Effects of price signal choices on market stability

Hideyuki Mizuta [a,*], Ken Steiglitz [b], Erez Lirov [b]

[a] *Tokyo Research Laboratory, IBM Japan, 1623-14 Shimotsuruma, Yamato, Kanagawa 242-8502, Japan*
[b] *Computer Science Department, Princeton University, 35 Older Street, Princeton, NJ, USA*

Abstract

Using simulation and analysis we show that agent-based auction-cleared automated markets can be stabilized using *only* completely myopic agents (without value traders), if these naïve agents are provided with a price signal that reflects order book information. This demonstrates that information in the order book is extremely valuable, that prediction can be replaced by better instantaneous information about others' bids, and suggests new, more stable algorithms for market-based control. © 2003 Elsevier Science B.V. All rights reserved.

1. Introduction

1.1. The problem of market stability

The stability of prices in asset markets is clearly a central issue in economics. From a systems point of view markets inevitably entail the feedback of information in the form of price signals and, like all feedback systems, may exhibit unstable behavior. Under varying circumstances we might expect convergence to some fundamental value, more or less regular oscillations, chaotic oscillations, sharp rises or falls followed by crashes or recoveries, and so on. Many writers have studied the effects of trading institutions, trader behavior, and feedback signals on such complex dynamic behavior, but the general problem remains poorly understood. A classic dialogue about this issue can be seen, for example, in the views of Friedman (1953) who argues that rational profit-seeking trading will always tend to stabilize a free market, and a long succession of others (see, for example, Baumol, 1957

---

* Corresponding author. Tel.: +81-46-215-4965; fax: +81-46-273-7428.
E-mail address: e28193@jp.ibm.com (H. Mizuta).

and de Long et al., 1990) who present models and accompanying arguments supporting the idea that speculating traders who seek to maximize their profit can in some natural circumstances destabilize a market.

In this paper, we study an agent-based simulation and focus on one particular question: How is dynamic behavior affected when the price signal supplied to the agents is changed? Briefly stated, our main result is that a signal that is apparently only slightly richer in information than the ticker price can dramatically stabilize our market—even when traders operate with no planning or foresight whatsoever.

In the next subsection we will briefly summarize the methods of attack on general questions of market stability and review previous work using what are called agent-based (or microscopic) simulations. We will then describe the construction and general characteristics of our own model.

1.2. Review of related work

The study of price movements in asset markets is remarkably complex: it combines the problems of modeling human behavior with those of predicting the dynamic behavior of very large, very nonlinear systems. Current approaches to the problem can be roughly classified as follows:

(a) Theoretical (analysis of mathematical models, usually using difference and differential equations, and usually using aggregate variables).
(b) Empirical (econometric studies using real data).
(c) Experimental (laboratory studies using human subjects).
(d) Computational (simulations modeling the actions of individual agents—the approach of the present paper).

Each has its advantages and disadvantages, and in some sense they are complementary, contributing different and overlapping pieces to the puzzle. We next briefly summarize previous work in these areas with the goal of putting our own work in context.

Theory, the first approach, is the oldest and most traditional in economics. It has the important advantages of generality, and as all theory, it can guide intuition as well as provide special tools for prediction and institutional design. The limitations of theory are equally clear. It is all too easy to formulate reasonable equations that are beyond the reach of current solution techniques. This is especially the case when studying markets with heterogeneous agents and highly nonlinear trading rules. It is often necessary to simplify and aggregate behavior to get results. The work of Caginalp and Balenovich (1994, 1996), which uses a set of coupled nonlinear differential equations, is representative of this approach applied to the study of market dynamics.

The second approach, empirical studies of asset prices, uses both conventional statistical approaches and nonlinear dynamic models. The work centers on testing for the existence of predictable structures in all kinds in time series. For a good review, especially of the work on chaotic models, see Brock et al. (1991). Specifically, a number of studies in econophysics (for example, Mantegna and Stanley, 2000) have used concepts from statistical physics and critical phenomena to study self-similarity and fat-tail distributions in empirical data.

The third approach, experimental economics, has the advantage of addressing more directly questions of human behavior. However, it is expensive, time-consuming, and it is difficult to ensure that people behave the same way under laboratory conditions as they do in real markets. Perhaps the most influential work is that of Smith, Plott, and their coworkers (Forsythe et al., 1982; Smith et al., 1988; Smith, 1989; Porter and Smith, 1994; Caginalp et al., 1998), which centers on the reproducibility of price bubbles. Along the same lines, the collection of papers edited by Stiglitz (1990) on price bubbles is revealing in its diversity of perspectives on just how a price bubble might be defined and whether in fact one can exist at all.

Large-scale agent-based simulation, the fourth approach and the one used in this paper, has become possible only relatively recently with the advent of fast, cheap, and readily available computers. It has been championed by physicists using the paradigm of computational statistical physics. For example, de Oliveira et al. (1999) review several papers over the past few years that exemplify the methodology, especially the work of Levy et al. (1994). The reader is also referred to the recent paper of LeBaron et al. (1999), which also contains many references to other work in this emerging field. The defining characteristic of the methodology is that the actions of *individuals* are simulated, explaining the term *microscopic*. This opens the door to the study of the interaction of large numbers of heterogeneous, interacting agents.

An important theme that runs through much of the work in market dynamics is the interaction between two kinds of traders: those who trade on "fundamentals" and those who trade on "technical" information. The former are often called *value traders*, and the latter *noise traders*, which include trend chasers (also called chartists). This interaction accounts for the appearance of price bubbles in the simulations of Levy et al. (1994), Youssefmir et al. (1996), and Steiglitz and his coworkers (Steiglitz et al., 1996; Steiglitz and O'Callighan, 1997; Steiglitz and Shapiro, 1998), for example, as well as the aggregate models of Caginalp and Balenovich (1994, 1996).

We mention important applications of agent-based simulations that are not directly economic in nature: they can be translated literally into algorithms for distributed control of resources (see, for example, the book edited by Clearwater, 1996). In these cases the agents may well be distributed software agents instead of humans. Examples include computing cycles (Waldspurger et al., 1992), network bandwidth, computer memory, electric power (Ygge, 1998), or even thermal energy in a building. These applications need not necessarily model realistic markets, but stability is obviously a key issue. More recently, Kephart et al. (1998) anticipate the emergence of an open, free-market information economy of automated agents buying and selling a rich variety of information goods and services on the Internet. To characterize and understand the dynamic behavior of such information economies, they very naturally employ agent-based simulation, and also use game theoretic analysis to investigate strategies and competition of software agents. As before, these markets do not necessarily behave the way human markets do, but an understanding of stability is crucial.

