
ARTICLES

Can Markets Learn to Avoid Bubbles?

Ross M. Miller

One of the most striking results in experimental economics is the ease with which market bubbles form in a laboratory setting and the difficulty of preventing them. This article re-examines bubble experiments in light of the results of an earlier series of market experiments that show how learning occurs in markets characterized by an asymmetry of information between buyers and sellers, such as found in Akerlof's lemons model and Spence's signaling model.

Markets with asymmetric information are incomplete because they lack markets for specific levels of product quality. Such markets either lump all qualities together (lemons) or using external indications of quality to separate them (signaling). Similarly, the markets used in bubble experiments are incomplete in that they are lacking a complete set of forward or futures markets, depriving traders of the information supplied by the prices in those markets. Preliminary experimental results suggest that the addition of a single forward market can sometimes mitigate bubble formation and this article suggests more extensive research in this direction is warranted. Market bubbles outside of the laboratory usually are found in markets in which forward and futures markets are either legally restricted or otherwise limited.

Experimentation in markets with asymmetric information also indicates that the ability of subjects to learn how to send and receive signals can be enhanced by changing the way that market information is presented to them. We explore how this result might be used to help asset markets learn to avoid bubbles.

The bubble in Internet-related stocks that formed in the late 1990s and then ultimately burst illustrates how markets can exhibit unstable behavior that would appear to impugn both their rationality and efficiency. While some defenders of the efficient-market theory, most notably Peter Garber [1990], claim that the great speculative bubbles of history—Dutch tulip mania, the South Sea bubble, John Law's Mississippi scheme, etc.—were not really bubbles but rather reflect a rational market response to conditions of low supply and high demand, the experimental examination of markets within a controlled laboratory setting beginning with the Vernon Smith, Gerry Suchanek, and Arlington Williams [1988] study has demonstrated the ease with which bubbles can form. While the true supply and demand in a naturally-occurring market can often only be roughly approximated at best, within the confines of a market laboratory, experimenters can control market conditions by the values they induce with payments to

their market subjects. With the parameters of supply and demand under his or her control, the experimenter can determine the degree to which market prices exceed the "rational" level of a competitive equilibrium. One can reasonably conclude that a bubble exists when the observed price greatly exceeds equilibrium and moves away from it rather than converging towards it.

Given the current body of experimental evidence it appears that bubbles arise from a confluence of factors; no single factor appears to be sufficient to generate a bubble in a given market. Of the many ingredients that go into a market bubble the one that has captured the popular imagination is "irrational exuberance," a phrase made famous by Federal Reserve Board Chairman Alan Greenspan and further popularized by Yale economist Robert Schiller [2000]. Irrational exuberance can be viewed as a mass delusion in which an item's "true value" becomes irrelevant as market participants increasingly believe that prices will continue to rise forever, or at least until they can sell their holdings to someone else. While such beliefs may be reinforced, the practical impossibility of fueling the bubble forever eventually leads prices back in line with rational valuations.

Ross M. Miller is the president of Miller Risk Advisors.

Requests for reprints should be sent to: Ross M. Miller, Miller Risk Advisors, 2255 Algonquin Road, Niskayuna, NY 12309-4711. Email: millerrm@alumni.caltech.edu

Obvious forms of intervention aimed at preventing bubbles or moderating their formation appear to be ineffectual in the laboratory. Daily price limits, circuit breakers, and restrictions on short selling may not only fail to inhibit bubbles, they can actually help to promote them by apparently providing traders with a false sense of security that helps fuel the bubble. While limitations on price declines can draw out the time it takes for the bubble to burst, the absence of market liquidity during this extended decline may make a bad situation even worse. It appears that the prevention of bubbles cannot be externally imposed on the market mechanism in an effort to put limits on its operation. Instead, a more fruitful approach might be to enhance the market mechanism in a way that harnesses its naturally tendency towards efficiency under normal circumstances as first demonstrated by Vernon Smith [1962] and reproduced countless times thereafter.

The experimental evidence pointing to the formation and persistence of market bubbles did not arise in the vacuum, but is the result of a logical progression of market experiments dating back to the original market experiment conducted by Harvard University's Edward Chamberlin in the 1940s. It took 40 years for experimental economists to introduce the necessary market features—order books, multiple periods with asset carryover, and speculative traders—to create a laboratory setting sufficiently rich to produce (and reproduce) bubbles.¹ While the road to isolating market bubbles in the laboratory followed an orderly and relative direct progression, market experimentation has spun off several alternative threads in the 1970s. One of the more notable threads, which received little attention at first but has been the source of renewed interest among experimental economists in recent years, is how learning (in a collective sense) occurs in a market setting.

A number of major advances in economic theory in the 1970s—recognized by the 2001 Nobel prize awarded to George Akerlof, Michael Spence, and Joseph Stiglitz—concern markets in which information was asymmetrical distributed, usually in such a way that the seller of an item knew more about its value than a potential buyer and had no way to convey that knowledge directly. Because informational asymmetries can seriously undermine the market mechanism, it often pays sellers to discover a way to signal the value of an item to buyers and buyers had to learn how to decode this signal.

