Miller
Risk
Advisors

Don’t Let Your Robots Grow Up To Be Traders: Artificial Intelligence, Human Intelligence, and Asset-Market Bubbles

by

Ross M. Miller

September 2002
Revised, April 2003

Working Paper 2002-02

2255 Algonquin Road   Niskayuna, NY 12309-4711 USA   millerrisk@earthlink.net

ABSTRACT

Don’t Let Your Robots Grow Up To Be Traders:
Artificial Intelligence, Human Intelligence, and Asset-Market Bubbles

by

Ross M. Miller
Miller Risk Advisors

Researchers who have examined markets populated by “robot traders” have claimed that the high level of allocative efficiency observed in experimental markets is driven largely by the “intelligence” implicit in the rules of the market. Furthermore, they view the ability of agents (artificial or human) to process information and make rational decisions as unnecessary for the efficient operation of markets. This paper presents a new series of market experiments that show that markets populated with standard robot traders are no longer efficient if time is a meaningful element, as it is in all asset markets. While simple two-season markets with human subjects reliably converge to an efficient equilibrium, markets with minimally intelligent robot traders fail to attain this equilibrium. Instead, these markets overshoot the equilibrium and then crash below it. In addition to firmly establishing the role of trader intelligence in asset-market equilibrium, these experiments also provide insights into why bubbles and crashes are consistently observed in many asset-market laboratory experiments using human subjects.

A particularly striking result from the early days of experimental economics was Vernon Smith's [1962] discovery that a simple auction mechanism patterned after the one used on the floor of the New York Stock Exchange leads to the efficient competitive equilibrium market allocation under a wide range of circumstances. The only major difference between the mechanism that Smith used, which he referred to as a double-oral auction, and the one used on the floor of the exchange was that the "specialist" (a role played by the experimenter) served only to maintain the order book and could not trade for his own account. In experiment after experiment spanning several years and using subjects ranging from high school students to seasoned commodities traders, virtually identical results have been obtained with the market quickly converging to the competitive equilibrium price and quantity (where the supply and demand curves induced by the experimenter cross). Furthermore, the resulting allocation of resources tended toward 100 percent efficiency as determined by the total surplus in the market, which could be measured by the amount of money that the subjects took away from the experiment. Furthermore, all this happened without the requirement of "perfect knowledge" that most neoclassical economists had come to assume was necessary for markets to operate efficiently as each subject in these experiments knew only his own conditions for supply and demand. This previously neglected feature of markets was the subject of a series of papers written by F.A. Hayek in the 1930s and 1940s and so Vernon Smith [1982] called this ability of markets to aggregate individual information efficiently the "Hayek Hypothesis."

As Vernon Smith and his colleagues gained more experience running experimental markets it became clear that the ability of the markets to converge rapidly to the equilibrium competitive equilibrium did not require that subjects involved in the experiment have an understanding of what they were doing. Indeed, it was common for one or more subjects in an experiment to behave in an apparently erratic manner without materially affecting the results of the experiment.

Inspired by a tournament of "robot traders" sponsored by the Santa Fe Institute (see John Rust, John Miller, and Richard Palmer [1992]) and patterned after Robert Axelrod's [1984] prisoners' dilemma tournament, Dhananjay Gode and Shyam Sunder [1993] examined what would happen in a double-oral auction experiment if all the roles of the human subjects were assumed by computer programs designed to act in the least sophisticated fashion possible. While markets consisting of such traders, known as "zero-information agents," were significantly more volatile than those with human subjects, Gode and Sunder found that they were nearly as efficient and generated average prices and quantities that approached the competitive equilibrium. These agents were programmed to generate bids and offers randomly selected from a uniform distribution function and constrained so that no trader could enter into a transaction that when taken by itself would generate a loss. Furthermore, random distribution of bids and offers never changed; hence, the agents could not learn from their past experiences in the market. Shyam Sunder, with Dhananjay Gode and other collaborators, later obtained similar results when zero-information agents could trade in a series of connected markets; however, these allocations were significantly less efficient than what was typical for comparable experiments conducted with human subjects. In summarizing the results of these experiments, Shyam Sunder [2002] writes in a working paper:

> When seen as human artifacts, a science of markets need not be built from the science of individual behavior. We outline how, in the recent decade, computer simulations enabled us to discover that allocative efficiency—a key characteristic market outcomes—is largely independent of variations in individual behavior under classical conditions.

