

Simulating the Madness of Crowds: Price Bubbles in an Auction-Mediated Robot Market

KEN STEIGLITZ and DANIEL SHAPIRO

Department of Computer Science, Princeton University, Princeton, NJ 08544-2087, USA
e-mail: ken@cs.princeton.edu

(Accepted in final form: 30 May 1997)

Abstract. We simulate a multiagent market with production, consumption, and exchange mediated by a sealed-bid double auction. Marked price bubbles and subsequent crashes occur when value-based (fundamentals-driven) and trend-based traders are both present, and the market equilibrium price is ramped up exogenously. Similarly, negative price bubbles and recoveries occur when the equilibrium price is ramped down. Because the simulated market is auction-mediated, we can observe the operations of traders during these events, and study the interactions that produce and resolve bubbles. Some preliminary circuit-breaker experiments are described, in which bubbles are interrupted during their formation.

Key words: negative price bubbles, price bubbles, auction-mediated robot market.

1. Introduction

Markets are subject to the formation and bursting of price bubbles, and to other seemingly autonomous violent price movements, but classical economic theory is hard-pressed to explain this behavior. Two fundamental reasons are at the root of this failure: First, conventional theory deals largely with equilibria or small excursions from equilibria, and not with essentially dynamic, often dramatic but ephemeral phenomena. Second, traditional economic analysis has not developed models of market agents who use algorithms of any sophistication or complexity – or who may even behave irrationally.

At the same time, buying frenzy, panic, and mob psychology are widely accepted as a fundamental cause of large price swings in the absence of exogenous perturbations. Charles Mackay reflects the popular wisdom in his description of the Dutch tulipmania of the 1630's [7]:

Nobles, citizens, farmers, mechanics, seamen, footmen, maid-servants, even chimney-sweeps and old clotheswomen, dabbled in tulips. People of all grades converted their property into cash, and invested it in flowers. Houses and lands were offered for sale at ruinously low prices, or assigned in payment of bargains made at the tulip-mart. Foreigners became smitten with the same frenzy, and money poured into Holland from all directions.

The bubble's burst is subsequently described in equally familiar terms:

At last, however, the more prudent began to see that this folly could not last for ever. Rich people no longer bought the flowers to keep them in their gardens, but to sell them again at cent per cent profit. It was seen that somebody must lose fearfully in the end. As this conviction spread, prices fell, and never rose again. Confidence was destroyed, and a universal panic seized upon the dealers.

Youssefmir, Huberman, and Hogg [13] provide a good review of recent economic literature concerning price bubbles. They point out the difficulty in explaining bubbles within the standard rational expectations theory, and go on to adopt the ‘noisy trader’ approach, modeling the formation of bubbles with trend-chasing traders. Their dynamic model captures an ‘interplay between fundamentalists and positive feedback trading and the relevance to the formation and breaking of bubbles [13, p. 5].’

Caginalp and Balenovitch [5] also construct a differential equation that reflects the operation of both fundamentalist and trend-chasing traders. Their work is similar to that in [13] in that it is aggregate, and does not model actual exchange, but assumes that prices move in response to buy and sell pressures. They attempt to calibrate their model to the laboratory experiments of Porter and Smith [9] and with stock market crashes since 1929.

Youssefmir et al. [13] verify the behavior of their dynamic model, a stochastic differential equation, with a multiagent simulation. The model adjusts market price when a buyer or seller ‘wakes up’ by a proportionate price movement in the appropriate direction. Thus, as Youssefmir et al. point out, their model is limited because they have not ‘incorporated the notion of variations in volume as determining current prices nor do we have any auction mechanism to mediate the purchase or selling of assets’.

The main aim of our paper is to extend the results in [13] to a market simulation in which trade is mediated in a central auction. Three kinds of agents will operate: consumer-producers (whom we term *regular agents*), fundamentalist traders (whom we term *value traders*), and speculators (whom we term *trend traders*). Although we recognize that agents in markets learn, and continually evolve, we hope our work moves a step closer to a model that reflects the operation of a real market with realistic traders. We will see that the operations of the three different kinds of agents – as reflected in supply and demand presented to the auction – will reveal the structure of both bubbles and corresponding price depressions, which we will call *negative bubbles*. The simulation also makes possible experiments with ‘circuit breakers,’ and we will present preliminary results along these lines.

Our model also includes production and consumption, and this has the advantage of defining a ‘fundamental value’: the unique equilibrium price that balances the two (given the predetermined production skill levels of the regular agents). Thus, bubbles and negative bubbles are clearly defined as excursions from fundamental value.

We mention several recent papers that are closely related to this work, and which add to accumulating evidence that heterogeneity of agents can produce significant endogenous price movements in a variety of market types.