This new thread of market experimentation that explored the specific kind of learning required to make signaling and related informational transfer mechanisms work; however, all markets can be viewed as implicitly involving learning. Even the simplest supply-and-demand experiments of Edward Chamberlin [1948] and Vernon Smith [1962] require the market to learn the proper price for an item that, in turn, determines the quantity traded at that price. As Smith discovered when

he modified Chamberlin's multilateral bargaining arrangement into an organized market patterned after the New York Stock Exchange, the ability of the market to learn prices depends on how the market is organized—a discovery that influenced the later experimental work of Smith and his many collaborators.

This article re-examines the results from the experimental studies of market bubbles in light of what we now know about how learning occurs in the context of asymmetric information. This way of viewing market bubbles suggests several new lines of bubble experiments that might be conducted to help determine what causes bubbles and how to prevent them. (The exercise of actually running these experiments is left to the reader: no new experimental results appear in this article.)

The basic answer to the question posed by the title of this article is an unsatisfying "Yes." Both inside and outside the laboratory direct personal involvement in a bubble appears to be a potent means by which the market participants can learn to avoid them in the future having, in effect, undergone a collective form of aversion therapy. There is anecdotal evidence that such learning goes on outside of the laboratory. For example, the stock market crash of October 1929 changed attitudes about investing for an entire generation, with exuberance only returning to the market in full flower during the "Go-Go Years" of the Sixties. Experimental subjects who had been through one or two experiments in which bubbles have formed and then burst similarly learn to avoid them.

Following Ben Franklin's advice that "experience keeps a dear school, but the fool will learn in no other," it would be highly beneficial to find a way to prevent bubbles without having to experience the deflation of even a single bubble first, much less the two or more that some experimental subjects appear to need to experience. Merely accelerating the existing learning process may reduce the number of crashes required to teach subjects not to collectively start a bubble in motion, it appears that something more is required if bubbles are to be eliminated entirely.

While learning to avoid bubbles poses an impressive challenge for the market mechanism, experiments involving asymmetric information indicate that very sophisticated learning is possible in a market environment. This type of sophistication is demanded by economic theories that depend on extended notions of rationality, such the "rational expectations" invoked in the monetary theories of Robert Lucas [1972]. Nonetheless, sophisticated learning does not occur automatically and how the markets are organized can play a critical role in aiding or inhibiting the type of learning required to mitigate bubbles. A careful reappraisal of the existing experimental evidence can aid in the development of ways to avoid bubbles both inside and outside of the laboratory.

How Markets Learn to Signal

Ross Miller and Charles Plott [1985] conducted the earliest market experiments that required any explicit learning by market participants.² (The “learning” of the competitive price and quantity in standard market experience is more an emergent property of the market—or “spontaneous order” as F. A. Hayek referred to it—and can occur without any apparent conscious effort of the part of subjects.) In contrast to the traditional market experiments pioneered by Vernon Smith [1962], in which buyers exchanged identical items that could be viewed as perfect substitutes for one another, the Miller–Plott experiments were the first market experiments in which items of different “qualities” were traded in the same market. In these experiments, only the seller knew the quality of an item at the point of sale, the buyer did not discover the quality until after the trade had been consummated. Without any loss of theoretical generality, these experiments limited sellers to two specific quality levels, using low-quality items called “Regulars” and high-quality items called “Supers.” Furthermore, in each period the quality that a seller could produce was exogenously determined, so that sellers assigned to produce Supers could not intentionally “rip off” buyers by delivering Regulars instead.³

Sellers of Supers were able to distinguish themselves by “signaling” that their item for sale was a Super by attaching “stripes” to it. The parameters for these signaling experiments were designed so that the unit cost of stripes was sufficiently less for sellers of Supers than it was for sellers of Regulars enabling the market to effectively separate the two groups, with Regulars selling at a lower price and with fewer stripes than Supers. These experiments were designed as a direct test of the signaling model developed by Michael Spence [1974], which has multiple Nash-like equilibria in which the signal (stripes) are used to distinguish the two qualities. Labor markets in which education was the signal inspired this model; however, the Miller–Plott experiments changed the setting to a consumer product market in order to avoid any prior associations or special expectations that subjects might attach to the roles of employer and employee.

Spence’s model and the later refinements of it do not specify how sellers learn to send the signal and buyers learn to recognize it, only that a signaling equilibrium is consistent with buyers associating the amount of the signal transmitted by the seller with the ultimate quality of the item purchased. Although the simplest form of Spence’s model provides the signal with no intrinsic value, the stripes used in the signaling experiments have enough value to buyers that sellers would provide them in a competitive equilibrium even if they were unnecessary for signaling quality. The experiments were parameterized so that the quantity of stripes required to distinguish Supers from Regular

significantly exceed the amount that would be provided in a market where buyers could know the quality of an item at the point of sale.