This is a natural conclusion to reach if markets whose robot subjects have only a bare minimum of machine intelligence perform almost as well as the same markets with intelligent human subjects. Cliff and Bruten [1997] demonstrate that if the (arbitrary) random distribution used by Gode and Sunder does not converge to the competitive equilibrium for every parameterization of supply and demand and propose a remedy for this anomalous case that constrain agents to make "reasonable" bids based on past prices. The conclusion that robots can trade as well—an even, in certain circumstances, better than humans—can be refuted, however, if there are markets that require the intelligence of human subjects or highly sophisticated robotic agents in order to operate efficiently.

We will demonstrate that the element of time separates man from machine when it comes to market efficiency. That the Gode-Sunder results would fail to carry over to such markets is of great significance because time is critical to the operation of almost all naturally occurring markets. Furthermore, time is a defining characteristic of all assets markets, including Vernon Smith's model for his earliest experimental designs—the New York Stock Exchange.

In order to operate effectively in a market where the cost or value of an item is time-dependent, an agent must be capable of planning. Indeed, the very notion of intelligence can be is often equivalent to the ability to create and execute plans. While economists in the field of experimental economics have tended to focus on F.A. Hayek's influential article, "The Use of Knowledge in Society," [Hayek, 1945] which expounds on the ability of markets to aggregate private individual knowledge, Hayek in an earlier work, "Economics and Knowledge," [Hayek, 1937] examines how time and planning figure into the operation of markets. Hayek makes his case for the virtue of markets based on the observation that the decisions that stem from central planning are inferior to those derived from the aggregation of individual plans. Hayek explicitly views time as a central element of the market's operation, and states early on:

> [S]ince equilibrium is a relationship between actions, and since the actions of one person must necessarily take place successively in time, it is obvious that the passage of time is essential to give the concept of equilibrium any meaning. This deserves mention, since many economists appear to have been unable to find a place for time in equilibrium analysis and consequently have suggested that equilibrium must be conceived as timeless. This seems to me to be a meaningless statement.

This ability to take multiple actions that occur successively in time requires intelligence that far exceeds that of a zero-information agent and in sufficiently complex settings may exceed that of any existing machine intelligence. Indeed, a detailed statistical examination of past market activity may be insufficient for the efficient operation of market—as Van Boening and Wilcox [1996] observe markets may require that agents in the economy possess “forward-looking attention.”

Experimental economists began to conduct experiments in which time played a meaningful role during the middle of the 1970s. The first set of such experiments, conducted by Ross Miller, Charles Plott, and Vernon Smith [1977], adapted Vernon Smith’s basic market design to incorporate two “seasons” where demand increased from the first season to the second. Buyers could “plan ahead” by purchasing units for “consumption” in the second season during the first season. The intertemporal competitive equilibrium for this market was characterized by a uniform price across seasons (there were no carrying costs) and required some, but not all, subjects to carry over units from the first season to the second. These markets converged to the competitive equilibrium (and near-perfect allocative efficiency) almost as quickly as the static single-season markets. This result was enough of a surprise to experimentalists that Arlington Williams [1980] and others conducted the experiment using other subject pools—the original experiments used Caltech undergraduates— and confirmed that these markets reliably converged to the intertemporal competitive equilibrium.