1.3. Description of our model

The simulation model we use in this paper is a direct descendant of those described in Steiglitz et al. (1996), and we outline its features in this section. The philosophy is to build

the simplest possible system that can reasonably be thought of as a *complete* economy: in some sense a *minimal* economy. Trade requires at least two commodities, so we use the minimum of two, which we call *food* and *gold*. Gold plays the role of numeraire, and the price of food is therefore measured in units of gold.

In general situation there are three types of agents: regular agents, value traders, and trend traders. Regular agents can produce food or gold and consume food; value traders and trend traders are solely speculators and play the roles of value and noise traders mentioned above. The regular agents are completely *myopic*; that is, they exercise no foresight or planning.

One trading period of the market simulation is executed as follows. The central market sends to each agent a Request For Bid (RFB) containing price signals. Consider first the case when the price signal is simply the previous closing price. Based on this signal, the regular agents decide on their levels of production for that time step, the value traders update their estimate of fundamental value, and the trend traders update their estimates of price trend. The agents then send bids to sell or buy according to their food inventory (regular agent), the difference between the market price and estimated fundamental price (value trader), or the direction of the trend (trend trader). Finally, the market treats the submitted bids as a sealed-bid double auction and determines a single price that maximizes the total amount of food to be exchanged. This institution is sometimes called a *clearing house* or *call market* as opposed to an *open-outcry market* (Friedman and Rust, 1993). The market-clearing price (*ticker* price) becomes the next signal in the RFB. Note that in Steiglitz and O'Callighan (1997) and Steiglitz and Shapiro (1998) the auctioneer determines the price to maximize the total amount of gold to be exchanged. However, in practice this difference has little effect on the overall qualitative results. Fig. 1 shows the derivation of the supply–demand curves and market-clearing price in such an auction.

Consider next the regular agents. They follow a simple dichotomous algorithm: in each trading period they can produce either food or gold. They make this production decision to maximize profit, but in a shortsighted way, based only on the current price. Heterogeneity
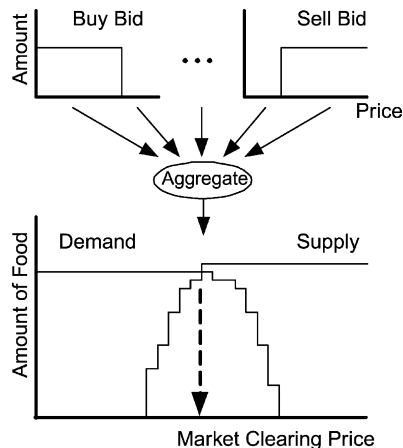

Fig. 1. Generation of the supply and demand curves and market-clearing price in the double auction.

is introduced by endowing agents with different "skills"—the amount of food and gold they can produce per period. In a similarly shortsighted way they determine their bids to maintain a fixed food inventory, based only on their current inventory. The regular agents therefore have no memory or foresight. Their strategy is so simple and myopic that it often throws the market into confusion, in a way reminiscent of the cobweb model (Carlson, 1967).

We note that our model has a natural equilibrium price, or fundamental value, determined by the equilibrium condition that total food produced is equal to the total food consumed. This is one way that our model is distinguished from that of Levy et al. (1994), which gives agents a choice between investments with certain and uncertain returns.

The remainder of the paper is organized as follows: in Section 2 we describe the results of simulations using the original model, with market-clearing price as the signal, illustrating the stabilizing effect of value traders and the destabilizing effect of trend traders. In Section 3 we describe the effects of using other price signals, specifically stabilization without traders using unweighted and inversely weighted bid averages. Then, after some concluding remarks, we present in Appendix A simplified model and its analysis, confirming the results of the simulations.

2. Simulations

Markets with only such simple regular agents exhibit large price oscillations (see Fig. 2). In these markets there is low trading volume, and most of the time there is a large overall
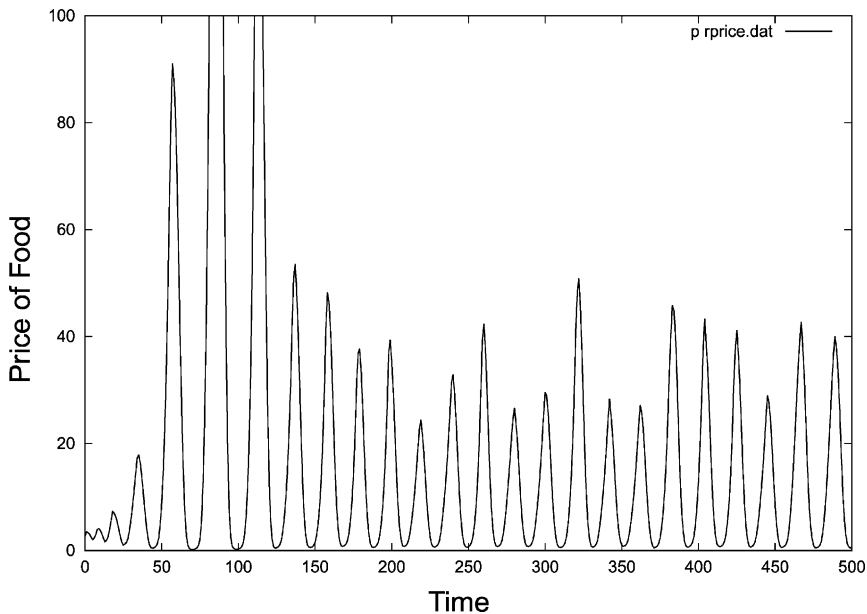