Arthur et al. [1] describe a model and simulations in which heterogeneous agents in the Santa Fe Artificial Stock Market [8] adapt to each other's behavior using a genetic algorithm. They argue that deductive reasoning fails with heterogeneous traders, forcing the agents to use inductive reasoning. This in turn leads to what the authors term an 'ecology of beliefs' that 'co-evolves over time'. Their actual market structure and market-clearing mechanism differ from ours in many respects. Agents trade in two assets, one a security with a stochastic payoff, and the other risk-free with constant payoff. The market is cleared by a mechanism that assumes knowledge of the agents' demand functions, rather than by a sealed-bid auction. The fundamental value is defined by the rational expectations equilibrium computed from the model, rather than by long-term equilibrium of production and consumption. But what emerges in the descriptors used in the genetic algorithm is a competition between agents using 'fundamental' information, analogous to our value traders, and those using 'technical trading' information, analogous to our trend traders. Their results show that when the rate of exploration in the genetic algorithm is low, fundamental information prevails, and the market remains close to the rational expectations equilibrium – but that when the rate of exploration is high, the market enters a 'rich psychological regime' characterized by endogenous bubbles and crashes in the price difference between the observed price series and fundamental value.

In [3], Brock and Hommes study a cobweb-type demand-supply model in which agents can choose between naive expectations that have zero cost, and rational expectations with positive information costs. They show that a high propensity for switching in this model leads to highly irregular equilibrium prices converging to a strange attractor. In [4], Brock and Hommes discuss adaptive belief systems in a simple present-value asset pricing model and arrive at similar conclusions. Different bifurcation routes to complicated price fluctuations are observed with different types of traders in the market.

Bak, Paczuski, and Shubik [2] construct simple models of stock markets, including traders with imitating behavior. They show in yet another context that the interaction of fundamental value buyers and noise traders can result in large price variations not justified by fundamental market analysis.

2. The Minimal Market

2.1. OUTLINE

Our starting point is the market described in [11], and the reader is referred to that source for details. We will consider it in some sense a *minimal market*, and it has the following elements:




Figure 1. Supply curve (gold vs. price) of one regular agent who decides to sell, corresponding to the resulting single bid. The agent submits a minimum acceptable price and an amount of food available for sale.

Regular agents Each consumes one unit of food per trading period, and is endowed with two skills: the amount of food that he can produce in one trading period, and the amount of gold that he can produce in one trading period. The skills are randomly assigned from a given range at the beginning of the simulation. Each regular agent decides each trading period, on the basis of current market price, whether to produce food (farm) or gold (mine). He also decides on the basis of reserves of food and gold whether to submit a bid to the auction, and the bid is a function of his food and gold reserves.

Central auction A sealed-bid double auction. Bids to buy and sell are received from regular agents and value traders. A single market-clearing price is determined each trading period, and trades implemented at that price.

Value traders Market participants who do not produce or consume, but who estimate fundamental value using an exponential smoothing filter, and who attempt to profit by buying low and selling high. They are endowed with a distribution of buy-sell thresholds and smoothing parameters.

2.2. AUCTION

The auction maximizes total gold exchanged at each trading period.* Regular agents or traders can submit bids to the auction, and bids are of the following form:

* This is a variation of the auction in [11], where volume of exchanged food is maximized. Maximizing gold volume seems more consistent with the incentives of an auctioneer, who might well receive commissions based on sales volume in gold. But variations in the auction mechanism are found to have little effect on the general behavior of our simulated markets.




Figure 2. Demand curve (gold vs. price) of one regular agent who decides to buy, corresponding to the resulting single bid. The agent submits a maximum price, a desired amount of food, and a limit on the amount of gold available for spending. This shows the most general situation, but it may happen that the maximum price is hit before the expenditure limit, in which case there is no plateau.




Figure 3. Aggregate supply-demand curve (gold vs. price) for one auction.

Sell bids If an agent decides to submit a bid to sell, he submits a minimum price, below which he is not willing to sell, and an amount of food offered for sale. Figure 1 shows the shape of the contribution to the supply curve due to one such seller.

Buy bids If an agent decides to submit a bid to buy, he submits an amount of food desired, a maximum acceptable price, and a limit on the expenditure, also in units of gold. Figure 2 shows the resulting contribution to the total demand curve seen at the auction.

Figure 3 shows the aggregate supply-demand curve, in effect constructed by the central auctioneer. His task is to choose a price that maximizes the quantity of gold that changes hands. The supply curve consists of linearly increasing segments with increasing slopes, punctuated by upward jumps. It is monotonically strictly increasing, because each of its components is. The demand curve consists of linear segments punctuated by downward jumps. The demand in units of gold starts and ends at zero, but moves up and then eventually down in fits and starts, as individual buyers reach plateaux and price limits in their bids. It is in fact possible for there to be more than one intersection at which supply and demand are equal, but it is clear – since the supply curve is monotonically increasing – that it is the rightmost intersection that determines the desired market-clearing price. The price optimal in this sense is found by finding the points at which the aggregate supply and demand curves have points of changing slope (*critical points*), sorting these points, and then scanning the lists while searching for intersections. The algorithm we have implemented is dominated in asymptotic time complexity by the time required to sort the lists of critical points, and in practice runs in time almost linear in the number of critical points, which is in worst case proportional to the total number of agents.