This article examines how markets that consist of zero-information agents operating under the protocols established by Sunder and his collaborators perform in the Miller-Plott-Smith setting. As one might expect, without the intelligence necessary to follow even the most rudimentary plan, the zero-information agents cannot use price signals to decide on the subset of agents that will carry units over to the second period. Instead, every agent who can carry units over to the second season will do so. As a result, the market fails to converge to the intertemporal competitive equilibrium and exhibits substantial allocative inefficiencies. Furthermore, because the agents are incapable of learning from their mistakes, this behavior is repeated indefinitely.

A surprising result is that what one might consider to be a market bubble with prices remaining stubbornly above the expected competitive price develops in the first season.

This apparent bubble is followed by a crash when new supply comes on line in the second season. Furthermore, reparameterization of the Miller-Plott-Smith markets to make them more consistent with previous experiments on zero-information agents serves to accentuate both the bubble and the crash. These results are of particular interest because when Vernon Smith, Gerald Suchanek, and Arlington Williams [1988] extended the two-season design to an asset market with as many as fifteen seasons, they consistently found that inexperienced subjects generated bubbles while sufficiently experienced subjects traded near competitive equilibrium prices. To the extent that naïve human experimental subjects faced with the novelty and complexity of a laboratory experience might behave in some measure like zero-information agents, the ease of creating bubbles with these simple robots provide an elegant model of how market bubbles can develop.

We begin with an overview of the experimental design and results of the Miller-Plott-Smith experiments and shows how they can are easily translated into the Gode-Sunder framework for trading with zero-information agents. The results of experiments run with these agents, along with a number of variants and extensions of the experiment are described. Finally, these results are compared with the results of experiments on human subjects in which bubbles and crashes are consistently observed.

Redesigning an intertemporal market for robot traders

The Miller-Plott-Smith experiments divided subjects into buyers and sellers and required the subjects to keep the same role throughout an experiment. In addition, one experimental design included “traders” whose function was to buy in the first season (known in these experiments as the “Blue Season”) for sale in the second or “Yellow” season. Traders could also buy and sell units within a period for purely speculative purposes. Because of the difficulty of designing zero-information agents who can speculate, an issue that we will address later on, “traders” were not employed in the computerized market described here.

	Unit Redemption Values			
	Blue		Yellow	
Buyer 1	2.00	1.10	2.45	1.55
Buyer 2	1.85	0.95	2.60	1.70
Buyer 3	1.70	0.80	2.75	1.85
Buyer 4	1.55	0.65	2.90	2.00
Buyer 5	1.40	0.50	3.05	2.15
Buyer 6	1.25	0.35	3.20	2.30

	Unit Costs	
	Both Seasons	
Seller 1	0.80	2.45
Seller 2	0.95	2.30
Seller 3	1.10	2.15
Seller 4	1.25	2.00
Seller 5	1.40	1.85
Seller 6	1.55	1.70

Table 1. Dollar unit redemption values and unit costs in Miller-Plott-Smith experiment with 6 buyers and 6 sellers.

Markets were, however, able to equalize prices between the seasons because buyers could purchase unit for redemption in the Yellow Season during the Blue Season. If prices tended to be higher in one season, a shift in the total number of units carried over by buyers is all that is necessary to bring the two seasons back into equilibrium.

The supply and demand parameters used in these experiments are given in Table 1. In order to economize on the number of subjects used, each buyer and seller was allotted two units for each season. For example, the first buyer has units with "redemption values" of \$2.00 and \$1.10 for the Blue Season and \$2.45 and \$1.55 in the Yellow Season. During the Blue Season, this buyer can make purchases for either season, with the restriction that the higher-valued unit in a season must be purchased first. Hence, if the Buyer purchases a unit for \$1.40 in the Blue Season, he can either redeem his first Blue Season unit for \$2.00 (yielding a profit of \$0.60) or his first Yellow Season unit for \$2.45 (yielding a profit of \$1.05). Only after the first unit in a season has been redeemed does the second unit's redemption value come into play. Sellers are given the same cost parameters in each season and can only sell their blue units in the Blue Season and their yellow units in the Yellow Season. Sellers are also required to sell their most profitable units first in each season before they can sell the second unit. Under no circumstances can any buyer pay more than unit's redemption value nor can any seller receive less than its cost. Also, buyers can