Fig. 2. Price vs. trading period with regular agents only and using closing price as a signal.
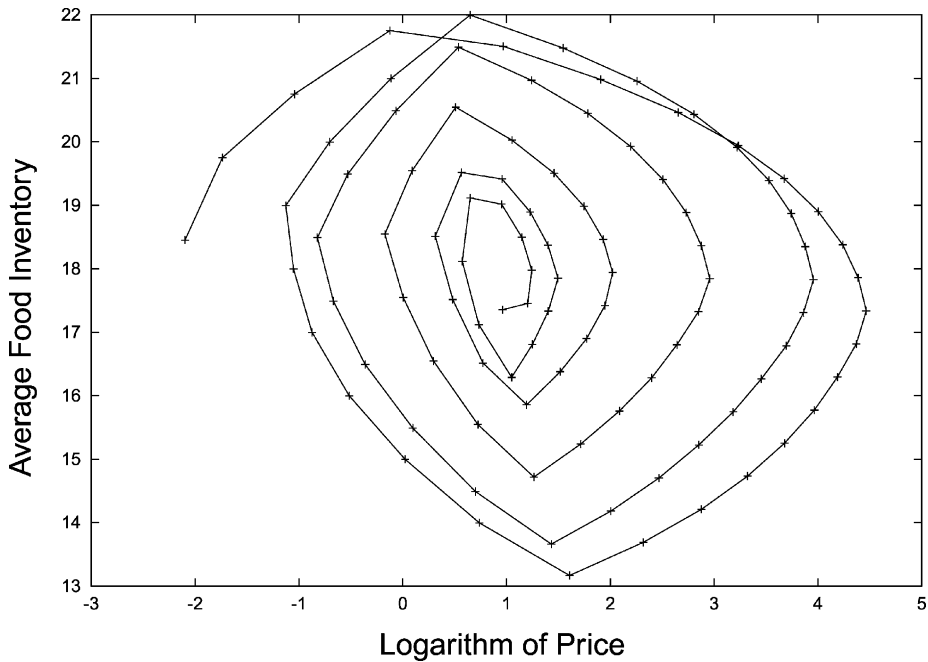
Fig. 3. Average food inventory vs. logarithm of price in the same simulation as the previous figure, illustrating the oscillation.

surplus or shortage of food. This oscillation can be visualized effectively by plotting a two-dimensional graph of average food inventory versus log-price. The result is a diamond-shaped cycle whose center is the ideal (equilibrium) price and ideal (desired) reserve (see Fig. 3). This cycle starts close to the center and rotates counterclockwise with gradually increasing radius. We cannot expect efficient resource allocation in such markets.

Fig. 4 shows a typical cycle of the oscillation, sketched diagrammatically in the food inventory–price plane. We divide the cycle into four regions. In region I, the low price prevents agents from producing food and the resulting deficiency of food causes the price to rise. In region II, when the price gets high enough, agents begin to produce food, but the price keeps rising since there still is not enough food to satisfy demand. In region III, agents now have enough food and the price begins to fall. However, they continue to produce food because the price remains high for a time. In region IV, agents stop producing food because the price finally becomes low. But the price continues to fall because of food surplus. It is therefore the delay between the price movement and the size of the food inventory that brings the system into oscillation, as in the cobweb model. However, this intuitive explanation only goes so far and does not enable us to predict, for example, the radius of the cycle or in fact whether a given system will be stable or unstable. One way to stabilize this market is to introduce value traders who estimate the fundamental price (Steiglitz et al., 1996), thus bringing a kind of foresight to market operations (see Fig. 5). As discussed above, the introduction of trend traders can produce price bubbles, as illustrated in Fig. 6.
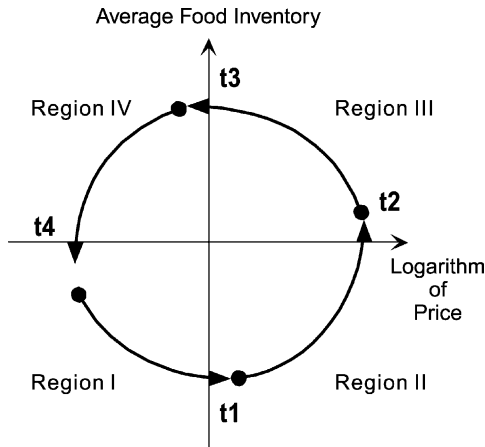
Fig. 4. Diagrammatic representation of price oscillations in an unstable market in the plane of food inventory vs. price.
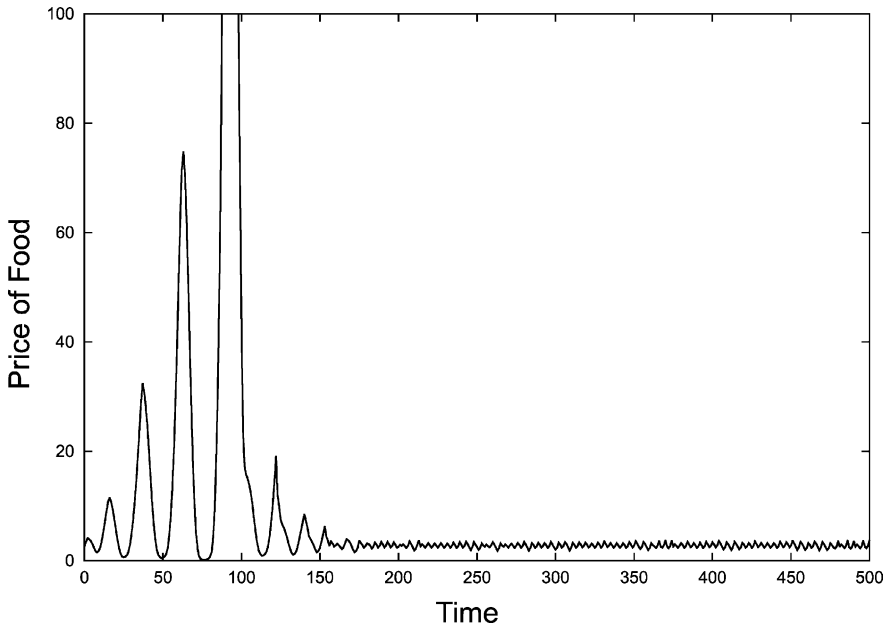