Once the new price is established the bidders trade at that price in the following way: All sellers whose minimum price is above the new price and all buyers whose maximum price is below the new price are discarded. The remaining bidders are then sorted by the magnitude of their bids. The trading then begins matching the buyer with the highest bid to the seller with the lowest bid. Once a seller or a buyer has been satisfied, the trading moves on to either the buyer with the next highest bid or the seller with the next lowest bid. In this way the market is cleared until either no more selling or buying is possible. At this point the auction for the current trading period concludes and a new trading period starts.

2.3. BIDDING FUNCTIONS

The algorithm that determines the bids of value traders is exactly as given in [11], and we give it here for completeness. Each agent's particular bid is determined by multiplying his bidding function by the most recent auction closing price, P_{t-1} , the result of yesterday's auction, which is known to everyone. The bidding function of an agent depends on his food and gold inventories. For simplicity we take the bidding function to be the same for all agents.

The important characteristics of the bidding function are that the bids increase (resp. decrease) when food inventory is low (resp. high). We also want the bidding function to increase with gold inventory when bidding to buy, and to decrease with gold inventory when offering to sell. That is, as the bidder gets richer, he is willing to pay more for food when buying, and to sell food for less when he is selling. For the examples considered it was found that the particular shape of the curve is not critical, and the particular function was chosen as follows, parameterized by only

three quantities. Let $\bar{f} = f[i]/r[i]$, the food inventory of agent i normalized by his reserve; and $\bar{g} = g[i]/(P \cdot r[i])$, the gold inventory normalized by the current value of his reserve. We then choose the values of the bid function $B(\bar{f}, \bar{g})$ at $\bar{f} = 0$ and $\bar{g} = 0, 1$, and ∞ as follows:

$$\begin{aligned} B(0, 0) &= b_{00}, \\ B(0, 1) &= b_{01} \\ B(0, \infty) &= b_{0\infty} \end{aligned} \tag{1}$$

and define the bid function at $\bar{f} = 0$ by the following exponential function of \bar{g}

$$B(0, \bar{g}) = b_{0\infty} - (b_{0\infty} - b_{00})e^{-\gamma\bar{g}}, \tag{2}$$

where

$$\gamma = \ln \left(\frac{b_{0\infty} - b_{00}}{b_{0\infty} - b_{01}} \right). \tag{3}$$

The complete bid function is then taken to be the exponential function of \bar{f} that passes through the point $\bar{f} = 1, \bar{g} = 1$

$$B(\bar{f}, \bar{g}) = (B(0, \bar{g}))^{(1-\bar{f})}. \tag{4}$$

Thus at the point $\bar{f} = 1$, which corresponds to the food inventory being exactly at reserve, the bid function is always one, which means that the bid is precisely equal to the current market price P_{t-1} . When the food inventory is below reserve, the bid function yields an offer price above P_{t-1} , and when the food inventory is above it yields an asking price below P_{t-1} . Again, we remark that the details have proven not to be critical, and this function has been chosen for simplicity and transparency. Finally, the amount bid is simply the difference between the agent's current food inventory, and his reserve level.

Value traders use adaptive expectations [6] to predict the average price. The price forecast is given by

$$F_t = \alpha P_{t-1} + (1 - \alpha)F_{t-1}, \tag{5}$$

where F_t is the forecast price at time t , and α is a weighting coefficient. A decision to buy is made by value trader j when

$$P_{t-1} < F_t \cdot (1 - \text{margin}[j]) \tag{6}$$

in which case a bid of $P_{t-1} \cdot (1 + \text{margin}[j])$ is posted. The use of the same margin to trigger trade decisions and to set the bid is arbitrary, but adopted for simplicity. The logic is quite simple: if the previous period's price is sufficiently below the

forecast average, then the speculator decides to buy. Conversely, speculator j sells when

$$P_{t-1} > F_t \cdot (1 + \text{margin}[j]), \quad (7)$$

in which case an offer of $P_{t-1} \cdot (1 - \text{margin}[j])$ is posted.

3. Strategy and Parameter Space

Before we describe some simulation results, we should discuss our general approach. We can hold no reasonable hope at this point of calibrating our model with a real economic market (such as done in Caginalp and Balenovich [5]). Our model doesn't incorporate the complexity of human behavior that is evidently at work in real markets, has only the minimum number of commodities necessary for exchange (two), and is provided with no real-world data.

What we hope to accomplish is to gain insight into the qualitative dynamics of markets, and to verify some commonly held intuitions about how markets do and don't react to perturbations. We aim at what might be termed a *phenomenological* model.