As in the single-product market experiments, the signaling experiments were run using a sequence of consecutive periods in order to see if repetition was sufficient to allow the market to converge towards one of the signaling equilibria. In the signaling experiments, it was necessary to assign Regulars and Supers to sellers at random each period so that buyers could only use the number of stripes and not which seller sold the unit as a signal of its quality. Furthermore, while single-product experiments, such as those pioneered by Vernon Smith, required the market to discover the values of two variables (price and quantity), signaling experiments placed a heavier information burden on the market, requiring not only that it determine the value of six variables (price, quantity, and stripes for both Supers and Regulars) but also that it establish the functional relationship between stripes and quality.

The first few signaling experiments were conducted using California Institute of Technology (Caltech) undergraduates as subjects and these initial experiments confirmed (or at least failed to refute) Spence’s signaling theory. In these experiments, the market showed a strong tendency to approach the most efficient of the signaling equilibria, i.e., the one that required the least possible number of stripes that would allow sellers of Supers to distinguish themselves from sellers of Regulars. This result was predicted by later refinements of Spence’s theory.⁴ The Caltech subject pool, however, consisted largely of individuals—aspiring scientists and engineers—who had been selected, in part, on the basis of their pattern-recognition ability, which is an important element of effective market signaling. In other subject pools, such as those recruited at community colleges, the market had a more difficult time learning that stripes could serve as a signal of quality.

The original signaling experiments were conducted in the early 1980s using blackboards and white chalk, these experiments were too complex for the computer-based experimental software that was still under development. The key technological innovation in these experiments was the use of intercoms and citizen band radios to transmit orders between buyers and sellers who were segregated from each other in different classrooms. This “channeling” of information prevented the subjects from using their voices and physical gestures to transmit information, so that any signaling that occurring was limited to market orders and transactions. The other major difference from the market experiments conducted with a single market was that instead of maintaining a single “market book” on the blackboard a separate one was kept for each possible level of stripes (usually the integers from zero to thirty). The blackboard was arranged so that the number of stripes increased as one moved from left to right. As in previ-

ous market experiments, when a transaction was consummated it was indicated by circling the relevant bid or offer on the blackboard.

There was a strong tendency for Supers to contain more stripes than Regulars regardless of subject pool or experimental parameters. This did not always reflect a conscious effort of the part of the sellers of Supers to distinguish their items, but was instead a consequence of their cost advantage for supplying stripes to the market. Furthermore, sellers did not know exactly how buyers valued stripes—the only feedback that they received was that supplied through the market mechanism. Indeed, the value of stripes to buyers was structured so that the higher prices required for sellers to sell units with several stripes at a profit would only be worthwhile to buyers if the unit was almost certain to be a Super.

It required only a minor institutional change to get the market to learn how to signal and converge toward equilibrium. At the end of each period when the quality of each unit was revealed, Regulars and Supers were circled with contrasting colors of chalk. (Previously, qualities had simply been indicated using an “R” or “S” written next to the circled transaction in white chalk.) Buyers who saw the Super color on the right side of the blackboard and the Regular color on the left would often make the connection between stripes and quality instantaneously. (This discovery was so dramatic that it called to mind the image of a bulb lighting up above their heads.)

In the absence of colored chalk, some of the signaling experiments failed to exhibit the stability that characterizes single-market experiments. Rather than converge to a signaling equilibrium, the market would simply move from one unstable and inefficient allocation to another without converging. Although these signaling experiments were parameterized so that they would each have a unique efficient signaling equilibrium and several inefficient ones, some experiments exhibited behavior that suggested the nonexistence of equilibrium. Such a situation is possible only when the signaling cost advantage for sellers of Supers is sufficiently small.⁵

The stability issues that can arise in signaling experiments can be traced to the additional burden that signaling places on the subjects. In a single-market experiment, only the bare minimum of individual economic rationality appears necessary for convergence to equilibrium—the auction mechanism itself provides so much of the impetus towards equilibrium that convergence can withstand significant deviations from rationality by the subjects.⁶ Convergence to a signaling equilibrium requires that some of the sellers learn how to send signals and that some of buyers learn how to receive them. While not every buyer and seller needs to “get the signal” in order for the market to approach a signaling equilibrium, without a critical mass of savvy

buyers and sellers driving the market it is more difficult to reach an equilibrium.

It can be useful to think of buyers and sellers in a signaling market as looking for a path that uses the signal to connect them. By altering the parameters that determine the cost of signaling and the value of quality, the experimenter can make the path wider (and easier to discover) or narrower (and more difficult to discover) or can even make it disappear entirely in the case where no equilibrium exists. The Miller–Plott experiments found that the market converged to a signaling equilibrium more quickly and with greater frequency when the path was wider. In addition, the use of colored chalk aided subjects who had become temporarily “lost” in ultimately finding the path to equilibrium.