Fig. 5. Price vs. trading period with value traders, showing how speculators can stabilize the market. Value traders are introduced after 100 trading periods.
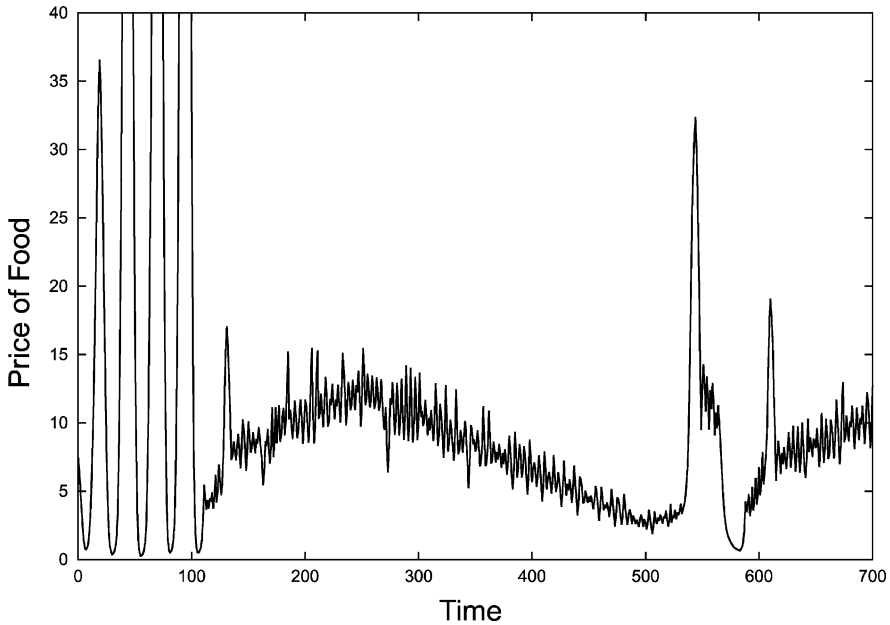
Fig. 6. Price bubble caused by the introduction of trend traders. The fundamental value is exogenously driven up and down to produce a trend. Value traders are introduced at period 100, after which the trading price remains close to the fundamental value until bubbles appear near trading periods 530 and 610.

Until now we have described simulations with previous models, which made available to the agents only the auction market-clearing price (ticker price) as a signal. This evidently does not communicate enough information to stabilize the market without some memory and foresight, which is invested in the value traders, who use an exponentially smoothed estimate of fundamental value. We next consider the possibility of using signals other than the market-clearing price to achieve stability.

3. Using other price signals

Consider again the market with only regular agents. After consuming one unit of food, each agent sends a bid p_i and a quantity a_i to be traded, both depending on the price signal as well as the difference between the agent's food inventory and his desired reserve. This bidding process generates at any given trading period an *order book*, comprising the agents' bid prices p_i and amounts a_i. This order book contains considerably more information about market conditions than simply the most recent closing price. This suggests that we can derive signals from the order book that can be more effective in stabilizing prices than the closing price. In practice it is this information that gives commodity traders in the pit an advantage over remote traders.
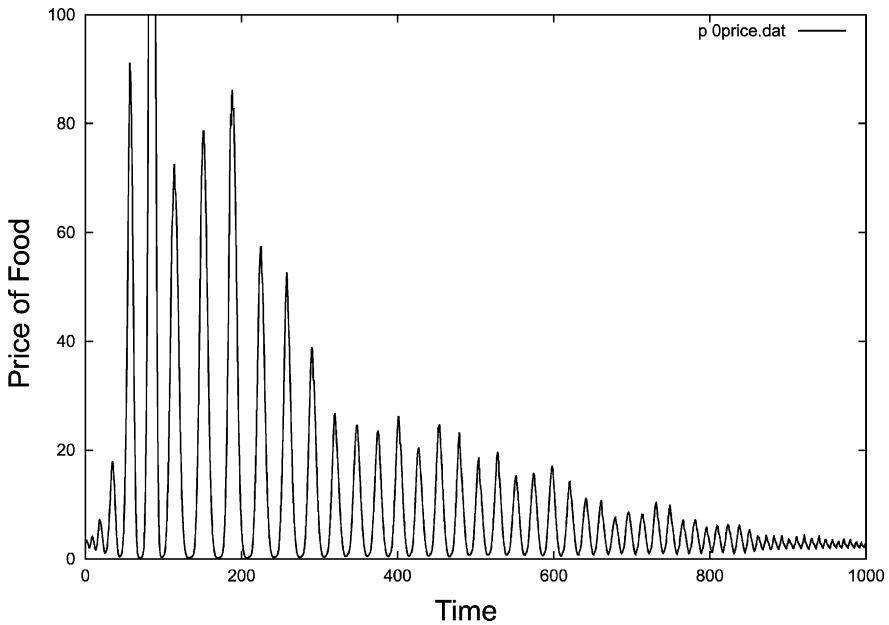
Fig. 7. Price vs. period with no traders, but using average bid, P_0, as the signal.

Consider first the simplest possibility: define the new signal P_0 to be the unweighted average of all the bid prices:

$$P_0 = \frac{1}{n} \sum_i p_i \tag{1}$$

Fig. 7 shows that the price is stabilized quite well, although the time to convergence is longer than with value traders.

Having observed the effectiveness of the mean bid as a signal, it is natural to try to improve it further, and a natural choice is the average of the bids weighted by the amounts P_1:

$$P_1 = \frac{1}{\Sigma a_i} \sum_i a_i p_i \tag{2}$$

Fig. 8 shows the result, which is perhaps surprising: weighting the bids by the amounts has the effect of destabilizing, rather than further stabilizing the market.

