and Neil Wallace (eds.), 1980, *Models of Monetary Economies*, Federal Reserve Bank of Minneapolis.

Kindleberger, Charles, 1996, *Manias, Panics, and Crashes: A History of Financial Crises*, 3rd ed., New York: John Wiley & Sons.

King, Ian, and Don Ferguson, 1993, "Dynamic inefficiency, endogenous growth, and Ponzi games," *Journal of Monetary Economics*, Vol. 32, No. 1, August, pp. 79–104.

Lei, Vivian, Charles Noussair, and Charles Plott, 2001, "Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality," *Econometrica*, Vol. 69, No. 4, July, pp. 831–859.

Lipsey, R. G., and Kelvin Lancaster, 1956–57, "The general theory of second best," *Review of Economic Studies*, Vol. 24, No. 1, pp. 11–32.

McGrattan, Ellen, and Edward Prescott, 2004. "The 1929 stock market: Irving Fisher was right," *International Economic Review*, Vol. 45, No. 4, pp. 991–1009.

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

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

Oliver, Jacques, 2000, "Growth-enhancing bubbles," *International Economic Review*, Vol. 41, No. 1, February, pp. 133–151.

Saint-Paul, Gilles, 1992, "Fiscal policy in an endogenous growth model," *Quarterly Journal of Economics*, Vol. 107, No. 4, November, pp. 1243–1259.

Samuelson, Paul, 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.

Santos, Manuel, and Michael Woodford, 1997, "Rational asset pricing bubbles," *Econometrica*, Vol. 65, No. 1, January, pp. 19–58.

Smith, Vernon, Gerry Suchanek, and Arlington Williams, 1988, "Bubbles, crashes, and endogenous expectations in experimental spot asset markets," *Econometrica*, Vol. 56, No. 5, September, pp. 1119–1151.

Tirole, Jean, 1987, "Intertemporal efficiency, intergenerational transfers, and asset pricing: An introduction," in *Essays in Honor of Edmond Malivaud, Volume 1: Microeconomics*, Cambridge, MA: MIT Press, pp. 116–142.

_____, 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. 1163–1182.

Townsend, Rob, 1980, "Models of money with spatially separated agents," in *Models of Monetary Economies*, John Kareken and Neil Wallace (eds.), Federal Reserve Bank of Minneapolis.

Velde, François, 2003, "Government equity and money: John Law's System in 1720 France," Federal Reserve Bank of Chicago, working paper, No. WP-2003-31, November.

Weil, Philippe, 1990, "On the possibility of price decreasing bubbles," *Econometrica*, Vol. 58, No. 6, November, pp. 1467–1474.

_____, 1989, "Overlapping families of infinitely lived agents," *Journal of Public Economics*, Vol. 38, No. 2, March, pp. 183–198.

_____, 1987, "Confidence and the real value of money in an overlapping generations economy," *Quarterly Journal of Economics*, Vol. 102, No. 1, February, pp. 1–22.

White, Eugene, 1990, "Are there any lessons from history?," in *Crashes and Panics: The Lessons from History*, Eugene White (ed.), Homewood, IL: Dow Jones-Irwin, pp. 235–240.