The instability seen in the signaling experiments that arose from the failure of an effective signaling mechanism to emerge spontaneously may be less dramatic than the instability of a market bubble; however, the fact that a minor institutional change—the use of colored chalk to distinguish Supers from Regulars—can help stabilize the market and guide it towards an efficient allocation indicates that relatively trivial changes in market institutions might help prevent bubbles from forming. Neoclassical economic theory makes no provision for colored chalk or any other representation method to affect individual choices and the equilibrium allocations that are generated by them. In fact, the learning induced by colored chalk may itself be viewed as irrational because it significantly alters individual behavior without changing any of the variables considered relevant to the economic choices faced by individuals in the market.

In a signaling experiment, making the connection between stripes and quality is in an individual’s rational best interest only when the market collectively learns to make the connection. This learning can occur even if the individual subjects never become consciously aware that signaling is taking place. A seller of a Super-quality product who has figured out that Regular-quality sellers cannot economically produce units enough stripes to imitate his or her product will be unable to recoup the cost of the stripes if none of the buyers is willing to pay a premium price. It is, in fact, irrational for a seller to send a signal when there are no buyers capable of receiving it.

The instability caused by the failure of buyers and sellers to learn how to use stripes as a means of distinguishing Supers from Regulars does not appear to be related to individual deviations from rationality of the sort documented by psychologists and economists dating back to the pioneering work of Daniel Kahneman and Amos Tversky.⁷ While buyers may not be able to update the probability that a unit will be a Super conditional on the number of stripes it contains in a properly Bayesian manner, this failure comes not from shifting frames of reference, but rather from a failure to notice

the link between quality and stripes quickly enough to provide sellers with the feedback necessary to transmit the signal.

Bubbles Under the Microscope

Markets that bubble to excess and then crash exhibit their instability in a more dramatic way than markets in which signaling fails to take hold, but the underlying problem may be quite similar. For asset markets to be both efficient and stable, individuals must learn to make the connection between an asset's price and its expected future cash flows. For example, some of the most basic bubble experiments allow trade in an asset that pays a dividend of \$0.24 at the end of each of 15 periods. At the beginning of the 15-period experiment, the asset will generate a total cash flow of \$3.60 with certainty, which gives it a competitive market value of \$3.60. After each period's dividend payment of \$0.24, the "intrinsic value" of the asset based on its cash flows declines by exactly the dividend amount until it becomes worthless after the last dividend is paid and the experiment concludes. More intricate versions of this basic bubble experiment make the dividend payment uncertain in order to stimulate the natural trade that would arise between more and less risk-averse subjects, but the basic outcome of the experiment is unchanged: a bubble forms in virtually every instance.⁸

The path taken by the bubble usually follows the same general pattern. In the early periods of the experiment, the asset trades at a substantial discount to the value of its cash flows. Within a few periods, competition among subjects drives the price up to its intrinsic value, which by then has fallen from its initial level of \$3.60. For example, the price of asset might rise from \$2.20 in the Period 1 to its intrinsic value of \$2.88 in Period 4 (\$3.60 minus three dividend payments of \$0.24) at that same time that its intrinsic value has dropped from \$3.60 down to \$2.88.

Subjects who participate in this experiment for the first time appear to be learning the lesson that the price of the asset always increases by generalizing from how prices behave in the first few periods, and so they continue to bid its price up even as its intrinsic value declines steadily towards zero.⁹ The result is a bubble in which the price of the asset greatly overshoots its intrinsic value until it finally crashes below it when just a few periods are remaining in the experiment and subjects must finally face up to the reality of holding an overvalued asset. The initial rise in the asset price serves to mask the subjects' ability to learn that the market price of the asset should be near its intrinsic value. In the typical bubble experiment, it takes one or more 15-period repetitions before a given pool of subjects learns to avoid a bubble.

The instability found in bubble experiments, like that of the signaling experiments, cannot readily be traced to any inconsistency or irrationality in individual choices. The bubble experiments are designed so that the intrinsic value of the asset is never in doubt. Regardless of whether the dividend payments are deterministic or drawn from a known random distribution, most bubble experiments are conducted so that the experimenter knows that every subject has complete knowledge of the asset's intrinsic value throughout the experiment by posting the expected total payoff on each subject's computer monitor. Additionally, subjects may be asked to estimate the total value of the remaining dividend payments to make sure that they are receiving the message about the asset's value.