Finally, this suggests moving in the opposite direction: weighting the prices by some function that varies inversely with the corresponding amount. We therefore define P_2 to be

$$P_2 = \frac{1}{\Sigma 1/(c + a_i)} \sum_i \frac{1}{c + a_i} p_i \tag{3}$$

where c is a scaling parameter that determines the extent of inverse weighting. The value $c = 1$ was used in the simulations in this paper. Fig. 9 shows that the market with signal P_2 converges faster and better than with P_0.
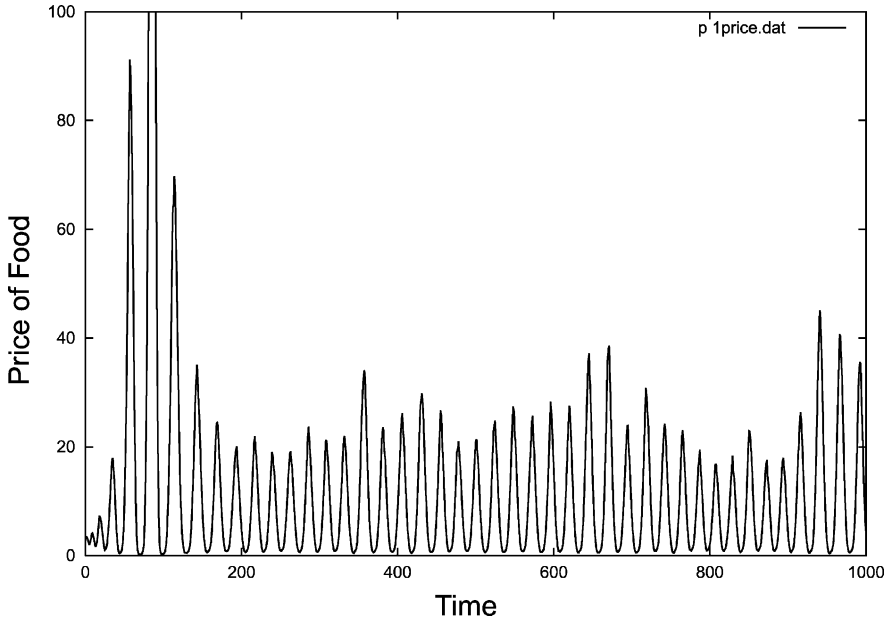
Fig. 8. Price vs. period with no traders, using the weighted average bid, P_1, as the signal.
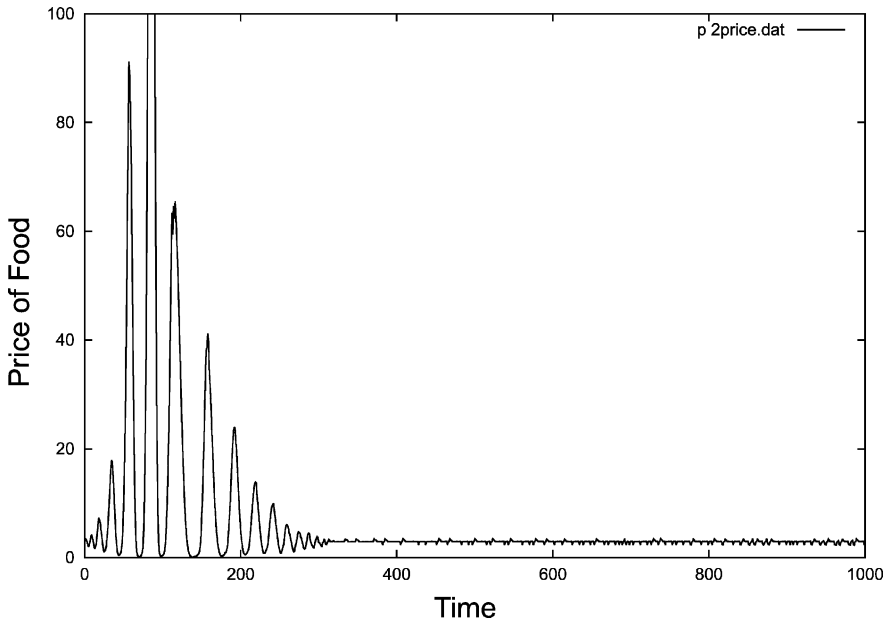

Fig. 9. Price vs. period with no traders, using the inversely weighted average bid price, P_2, as the signal.

The fact that weighting the bids with the amounts is destabilizing can be explained intuitively as follows: the agents bidding for large quantities are generally farther from their desired reserves, and their bids are therefore farther from equilibrium—farther above for buyers who have a severe deficiency, and farther below for sellers who have a severe surplus. Their bids are therefore more likely to be far away from the actual equilibrium than agents bidding for small quantities.

4. Concluding remarks

In this paper, we have considered the effects of different price signals on market stability using agent-based, microscopic simulations. Our models are practical for simulations of many hundreds of time steps, allow arbitrary, heterogeneous trading strategies and agent characteristics, and use a closed economy with a naturally defined equilibrium price that equates production and consumption. The simulations presented here were implemented using both Java and Java mobile agents *Aglets* (Lange and Oshima, 1998). Implementations in C run many times faster and make simulations for thousands of time steps practical if necessary.

Our results show that the average bid price signal P_0 stabilizes the market price effectively, and stable resource allocation is approached as well, all without predictive traders. What is perhaps counterintuitive is that supplying the agents with the weighted average P_1 neither increases stability nor improves resource allocation, but in fact achieves little improvement over using the closing price. Moreover, the inversely weighted average P_2 yields the greatest improvement in stability and resource allocation. It is noteworthy that this method for stabilization and control requires no increase in traffic, computation, or number of agents. These results suggest stabilization strategies for any applications that use agent-based technology, such as market-based distributed resource allocation or automated e-commerce on the Internet.

We have also investigated analytically the dynamics of a simplified model that concentrates on out-of-equilibrium price movements with small liquidity (see Appendix A). Figs. 11–13 show the reduction of the cycle after one period versus initial radii. The results are consistent with our simulations, showing that weighted, forward weighted, and inversely weighted price signals result in slowly converging, unstable, and strongly stable behavior, respectively. This analysis has so far been able to verify these results only for the two-agent case with idealized dynamics and price signals. Although these models are extremely simplified they still exhibit the complex behavior of the full system and retain its qualitative stability properties. More accurate analysis for more agents and more realistic versions of the weighted and inversely weighted price signals is left for future work. Finally, we can try to derive some insight into market mechanism from the qualitative results, one of the main motivations for agent-based simulation. It is somewhat surprising that an artificial market with no memory or foresight on the part of its agents can in some sense "learn" the equilibrium price and find an efficient equilibrium with only a slight amount of information beyond the most recent closing clearing price. This underscores the crucial importance of information flow in all markets, a fact well recognized of course by real-world traders. The fact that inversely weighting the bid information by quantity increases stability may be a

consequence of our particular market structure: those agents who bid to buy or sell small quantities are closer to their desired reserves, so their bids may reflect the true equilibrium price more accurately. Perhaps also there is a sense in which the new signal prices represent gradients when the market is viewed as an optimization problem, an idea related to the work of Ygge (1998) that deserves further study.