For example, we will describe the effects of a mixture of value-based and trend-based trading on price movements. Just what the mixture is of these two kinds of trading sentiments depends in a real market on a myriad of factors, and trader characteristics are clearly undergoing constant evolution and selection, with complex interactions. What we can accomplish is to test the general intuition that trend-chasing can produce a positive or negative price bubble in a commodity market, and that value-based traders will return the market to fundamental value.

We will also perform some experiments for which the results are not so easily foreseen: For example, if we interrupt the trend-chasing during a bubble (trip a circuit breaker), does the system return to the same basic region of operation as it would if the bubble played out (crashed)?

Given that we are not trying to calibrate our model with real-world data, we are left with the important problem of choosing parameters. In a typical simulation, even with all the algorithms in the system fixed in form, there are many such parameters, grouped in the following categories:

- Global
 - initial distribution of gold and food
 - numbers of regular agents and traders
- Regular agent parameters
 - bidding functions
 - production skills, per trading period
 - consumption, per trading period
 - desired reserve of food

- Trader parameters
 - smoothing parameters for price and trend estimates
 - thresholds for buy-sell decisions
 - factors for determining bidding or asking price
 - caps on size of trades

There is of course much room for experimentation guided by accumulated intuition, and no hope of exploring such a gigantic parameter space completely. We do adhere to one important principle: *The qualitative results should be robust with respect to moderate variations in any one parameter.* Otherwise it is easy to be trapped into adopting a ‘peculiar’ system – one whose operation depends critically on some particular combination of parameters, and which is therefore unlikely to be reflected in reality. Thus, the results we present are reasonably insensitive to the parameters listed below. For example, scaling the number of agents by a factor of two (or even ten), or changing skills by 10 percent results in similar qualitative behavior. We give below values for some parameters that are fixed for all the experiments described in this paper (unless otherwise noted).

- initial gold stock of all agents = 60 units
- food consumption = 1 unit per trading period
- initial food stock of regular agents = reserve, uniformly randomly distributed between 15 and 20 units
- number of regular agents = number of value traders = number of trend traders (when present) = 25
- bidding function uses

$$b_{00} = 4.0,$$

$$b_{01} = 8.0,$$

$$b_{0\infty} = 16.0,$$

as in [11]

- gold skills uniformly randomly distributed between 3.0 and 4.0 units per trading period
- food skills uniformly randomly distributed between 1.25 and 1.75 units per trading period
- value traders use a smoothing parameter $\alpha = 0.008$; they bid and ask at factors $(1 + \textit{margin})$ and $(1 - \textit{margin})$ times the current market price, respectively, where the *margins* are fixed at the beginning of the run by uniformly dividing the interval [0.0, 0.5].

Another principle we have adopted: We develop more complicated systems by building on simpler ones. Thus, we first make sure that the system with only regular agents behaves as expected (it oscillates strongly), then that the addition of value traders stabilizes the price to a large extent [11]. Only then do we dare add trend traders.




Figure 4. History of auction closing price in a baseline economy: only value traders are present, and they are introduced at trading period 500.




Figure 5. Closeup of the price history shown in Figure 4.

4. Value Traders and a Baseline Economy

In keeping with our intent to build from simpler to more complex systems, we begin with a *baseline* economy, which consists of regular agents and only value traders. As mentioned earlier, the simulation algorithm is the same as that described in [11], except that the auction has been re-coded to maximize gold instead of food volume. The results are very similar. Figure 4 shows the history of the auction closing price, with value traders added at the 500th trading period. The price is




Figure 6. Volume history corresponding to the price history in Figure 5.

stabilized in about 1000 trading periods, and fluctuates in a band between about 1.8 and 3.4, whereas the computed equilibrium price is 2.41.

From the closeup of the price history shown in Figure 5 we see a bursty, chaotic-looking behavior that is characteristic of all our simulations. Using the parlance of chaotic systems informally, we can say the price switches at irregular intervals between regions of phase space, which can be thought of as attractors. Quite naturally, the volume, shown in closeup in Figure 6, behaves in a way highly correlated with price. Note that there are two brief periods when the price and volume are both absolutely constant.

The general behavior is reminiscent of the results reported by Youssefmir and Huberman [12], who simulate large systems of agents who switch strategies to compete for limited resources. They report ‘bursts of activity superimposed on a background of small fluctuations around the predicted equilibrium value.’ A similar appraisal might be made of our volume history, except our system appears to switch among several regions of phase space, rather than between a near-equilibrium state and large excursions.