At the same time that subjects are learning about the asset's intrinsic value, the market teaches them two things that can undermine that knowledge. First, as the asset price moves towards equilibrium in the early periods, subjects see that prices tend to increase over time. Second, because this increase occurs as the intrinsic value is decreasing, subjects learn that the market price does not need to track the intrinsic value, at least over the short run. Until the markets crashes as the experiment nears its conclusion, subjects who learn to ignore the asset's intrinsic value are rewarded by speculative profits, while those who follow it are quickly priced out of the market. Indeed, in experiments that allow selling short, subjects who sell the asset short may not only lose money, should they liquidate their short positions too soon their purchases can help sustain the bubble.¹⁰

Filling Holes in the Market

A notable similarity between the signaling and bubble experiments is that both involve incomplete market systems. Although the signaling market facilitates trade in items with every possible number of stripes, it does not allow trade in markets for Supers and Regulars directly, creating a gap in the market system. In the presence of a mechanism that would enforce or guarantee quality, the necessary markets could exist and the signaling value of stripes would vanish. Signaling emerges as the market's way of dealing with incomplete markets for quality.

The bubble experiments are also missing keys markets, those for future delivery of the asset. In a perfect world of complete markets, subjects would be able to trade not only in the "spot" market that provides direct ownership of the asset, but also in forward markets that provide for the buyer of the contract to receive the asset from its seller in a specified future period of the experiment.¹¹ Hence, at the beginning of the experiment there would be 14 additional forward markets, one each for Period 2 through Period 15.

It is worth examining how prices in the forward markets should behave in the simple case of a fixed dividend of \$0.24 at the end of each period. Consider the price of a forward contract for delivery in Period 2 that trades during Period 1. Although prices in bubble experiments usually move higher during the early periods, it is unlikely that the Period 2 forward contract will exceed the spot price at any time during Period 1. Were such an opportunity to present itself, a subject could purchase the asset on the spot market and simultaneously sell a forward contract at a higher price, yielding not only an immediate profit, but also the \$0.24 dividend that is paid at the end of Period 1. When one fully takes the dividend into account, not only must the Period 2 forward contract be priced at least \$0.24 less than the spot price in Period 1, but also the Period 3 forward contract must be at least \$0.24 less than the Period 2 forward contract, and so on to Period 15. If all 14 futures contracts are actively traded—as we will see below this is a very big “if”—then the fact that the Period 15 forward contract cannot have a negative price means that the spot price during Period 1 must stay above \$3.36. (If the Period 15 forward contract trades at its intrinsic value of \$0.24, then Period 1 spot price is pushed up to at least its intrinsic value of \$3.60.)

With a properly functioning set of forward contracts, the downward pressure on prices from period to period will be apparent to subjects from the beginning of the experiment. It is still possible for a bubble to form that raises the price of the asset in the spot market and all of the forward markets above their intrinsic values; however, if the pattern of declining forward prices that is easily generated by the arbitrage activities of even a single rational subject is detected by other subjects, the illusion that prices will increase from period to period, which appears necessary for the formation of bubbles, will be difficult to maintain.

Bubble experiments are already among the most complex experiments that are conducted on a regular basis and adding a full complement of forward markets further complicates them. David Porter and Vernon Smith [1995] have conducted bubble experiments in which they have added a single forward market for the asset with delivery at the midpoint of the experiment in Period 8. While this additional forward market still leaves the market system substantially incomplete, it does appear to attenuate the bubble that forms in experiments with inexperienced subjects.

The difficulty with using forward markets to prevent bubbles in both experimental and naturally-occurring markets is that these markets can exert a stabilizing influence only if they transact enough business to generate meaningful prices. In their current design, bubble experiments already provide extremely limited economic incentives to trade on the spot market because the asset has the same intrinsic value of everyone; trade in the forward markets, especially a multitude of

them, has even less motivation. Indeed, on the major futures markets of the world, many contracts go for days or weeks without a trade and provide no useful price information to the market system.

The potential for futures contracts to limit asset-pricing bubbles works somewhat differently in naturally occurring markets than it does in the laboratory. In contrast to the experimental market, in which the asset price should decline from period to period in a competitive equilibrium, the futures price of an asset that pays little or no dividends can be expected to increase over time so that the capital gains from holding it provide a suitable return on the capital invested in it. Hence, futures prices for such assets (or indexes that consist of them) increase as the delivery date moves further into the future. If the future prospects of an asset are sufficiently promising, any increase in the futures prices will tend to drag the spot price up with it.

Consider, for example, the stock of a company, let us call it *hightech.com*, that currently trades at \$50/share and that the market believes will trade at \$200/share in six months. (Such highly optimistic projections were common during the run-up in Internet-related stocks.) If a futures contract for delivery of *hightech.com* in six months were publicly available, a uniform belief that a share would be worth \$200 then would drive the futures price up towards \$200/share.¹² Such a move would be inconsistent with a current stock price of \$50 because the simultaneous purchase of the stock on the spot market and the sale of it on the futures market would give a return of \$150/share for an investment of \$50/share made over six months, which is vastly beyond any realistic cost of capital for this investment. In such a situation, the existence of the futures contract means that either the present or future assessment of the stock value must be reappraised until the spot and futures prices are brought into equilibrium. As in the experimental market, there is still the possibility that a bubble would form in which all these prices simultaneously exceed the stock's intrinsic value, but this would presumably be more difficult than generating a bubble in the spot market alone. Furthermore, to the extent that a temporary scarcity of stock issues that provide speculators with a “technology play” helps fuel the bubble, the existence of stock futures provides an inexhaustible outlet for speculation that does not carry the time premium associated with other alternatives, such as options. While many stocks would not warrant trade in futures contracts, it is likely that speculative issues would attract sufficient volume to maintain a full complement of them.