Appendix A

Exact theoretical analysis of the markets studied here is very difficult because of the highly nonlinear nature of the agent interactions mediated by the auction. Still, it would be helpful to verify to any extent possible the general results we have obtained by simulation. To this end we present in this section an analysis of a highly simplified model that despite its simplicity retains the essential properties of interest.

A.1. The simplified model

In this appendix, we consider the following dynamical equations for log-price-signal q and the centered food inventory variables a_i of two agents:

$$q(t + 1) - q(t) = -bE_w(a_i(t));$$
$$a_i(t + 1) = \begin{cases} S_i(t + 1) + a_i(t) & (\text{if } a_0(t)a_1(t) \geq 0) \\ S_i(t + 1) & (\text{if } a_0(t)a_1(t) < 0), \end{cases} \tag{A.1}$$

where the food production and consumption S_i is given by

$$S_i(t + 1) = \begin{cases} \sigma_i & (\text{if } q(t + 1) \geq 0) \\ -1 & (\text{if } q(t + 1) < 0), \end{cases} \tag{A.2}$$

and E_w denotes the weighted average corresponding to the particular definition of signal. For concreteness we assume the constant $b = \log(16)/17.5$, but the results do not depend on b because we can eliminate it by scaling q.

This simple dynamic model was obtained using a number of approximations and assumptions:

- We omit any dependence of the bidding function and auction on gold inventories.
- We consider only two regular agents in the market, and define the centered inventory variables a_i to be the difference between actual inventory and the desired reserve. Thus, $a_i > 0$ or $a_i < 0$ depending on whether agent i experiences an excess or shortage of food. In each trading period agent i can either produce σ_i units of food (his "skill") or consume one unit of food.
- Agents trade as follows: if $a_0 < 0$ and $a_1 > 0$, there is a trade and we set $a_0 = a_1 = 0$.
- Finally, we set $\sigma_0 < \sigma 1$, and choose the initial condition so that $a_0 \leq a_1$.

Using these simplifications, the use of $-bE_w(a_i(t))$ for the log-signal can be justified by simulations as a good approximation, using the exponential bidding function which Steiglitz et al. (1996) introduced, with dependence on gold suppressed. Despite the somewhat drastic
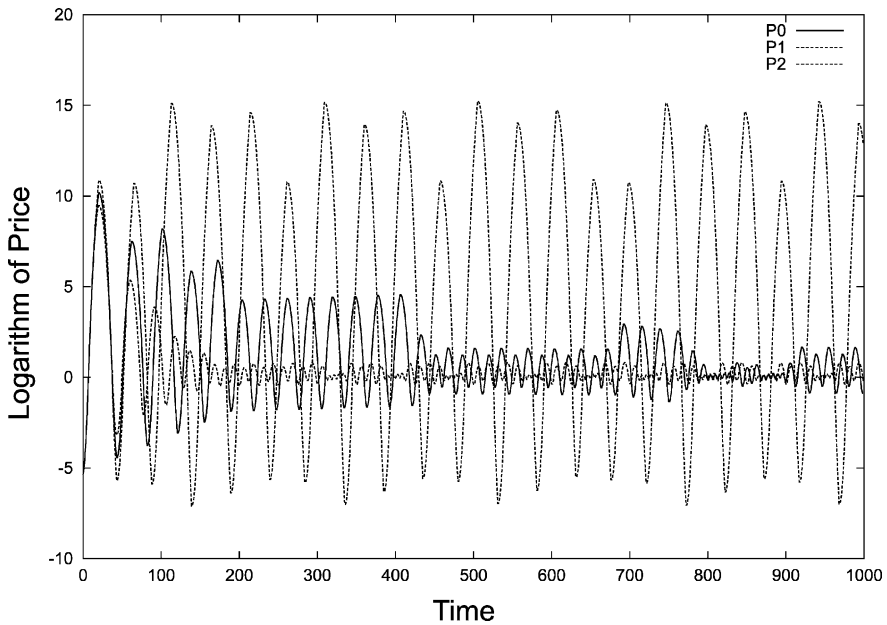
Fig. 10. Logarithm of price vs. trading period for signals P_0, P_1, and P_2 in simulations of the simplified model.

simplifications, this model exhibits behavior in simulations that is similar qualitatively to the full model: the price oscillates for P_0, diverges for P_1, and converges for P_2 (see Fig. 10). The goal of this analysis is to verify this analytically, and we discuss this in the subsections below for different price signals.

A.2. Simplified model with signal P_0

The price signal P_0 corresponds to the choice

$$E_w(a_i) = \tfrac{1}{2}(a_0 + a_1), \tag{A.3}$$

just the average of all bids. The initial conditions for time $t = 1$ are chosen as follows:

$$
\begin{aligned}
q(1) &= -Q < 0, \\
a_0(1) &= a_1(1) = -1.
\end{aligned}
\tag{A.4}
$$

We next calculate the price and the stock movement for the four periods in the cycle described in Fig. 4. The boundary of each period is given by $t = t_1$, t_2, t_3, and t_4. Let the average production $\hat{\sigma} = (\sigma_0 + \sigma_1)/2$ and the time interval $T_n \equiv t_n - t_{n-1}$. Our objective is to obtain the following state after one cycle:

$$
\begin{aligned}
q(t_4 + 1) &= -Q - \Delta Q < 0, \\
a_0(t_4 + 1) &= a_1(t_4 + 1) = -1.
\end{aligned}
\tag{A.5}
$$

After straightforward but lengthy calculations we obtain the following conditions for t_1, T_2, T_3, and T_4:

$$bt_1^2 + bt_1 - 2Q \geq 0,$$

$$T_2 > \frac{t_1}{\sigma_1},$$

$$T_3^2 + T_3 + \frac{2}{b\hat{\sigma}}\left[Q - \frac{t_1(t_1+1)b}{2} - \frac{(2bt_1 - b\hat{\sigma}(T_2+1))T_2}{2}\right] > 0,$$

$$T_4 > T_3\sigma_0.$$

(A.6)