Without the value traders the market price oscillates wildly, as can be seen from the first 500 trading periods in Figure 4. This is a consequence of the regular agents having no foresight, and the oscillation is analogous to that seen in a cobweb model. Oversupply alternates with undersupply, and agents are periodically forced to switch their strategy – farming or mining – to accommodate. Thus the ‘intelligence’ is allocated to the value traders, who forecast the price using an exponential smoothing algorithm, and who stabilize the price, and create a liquid market where agents can better specialize in the activity of higher skill. As we show in [11], the value traders increase overall productivity, even after their profit

is accounted for. The more efficient division of labor is, after all, the source of the ‘wealth of nations’ identified by Adam Smith [10].

5. Trend Traders

5.1. TREND-BASED TRADING STRATEGY

We next introduce trend-based traders, who are intended to model the behavior of traders motivated by price movement alone, rather than any estimate of fundamental value. We implemented them by computing a running estimate of price trend, and then triggering buying or selling when this estimate is above a positive threshold or below a negative threshold, respectively. The actual thresholds were randomly distributed across the trend traders, to avoid any artificial synchronization of action.

The details are as follows: An exponentially smoothed price estimate is already computed by value traders, using the adaptive expectations update formula

$$F_t = \alpha P_{t-1} + (1 - \alpha)F_{t-1}, \quad (8)$$

where P_t is the market price (the most recent auction closing) at time t , F_t is the forecast, and α is a smoothing constant, fixed at 0.008 for our experiments.

The trend traders compute an exponentially smoothed estimate of the trend, Δ_t , using the same method applied to first differences of the forecast F_t . That is, they use the update formula

$$\Delta_t = \alpha(F_t - F_{t-1}) + (1 - \alpha)\Delta_{t-1}. \quad (9)$$

5.2. FORCING TRENDS

We next faced an obstacle to creating a mixture of active value and trend traders: In a world with fixed skill levels and hence fixed equilibrium price, the trend traders could not compete with the value traders, and quickly went bankrupt. To ensure that they survived it was found necessary to introduce real trends in the equilibrium price. We did this by simply ramping the gold skills of all regular agents up and down in a triangular pattern, in our experiments completing one cycle every 5000 trading periods. This scales the equilibrium price, periodically inflating and deflating the value of food in terms of gold.

Figure 7 shows the forecast price and trend in a simulation over two periods of value-ramping. We see that the time constant used, $\alpha = 0.008$, provides a good tradeoff between smoothing and speed of response. The estimate of trend is a good deal noisier than the estimate of price, as expected. We will use the results of this particular simulation throughout the rest of this paper.

5.3. TUNING COMPETITIVENESS

Our goal is to create a simulation in which the two kinds of traders compete, and hence in which they both survive. The exogenous ramping of the equilibrium price




Figure 7. Forecast price and trend. The equilibrium price is being forced up and down every 5000 trading periods.

up and down provides some objective support for the trend traders’ strategy of buying when the price is rising and selling when the price is falling. We provided additional tuning of the relative effectiveness of the two types of traders by introducing caps on the size of transactions allowed. In the experiments described here we used caps of $40 \cdot P_t$ and $4 \cdot P_t$ gold units for value and trend traders, respectively, where P_t is the most recent closing price of food. That is, the cost in gold of 40 and 4 trading period’s worth of food.

This problem of ensuring the survival of both kinds of traders is rather delicate in an economy where price is determined by exchange in a market, and the commodity in question is both produced and consumed. In laboratory experiments the price and fundamental value of a commodity can simply be stipulated by the experimenter, and there is no need for traders to regulate production and the division of labor and to survive. And in simulations not mediated by auction and without exchange, there is again no problem of trader survival.

Figure 8 shows the total wealth of the two types of traders as time progresses through three cycles of the skill changes. ‘Wealth’ here is defined as the total of gold reserves and the value of food stocks at current market price. The wealth ramps up and down with the market price of food reserves, but also shows a consistent rise in absolute level with each cycle, demonstrating that the traders are profiting from their trades. They are also roughly equally successful. However, as the trend traders accumulate capital, they suffer successively greater sudden losses at the turnaround, the midpoint of the cycle when the direction of prices reverses. In fact, at the turnaround at trading period 1250 the trend traders get caught badly by the sudden reversal and the value traders profit. We see next that this epoch corresponds




Figure 8. Net worth of traders vs. time, evaluated at current market price in units of gold.




Figure 9. History of auction closing price for two cycles of the periodic ramping of equilibrium price.

to a negative bubble, with the selling pressure from trend traders depressing the market price.

6. Price Bubbles

Figure 9 shows the price history of the first two complete cycles in the experiment we are carrying forward. The band of chaotic price fluctuations appears to expand




Figure 10. Close up of the price bubble starting at trading period 10271 in Figure 9.

as the equilibrium price increases, but it actually stays approximately constant in width as a percentage of current market price. What is immediately evident is that a sharp positive bubble appears near the beginning of the ramp-up of the second cycle, and a similarly sharp negative bubble appears near the beginning of the ramp-down part of the cycle. The third cycle also has even larger bubbles, the fourth cycle has no pronounced positive bubble, but an even larger negative bubble, and the fifth cycle has smaller bubbles of both types. Experimentation has shown that the appearance of the bubbles depends strongly on the relative assets of the two types of traders, which must be substantial before they can precipitate large price excursions. The cycle between 10,000 and 15,000 trading periods shows typical bubbles and we chose it for an ongoing example.