During the Internet boom years, there was no organized mechanism for trading forward contracts in individual U.S. stocks, the best that one could do was to trade in relatively illiquid “equity swaps” created by investment banks. Trade in standardized futures or forward contracts on individual shares had been made

temporarily illegal by the 1982 Shad–Johnson Accord that delineated the securities under SEC and CFTC jurisdiction and left stock futures in limbo until their status was finally resolved in 2000. Some astute financial observers believe that futures on individual Internet stocks might have prevented the Internet bubble from forming at all.¹³

To the extent that futures contracts on stock indexes were available as an alternative way to cash in on the Internet boom, they appear to have been ineffectual in bringing the bubble under control. Futures contracts on the Nasdaq Composite Index and the Nasdaq 100 were heavily traded during the period; however, because they averaged the returns from many stocks, including a significant proportion with no connection to the Internet boom, they lacked the excitement and potential stratospheric returns of individual Internet companies. Internet index futures, such as the ISDEX contract on the Kansas City Board of Trade, appeared only as the boom was reaching its crest and they failed to attract much attention from the financial media.

Efforts to prevent bubbles by making markets more complete are further complicated by the difficulty of establishing an objective intrinsic value for securities that have speculative appeal. In the signaling experiments, both Supers and Regulars had well-defined intrinsic values, all the market had to do was figure out how to encode and decode that information with stripes. While the expected future cash flows from highly speculative investments may be highly variably and subject to vast differences of opinion, an objective market consensus is still possible. What is troubling in the case of the Internet bubble was a popular line of analysis that appeared to refute the fundamental economic belief that the value of an asset should be equal to the sum of its future discounted cash flows. The need to ground valuation in tangible future returns was often dismissed as “old economic thinking” that was irrelevant to the “new economy.” While such analysis constitutes a serious departure from rationality, the lack of a valuation methodology for new technology ventures that is truly objective helps to facilitate such flights of fancy.

It is also possible for bubbles to form in which the prices of assets can be individually rationalized, but their aggregate valuation is inconsistent with any plausible market outcome. When many companies are competing for dominance in the same market and only one or two winners are likely to emerge, it is possible that a bubble will form in which every firm is valued as if it will emerge as the ultimate victor. The development of markets that will not only highlight such inconsistencies but also allow arbitrageurs to profit from them and nip any nascent bubble in the bud requires better models than those currently available. Robert Shiller has championed the use of “macro markets,” such as future contracts based on a country’s GDP, as a

way of helping to promote more rational and informative markets. Given the difficulty of correlating macro variables with specific financial assets, it is difficult to see how a market-generated prediction of aggregate economic activity could control any bubble that was less than universal in scope.

Highlighting Bubbles

Just as colored chalk could be used to help guide the market to a signaling equilibrium, it is possible that a similar method could be employed to guide the value of an asset toward its intrinsic value and away from bubbly excess. Exactly what mechanism might be able to contain the market has yet to be determined; however, all possible methods are likely to face similar challenges. Existing methods that provide information about the intrinsic value of an asset to experimental subjects has not been successful at mitigating bubbles.

A possible way of providing rational guidance to the financial markets is to report the intrinsic value of an asset along with its market price in the listings provided in newspapers and over the Internet. There is substantial precedent for this practice in closed-end mutual funds, whose intrinsic value can be explicitly determined from their asset holdings. While the tendency for the price of closed-end mutual funds to diverge from their intrinsic (or net asset) value has been acknowledged by even the most ardent advocates of the efficient-market theory, such disparities are usually so small as to not approach constituting a bubble.

An assortment of relative valuation measures, such as book value and the ratio of price to earnings, are readily available to investors; however, from the end of the 1990s and into the 2000s indications that stocks in general, and technology stocks in particular, were historically overvalued seemed to have little impact on prices. The ability of major stock indexes such as the S&P 500 Stock Index and the Nasdaq Composite Stock Index to sport new record price/earnings ratios with each passing month tended to reward investors who ignored the early signs of overvaluation and punished those who heeded it. A similar phenomenon appeared during the Japanese stock market bubble that lasted into the early 1990s. In that bubble, U.S. brokers were known to reassure their customers that the absurdly high price/earnings ratios seen on Japanese stocks merely reflected differences in accounting practices and not a reflection of a market bubble that had gone out of control.