Using the minimum positive integers for t_1, T_2, T_3, and T_4 satisfying these conditions, we obtain the difference ΔQ in the log-price after one cycle,

$$-\Delta Q = \frac{t_1(t_1+1)b}{2} + \frac{[2bt_1 - b\hat{\sigma}(T_2+1)]T_2}{2} - \frac{b\hat{\sigma}(T_3+1)T_3}{2}$$
$$- bT_3T_4\hat{\sigma} + \frac{bT_4(T_4+1)}{2}.$$

(A.7)

The radius of the cycle decreases if $-\Delta Q > 0$ and increases if $-\Delta Q < 0$. Fig. 11 shows the dependence of $-\Delta Q$ on the initial value of Q with $\sigma_0 = 0.3$, and $\sigma_1 = 0.6$. For large values of Q, the variance tends to decrease. For relatively small values of Q, the variance may decrease or increase. For very small Q, the variance tends to increase. Thus, we can say
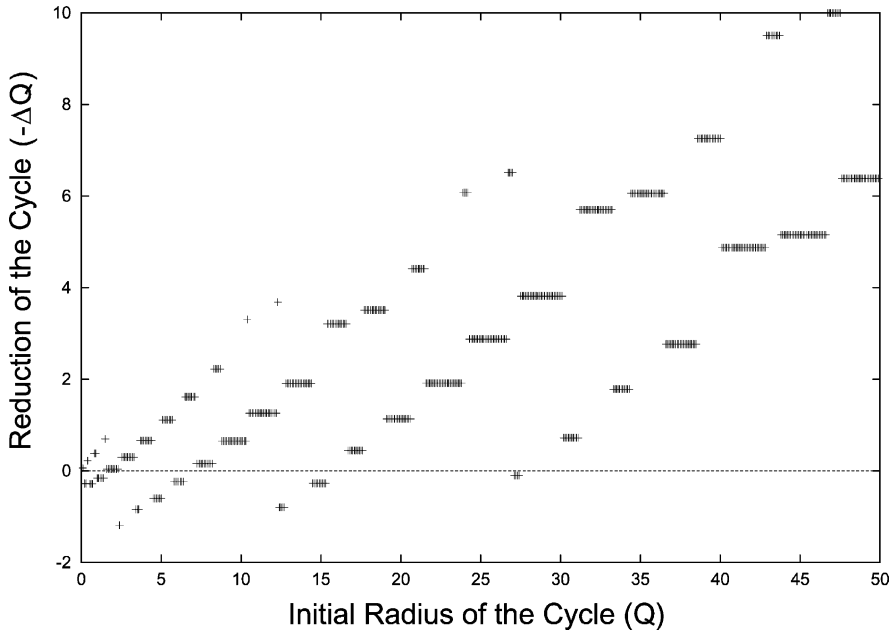

Fig. 11. Reduction of the cycle ($-\Delta Q$) vs. initial radius of the cycle (Q) using average bid, P_0, as the signal. If the initial radius is greater than 16, the cycle almost always shrinks. If the initial radius is less than 5, the cycle almost always expands.
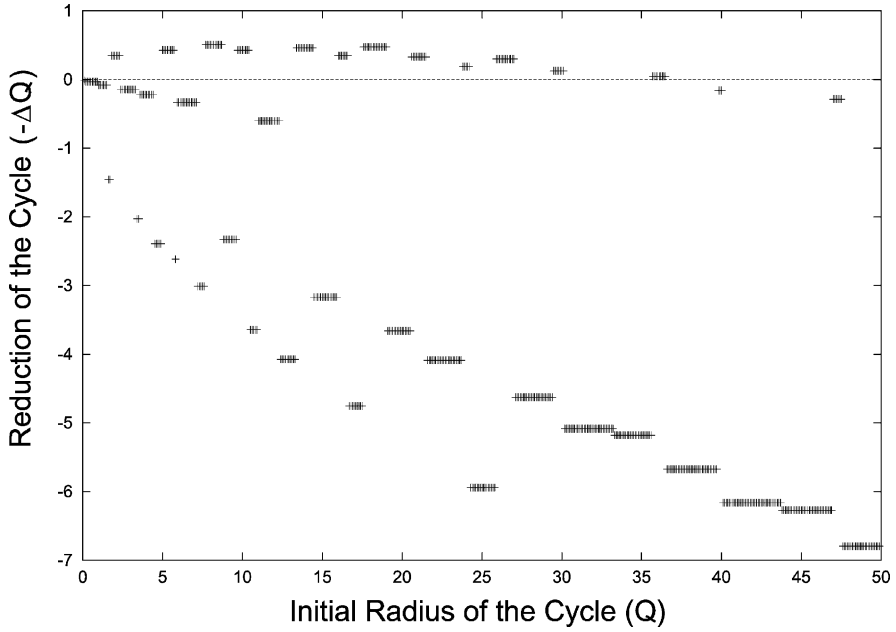
Fig. 12. Reduction of the cycle ($-\Delta Q$) vs. initial radius of the cycle (Q) using P_r as the signal. The predominantly negative values mean strong divergence.
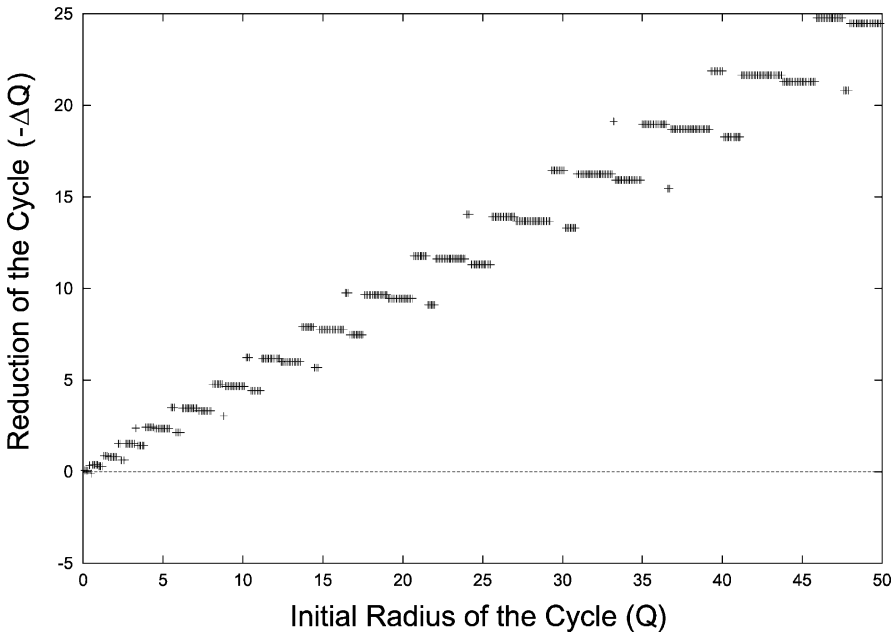