At this point we should note that we distinguish between a ‘crash’ and a ‘negative bubble’. We reserve the term ‘crash’ for a sudden return to fundamental value at the end of a positive price bubble, as determined by the equilibrium price; whereas we use the term ‘negative bubble’ to mean a sudden drop below fundamental value. We can call the return from a negative bubble a ‘recovery’. This terminology seems to be closest to common usage, where ‘bubble’ almost always means a positive price excursion followed by a crash.

Laboratory experiments (like [9], for example) and mathematical models (like [5] and [13], for example) concentrate on positive bubbles. In our model we observe both positive and negative price bubbles, because the trend traders sell on detecting a down-trend as readily as they buy on detecting an up-trend. We have not yet explored the effects of making the trend traders’ strategies asymmetric with respect to buying and selling. As we see below, our system is susceptible to run-downs as well as run-ups of price by the trend traders.




Figure 11. Volume corresponding to the price bubble in Figure 10.

We will concentrate in this section on the bubble that occurs at the beginning of the second driven price cycle, shown in detail in Figure 10. It rises sharply from trading period 10273 to period 10282, and drops even more sharply from 10283 to 10288, for a total duration of about 15 trading periods. It may appear that the working out of the bubble has temporarily decreased the price volatility, but it is dangerous to generalize because the normal price behavior has a strong chaotic component.

Given that we have built in trend-chasing traders, most readers will have a ready intuitive explanation for such a sudden excursion and correction in a commodities market, and we now set about verifying that the expected mechanisms are in fact at work. We do this by analyzing data collected during the simulation about transactions mediated by the auction. We'll see that implementing the agent-level actions of production, consumption, buying and selling permits an almost anthropomorphic view of the market operation.

The next most natural time series to look at is perhaps the total auction volume, the amount of food that changes hands each auction period, shown in Figure 11. This shows that there is volatile and high volume immediately before the climb of the bubble, that the volume remains flat during the rise, drops to zero during the bubble's collapse, then spikes sharply at the end of the correction. The total volume is not very illuminating by itself, however, since it doesn't reveal the nature of the transactions taking place.

Figures 12 and 13 show the supply and demand curves of the auction during the period around the bubble, broken down by type of agent and trader. This now tells us a great deal: Most important, we see that the trend traders, and they alone, do in fact provide demand, and hence upward price pressure, precisely during the




Figure 12. Supply (sell bids) in the auction during the bubble in Figure 10.




Figure 13. Demand (buy bids) in the auction during the bubble in Figure 10.

rise of the bubble. It is also clear that the regular traders provide the only supply, and hence downward price pressure, during the bubble's collapse. But these time series still do not reveal the details of the bubble's anatomy.

The critical information is provided by the history of agents' and traders' gold and food holdings ('stock'). Figures 14 and 15 show those curves, broken down by agent and type of trader. We see immediately from Figure 14 that in the period immediately preceding the bubble, the value traders are selling off to the trend traders, while the regular agents keep relatively stable food stock. This exchange




Figure 14. Food stocks during the bubble in Figure 10.




Figure 15. Gold stocks during the bubble in Figure 10. The gold stock of the regular agents is off scale and is not shown.

is also reflected in the gold stocks shown in Figure 15. Two critical events trigger the rise and fall of the bubble: The rise is triggered at trading period 10273 when the value traders run out of food to sell to the trend traders. This causes the demand of the trend traders to go unmet, and the price runs up. But the trend traders keep buying, so they must buy from the regular agents, who do offer a supply, as we saw in the discussion of the supply-demand curves.

The second critical moment is reached when the trend traders run out of gold, at trading period 10282. The demand suddenly disappears – one could say that the bottom drops out of the market – and a glut of food is produced, which is begun to be bought up by the gold-rich value traders. The overproduction of food is encouraged by the high price, since the regular agents look to the market price of food to decide whether to farm or mine. Thus, the bubble has deceived some regular agents into producing excess food, and this extra food can be seen as a triangular bump in the food supply of the regular agents. It is quickly bought up, however, by the value traders, when the price drops so suddenly.

We can thus verify that the market is doing what we think it should be doing, an important step in moving forward to more complicated simulations. To summarize, a bubble forms and collapses as follows:

- The trend traders, sensing an upward trend, exert buy pressure, but the price remains near fundamental value because the value traders are willing to sell at the rising price.
- The bubble begins at the critical moment when the value traders exhaust their supply of food, at which point the trend traders can keep buying from the regular agents, but at a lower rate.
- The bubble collapses when the trend traders exhaust their supply of gold, because demand disappears, and the regular agents overproduce because of the artificially high price.