If a respected financial publication were to institute its own version of colored chalk by highlighting overvalued stocks (e.g., ones with price/earnings or price/sales above a specified cutoff) much as it highlights stocks with unusually high volume or price movement, it would be purely a matter of chance whether its efforts

would be sufficient to prevent bubbles in those issues. Certainly, if a pattern emerged where any stock that rated a warning would immediately succumb to selling pressure that would drop its price back in line with its value, this mechanism could effectively limit stock prices. Indeed, given an objective and open method for determining overvaluation, anticipation of the selling that being a highlighted issue would bring might pre-empt most stocks from achieving this distinction.

Unfortunately, this heavy-handed approach to bubble prevention is likely to be ineffectual. There are almost certain to be stocks that might appear objectively overvalued, but deserve their high prices because of outstanding growth prospects or other special situations. Such issues could come to dominate the highlighted list since they would be able to survive the automatic selling that would materialize as they approach overvaluation. When a sufficient number of issues continued to rise once they had been highlighted, likely aided by the purchases of short-sellers covering their positions, the market would then learn that it should no longer avoid shares that appear to be overvalued. This sequence of events would effectively discredit the overvaluation method and provide the opening for truly overvalued shares to escape the discipline of the market.

While the financial media never explicitly drew the chalk lines that distinguished overvalued companies during the stock market boom of the late 1990s, the extreme overvaluation of many companies were repeatedly highlighted in only slightly less extreme manners. The ability of blatantly overvalued shares to move even higher provided positive feedback to speculative holders that masked any considerations of intrinsic value.

Conclusion

The experimental evidence compiled to date indicates that under normal market circumstances market bubbles may be difficult to eliminate. This article has examined experiments in which market efficiency relies entails learning at an aggregate level to try to gain insight into how markets might learn to avoid bubbles. Two of the more promising approaches involve patching holes in the market system with the appropriate forward or futures markets and supplementing the market with information (colored chalk) that guides prices in the right direction. These two remedies, either alone or in combination, are especially appealing because if they are shown to work in the laboratory the development of real-world policies that incorporate them can be straightforward, as opposed to the elimination of bubbles through past experience with comparable real-world crashes.

In both the business press and the legal system, much of the blame for the Internet bubble has been heaped on the stock analysts who placed exorbitant

valuations and price targets on the stock of Internet-related companies. While laboratory experiments had yet to incorporate these analysts into their design, it is clear for existing experimental results that their presence is not necessary for the creation of bubbles and that they may well simply make convenient deep-pocketed scapegoats after the fact. Although subjects in a laboratory setting can eventually learn not to get involved in bubbles, it is not so clear whether the legal and political system in which our markets exist can so easily learn to find a cure for bubbles rather than simply treat the symptoms.

Acknowledgments

The paper draws on ideas developed while writing *Paving Wall Street* [Ross Miller, 2002]. The author benefited tremendously from thoughtful discussions with Gunduz Caginalp and Vernon Smith.

Notes

1. Ross Miller [2002] contains a detailed account of the path from Edward Chamberlin's [1948] first experiments to the Vernon Smith, Gerald Suchanek, and Arlington Williams [1988] bubble experiments. Milestone experiments in development of bubble experiments are described by Vernon Smith [1962], Ross Miller, Charles Plott, and Vernon Smith [1977], and Robert Forsythe, Thomas Palfrey, and Charles Plott [1982].
2. Game-theoretic experiments in which a small group of subjects (often just two) would learn to cooperate predated experiments in which any learning was mediated by the market mechanism.
3. This moral hazard problem was examined in several experiments based on George Akerlof's [1970] lemons model by Michael Lynch, Ross Miller, Charles Plott, and Russell Porter [1991] that employ the same basic design as the signaling experiments.
4. See John Riley [1975] and Charles Wilson [1977] for more advanced version of the signaling model that eliminate many of the inefficient signaling equilibria.
5. Michael Rothschild and Joseph Stiglitz [1976] prove a theoretical demonstration of the possible nonexistence of a signaling equilibrium under normal circumstances. Although Miller and Plott had considered running such experiments, the observed behavior of markets in which a signaling equilibrium existed but could not be determined pointed out the likely futility of this avenue of investigation.
6. Dhananjay Gode and Shyam Sunder [1993] show that the rules of the standard auction-based market can guide a single market to equilibrium even with robot traders programmed to place random orders. Such results cannot be expected to carry over to markets where learning or the formation of expectations is required of the traders.
7. Amos Tversky and Daniel Kahneman [1987] provide a summary of this research and its relation to theories of economic rationality.
8. The general properties attributed to experimental bubble markets in this article drawn heavily on the papers by Vernon Smith, Gerald Suchanek, and Arlington Williams [1988], Porter and Smith [1995], and Gunduz Caginalp, David Porter, and Vernon Smith [1998 and 2000].