Fig. 13. Reduction of the cycle ($-\Delta Q$) vs. initial radius of the cycle (Q) using P_3 as the signal. The predominantly positive values mean strong convergence.

for the signal P_0 that the price oscillates to some extent but neither diverges nor converges. This coincides with the results obtained by simulating the simplified model.

A.3. Simplified model with other signals

We next consider ΔQ for two other simple signals: P_r, which corresponds to the closing market price, and P_3, a simpler version of the inversely weighted average P_2. In particular, for P_r the weighted average $E_w(a_i)$ is defined to be a_0 if $|a_0| > |a_1|$ and a_1 otherwise. For P_3, the weighted average $E_w(a_i)$ is defined to be a_1 if $|a_0| > |a_1|$ and a_0 otherwise. We also modify the dynamical equations for simplicity. If $a_0(t) < 0 < a_1(t)$, there is a trade: we set $a_i(t+1) = S_i$ and assume that there is no change in the price ($q(t+1) = q(t)$).

Simulation results shows that the price diverges with signal P_r and converges with signal P_3. We can study the system analytically for these signals P_r and P_3 in much the same way as we did for P_0. Fig. 12 shows the reduction values of $-\Delta Q$ versus the initial value Q. We see strong divergence from the dominating negative values. For P_3, the graph indicates strong convergence except when Q is very small (see Fig. 13).

References

Baumol, W.J., 1957. Speculation, profitability, and stability. Review of Economics and Statistics 39, 263–271.

Brock, W.A., Hsieh, D.A., LeBaron, B., 1991. Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence, MIT Press, Cambridge.

Caginalp, G., Balenovich, D., 1994. Market oscillations induced by the competition between value-based and trend-based investment strategies. Applied Mathematical Finance 1, 129–164.

Caginalp, G., Balenovich, D., 1996. Trend-based asset flow in technical analysis and securities marketing. Psychology and Marketing 13, 407–444.

Caginalp, G., Porter, D.P., Smith, V.L., 1998. Initial cash/asset ratio and asset prices: an experimental study. Proceedings of the National Academy of Sciences of the United States of America 95, 756–761.

Carlson, J.A., 1967. The stability of an experimental market with a supply–response lag. Southern Economic Journal 23, 305–321.

Clearwater, S. (Ed.), 1996. Market-Based Control: A Paradigm for Distributed Resource Allocation. World Scientific, Hong Kong.

de Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Positive feedback investment strategies and destabilizing rational speculation. Journal of Finance 45, 379–395.

de Oliveira, S.M., de Oliveira, P.M., Stauffer, D., 1999. Evolution, Money, War, and Computers: Non-Traditional Applications of Computational Statistical Physics. B.G. Teubner, Stuttgart.

Forsythe, R., Palfrey, T., Plott, C., 1982. Asset valuation in an experimental market. Econometrica 50, 537–567.

Friedman, M., 1953. Essays in Positive Economics. University of Chicago Press, Chicago.

Friedman, D., Rust, J. (Eds.), 1993. The double auction market. In: Proceedings of the Workshop on Double Auction Markets. Addison-Wesley, Santa Fe.

Kephart, J.O., Hanson, J.E., Levine, D.W., Grosof, B.N., Sairamesh, J., Segal, R.B., White, S.R., 1998. Dynamics of an information-filtering economy. In: Proceedings of Second International Workshop on Cooperative Information Agents (CIA-98).

Lange, D.B., Oshima, M., 1998. Programming and Deploying Java Mobile Agents with Aglets. Addison-Wesley, Massachusetts.

LeBaron, B., Arthur, W.B., Palmer, R., 1999. The time series properties of an artificial stock market. Journal of Economic Dynamics and Control 23, 1487–1516.

Levy, M., Levy, H., Solomon, S., 1994. A microscopic model of the stock market cycles, booms and crashes. Economics Letters 45, 103–111.

Mantegna, R.N., Stanley, H.E., 2000. An Introduction to Econophysics: Correlations and Complexity in Finance. Cambridge University Press, UK.

Porter, D.P., Smith, V.L., 1994. Stock market bubbles in the laboratory. Applied Mathematical Finance 1, 111–127.

Smith, V.L., 1989. Theory, experiment and economics. Journal of Economic Perspectives 3, 151–169.

Smith, V.L., Suchanek, G.L., Williams, A.W., 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. Econometrica 56, 1119–1151.

Stiglitz, J., 1990. Symposium on bubbles. Journal of Economic Perspectives 4, 13–101.

Steiglitz, K., O'Callighan, L.I., 1997. Microsimulation of markets and endogenous price bubbles. In: Proceedings of the Third International Conference on Computing in Economics and Finance, Stanford.

Steiglitz, K., Shapiro, D., 1998. Simulating the madness of crowds: price bubbles in an auction-mediated robot market. Computational Economics 12, 35–59.

Steiglitz, K., Honig, M.L., Cohen, L.M., 1996. A computational market model based on individual action. In: Clearwater, S. (Ed.), Market-Based Control: A Paradigm for Distributed Resource Allocation. World Scientific, Hong Kong, pp. 1–27.

Waldspurger, C.A., Hogg, T., Huberman, B.A., Kephart, J.O., Stornetta, W.S., 1992. Spawn: a distributed computational economy. IEEE Transactions on Software Engineering 18, 103–117.

Ygge, F., 1998. Market-oriented programming and its application to power load management. Ph.D. Thesis, Lund University.

Youssefmir, M., Huberman, B.A., Hogg, T., 1996. Bubbles and market crashes. Technical Report. Xerox Palo Alto Research Group, Palo Alto.