Finally, notice that the bubble can form in this way only if the resources of the value and trend traders are balanced in such a way that the value traders run out of food *before* the trend traders run out of gold. This explains why bubbles do not appear more frequently in our simulations.

7. Negative Price Bubbles

Figure 16 shows the negative bubble that occurs in the second half of the same driven price cycle as the positive bubble discussed in the previous section. It begins at trading period 12721, falls to a minimum at trading period 12729, and then recovers very quickly, with an overshoot at trading period 12734.

The anatomy of this negative bubble is partly symmetric to that of the positive bubble, but it differs in some interesting and important ways. Figures 16–21 show the volume, supply, demand, food stock, and gold stock for the negative bubble, in the same way as in previous section.

In this case there is no period of trading between the trend and value traders before the negative bubble; in fact, the trend traders have no gold, but only a large supply of food. In this part of the driven cycle, the equilibrium price is actually falling, and the value traders therefore buy at a moderate rate, from the regular agents, it turns out.




Figure 16. Close up of the negative price bubble starting at trading period 12721 in Figure 9.




Figure 17. Volume corresponding to the negative price bubble in Figure 16.

At some point (beginning at trading period 12721) the trend traders are able to unload a large quantity of food to the value traders. While this dumping is going on, the price falls sharply, forming the negative bubble. The critical event that signals the end of the bubble is the trend traders running out of food. (Compare this to the symmetrical event that ends the positive bubble: the trend traders running out of gold.) The low price has caused the underproduction of food, and the price springs back with an overshoot.




Figure 18. Supply (sell bids) in the auction during the negative bubble in Figure 16.




Figure 19. Demand (buy bids) in the auction during the negative bubble in Figure 16.

The conditions that trigger the start of the price collapse are not so clear as the value traders running out of food in the case of a positive bubble. Somehow the trend traders suddenly are able to sell off to the value traders. It must be that the two types of traders perceive the price as falling, because this encourages the trend traders to sell and the value traders to buy. In both kinds of bubbles, the positive feedback acting on the trend traders – buying pressure causes rising price causes buying pressure – or vice versa – fuels the sudden bubble formation.




Figure 20. Food stocks during the negative bubble in Figure 16.




Figure 21. Gold stocks during the negative bubble in Figure 16. The gold stock of the regular agents is off scale and is not shown.

8. Circuit-Breaker Experiments

We conclude with some preliminary experiments on interrupting bubbles during their formation. Figures 22 and 23 show the price histories when the positive and negative bubbles we've been studying are prematurely burst by turning off the trend traders. This provide more evidence, if that is necessary, that the trend traders in fact cause the bubbles. But such experiments are obviously interesting to those who wish to stabilize markets by intervention, and simulations such as ours may




Figure 22. The same bubble as in Figure 10 except the operations of the trend traders are interrupted during the formation of the bubble.

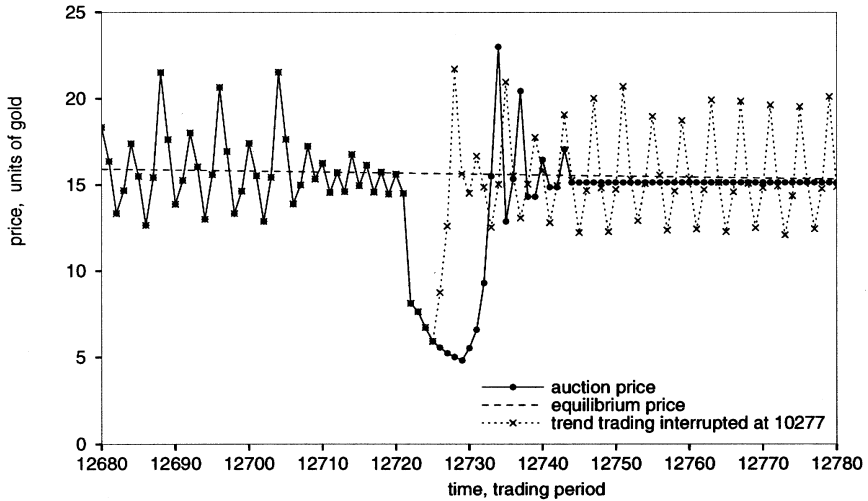


Figure 23. The same negative bubble as in Figure 16 except the operations of the trend traders are interrupted during the formation of the bubble.

provide a tool for studying the efficacy of such policy. For example, the price histories in Figures 22 and 23 show that the intervention does return the price to fundamental value sooner than without the intervention in the two cases studied here. As mentioned before, it is tempting to conclude also that there is a price to be paid in volatility after the correction, but much more work needs to be done on the whole question of volatility before any conclusions like that can be drawn.