9. See G. Caginalp and D. Balenovich [1999] for a formal model of how momentum might drive an asset market bubble.
10. The “limits of arbitrage” issues facing short-sellers analyzed by Andrei Shleifer and Robert Vishny [1997] do not arise in these experiments.
11. The main difference between forward contracts and futures contracts is that futures contracts provide for periodic settlement of gains and losses in the value of the contract in advance of the delivery date as a way to reduce the possibility that the contract will be breached. The futures markets used in Porter and Smith [1995] are technically forward markets because settlement in cash occurs only at the end of the experiment.
12. In a world with relatively efficient futures markets, analysts’ projections that a stock will increase by 300% in six months might become pointless. Such a proclamation can be recast to the spot price by simple discounting; hence, any analyst’s statement that the market believes and is reflected in the futures price automatically implies a specific spot price.
13. Holman Jenkins [2000] advocates futures on individual stocks as a way of preventing market bubbles.

References

- Akerlof, George A. “The Market for ‘Lemons’: Qualitative Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics*, 84, (1970), pp. 488–500.
- Caginalp, G. and D. Balenovich. “Asset Flow and Momentum: Deterministic and Stochastic Equations.” *Phil. Trans. Royal Soc. London A*, 357, (1999), pp. 2119–2133.
- Caginalp, Gunduz, David Porter, and Vernon Smith. “Initial Cash/Asset Ratio and Asset Prices: An Experimental Study.” *Proceedings of the National Academy of Sciences of the United States of America*, 95, (1998), pp. 756–761.
- Caginalp, Gunduz, David Porter, and Vernon L. Smith. “Momentum and Overreaction in Experimental Asset Markets.” *International Journal of Industrial Organization*, 18, (2000), pp. 187–204.
- Chamberlin, Edward H. “An Experimental Imperfect Market.” *Journal of Political Economy*, 56, (1948), pp. 95–108.
- Forsythe, Robert, Thomas R. Palfrey, and Charles R. Plott. “Asset Valuation in an Experimental Market.” *Econometrica*, 50, (1982), pp. 537–568.
- Garber, Peter M. “Famous First Bubbles.” *Journal of Economic Perspectives*, 4, (1990), pp. 35–54.
- Gode, Dhananjay K., and Shyam Sunder. “Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality.” *Journal of Political Economy*, 101, (1993), pp. 119–137.
- Jenkins, Holman W. Jr. “Business World: Future Shock for the Next Bubble Stock?” *Wall Street Journal*, May 31, (2000), p. A27.
- Keynes, John Maynard. *The General Theory of Employment, Interest, and Money*. New York: Harcourt, Brace, 1936.
- Lucas, Robert E., Jr. “Expectations and the Neutrality of Money.” *Journal of Economic Theory*, 4, (1972), pp. 103–124.
- Lynch, Michael, Ross M. Miller, Charles R. Plott, and Russell Porter. “Product Quality, Informational Efficiency, and Regulations in Experimental Markets.” In *Research in Experimental Economics*, Vol. 4 (1991). Edited by R. M. Isaac, Greenwich, Connecticut: JAI Press, pp. 269–308.
- Miller, Ross M. *Paving Wall Street: Experimental Economics and the Quest for the Perfect Market*. New York: John Wiley & Sons, 2002.
- Miller, Ross M., and Charles R. Plott. “Product Quality Signaling in Experimental Markets.” *Econometrica*, 53, (1985), pp. 837–872.
- Miller, Ross M., Charles R. Plott, and Vernon L. Smith. “Inter-temporal Competitive Equilibrium: An Empirical Study of Speculation.” *Quarterly Journal of Economics*, 91, (1977), pp. 599–624.
- Porter, David P., and Vernon L. Smith. “Futures Contracting and Dividend Uncertainty in Experimental Asset Markets.” *Journal of Business*, 68, (1995), pp. 509–541.
- Riley, John G. “Competitive Signaling.” *Journal of Economic Theory*, 10 (1975), pp. 174–186.
- Rothschild, Michael, and Joseph Stiglitz. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information.” *Quarterly Journal of Economics*, 90, (1976), pp. 629–650.
- Shiller, Robert J. *Irrational Exuberance*. Princeton, New Jersey: Princeton University Press, 2000.
- Shleifer, Andrei, and Robert W. Vishny. “The Limits of Arbitrage,” *Journal of Finance*, 52, (1997), pp. 35–55.
- Smith, Vernon L. “An Experimental Study of Competitive Market Behavior.” *Journal of Political Economy*, 70, (1962), pp. 111–137.
- Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams, 1988. “Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets.” *Econometrica*, 56, (September), pp. 1119–1151.
- Spence, A. Michael. *Market Signaling: Informational Transfer in Hiring and Related Processes*. Cambridge, Massachusetts: Harvard University Press, 1974.
- Tversky, Amos, and Daniel Kahneman. “Rational Choice and the Framing of Decisions.” In Robin M. Hogarth and Melvin W. Reder, eds., *Rational Choice: The Contrast between Economics and Psychology*. Chicago and London: University of Chicago Press, pp. 67–94, 1987.
- Wilson, Charles A. “A Model of Insurance Markets with Incomplete Information.” *Journal of Economic Theory*, 16 (1977), pp. 167–207.