9. Conclusions and Discussion

Our main contribution has been to confirm earlier work in accounting for price bubbles by trend-chasing traders – the ‘madness of crowds’ described so vividly by Mackay more than 150 years ago – in a simulation with auction-mediated transactions. Trading a commodity at auction makes the simulated market more difficult to analyze theoretically, but also makes it closer to reality. The auction in combination with the production-consumption economy also defines an equilibrium price, and hence a fundamental value, which in turn clarifies the meaning of the term ‘price bubble’ itself.

We have also demonstrated corresponding negative bubbles – sudden and violent excursions below fundamental value – when the price is generally descending instead of rising. One could say that our positive bubbles occur in bull markets and our negative bubbles occur in bear markets.

The fact that the trend traders are responsible for the observed bubbles is easily verified: the bubbles collapse when the trend traders are removed, and it is they who supply the upward (buying) price pressure during the formation of a positive bubble, and the downward (selling) price pressure during the formation of a negative bubble. The experiments with interrupted trading illustrate the flexibility of the simulation approach, and may ultimately provide some insight into the design of circuit breakers in real markets.

It is easy to think of ways in which this line of work can be elaborated and extended: one could try to model technological advance, for example, which would make gold have time value as an investment. That would lead to a natural interest rate. Or one could study the effect of set-up costs, or noisy information. The difficulty is not in finding interesting directions to explore, but in finding methodologies that can make it easier. At this point the construction and evaluation of models like the one described in this paper requires constant verification that qualitative results are not overly sensitive to arbitrary decisions. The process thus demands many trials to develop intuition about, and confidence in, any particular model, and that intuition and confidence is difficult to quantify and to convey to others.

Perhaps the most promising and interesting direction is the development of adaptive and learning agents. The hope is that they can replace much of the trial and error in model construction by artificial intelligence, and at the same time (and this is more to the point) come closer to the real-world behavior of economic actors.

With all its difficulties and subtleties, the field of multiagent simulation is nevertheless emerging as a flexible and powerful tool for studying economic phenomena, complementary to – and in some sense between – theory and laboratory experiment.

Acknowledgements

The authors are first of all indebted to the co-authors of [11], Leonard Cohen and Mike Honig. A succession of students contributed to the development of code and

the auction algorithm: Isaac Chang, Matt Lawson, and Michael Weiksner. We also thank Andrew Appel, Gunduz Caginalp, Wenjia Fang, Bernardo Huberman, and Tim Van Zandt for helpful conversations along the way.

References

1. Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and Taylor, P. 1996, Asset pricing under endogenous expectations in an artificial market, Santa Fe Institute.
2. Bak, P., Paczuski, M. and Shubik, M. 1996, Price variations in a stock market with many agents, Santa Fe Institute Working Paper 96-09-075.
3. Brock, W.A. and Hommes, C.H. 1996, A rational route to randomness, to appear in *Econometrica*.
4. Brock, W.A. and Hommes, C.H. 1996, Models of complexity in economics and finance, Report AE 5/96, Institute of Actuarial Science and Econometrics, University of Amsterdam.
5. Caginalp, G. and Balenovich, D. 1993, Market oscillations induced by the competition between value-based and trend-based investment strategies, Economic science association fall meeting, October 21–23, Tucson, Arizona.
6. Carlson, J.A. 1967, The stability of an experimental market with a supply-response lag, *Southern Economic Journal*, **23** (1967), 305–321.
7. Mackay, C. 1995, *Extraordinary Popular Delusions and the Madness of Crowds*, Wordsworth Editions Ltd. Hertfordshire, England, 1995 (originally published 1841).
8. Palmer, R., Arthur, W.B., Holland, J.H., LeBaron, B. and Taylor, P. 1994, Artificial economic life: A simple model of a stockmarket. *Physica D*, **75**, 264–274.
9. Porter, D.P. and Smith, V.L. 1994, Stock market bubbles in the laboratory, California Institute of Technology, preprint.
10. Smith, A. 1970, *The Wealth of Nations*, Books I-III, Penguin Books Ltd., Harmondsworth, Middlesex, England, 1970 (originally published 1776).
11. Steiglitz, K., Honig, M.L. and Cohen, L.M. 1996, A computational market model based on individual action, K. Steiglitz, M.L. Honig, and L.M. Cohen, Chapter 1 in *Market-Based Control: A Paradigm for Distributed Resource Allocation*, Scott Clearwater (ed.), World Scientific, Hong Kong.
12. Youssefmir, M. and Huberman, B.A. Clustered Volatility in Multiagent Dynamics, *Journal of Economic Behavior and Organizations*, in press.
13. Youssefmir, M., Huberman, B.A. and Hogg, T. 1994, Bubbles and market crashes, Dynamics of Computation Group, Xerox Palo Alto Research Center, Palo Alto, CA.