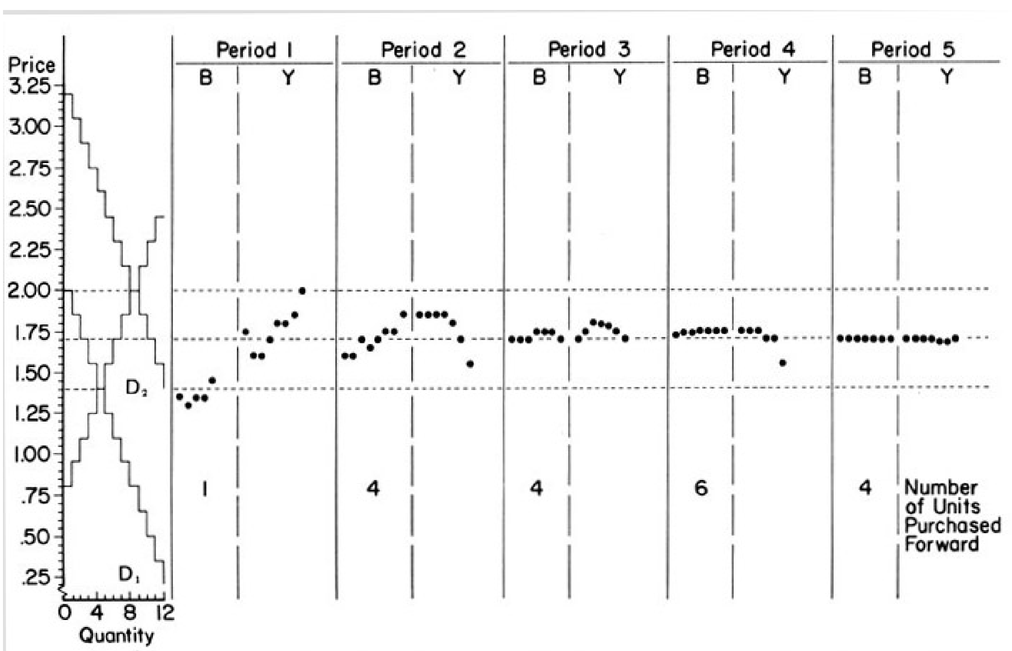

Figure 1. Results with human subjects from Miller, Plott, and Smith [1977]

only make bids that improve on the outstanding bid and sellers can only make offers that improve on the outstanding offer. After each unit is traded, the order book is reset; automatically canceling any untaken bids or offers. Finally, each buyer and seller receives a commission of $0.05 for every unit traded to provide an incentive to trade the marginal unit. All subjects in the Miller-Plott-Smith experiments received their earnings in cash at the end of each experiment. These earnings were calibrated to exceed the hourly rate for comparable campus jobs regardless of the experiment's outcome.

Figure 1 shows the equilibrium allocations and all the trades for a five-period experiment. Considered in isolation (the autarky case), the competitive equilibrium in the Blue Season is 5 units trading at a price of $1.40 and in the Yellow Season is 9 units trading at a price of $2.00. Combining the two markets by allowing buyers to purchase units in the Blue Season and carry them over to the Yellow Season, the price in the Blue Season rises and the price in the Yellow Season falls until they reach equilibrium at $1.70. In this equilibrium, 7 units are sold in both season with 4 units are carried over from the Blue Season to the Yellow Season.

Human subjects tend to converge to equilibrium in this market design somewhat more slowly than in the traditional double-oral auction markets in which periods are not divided into seasons. In Period 1, only a single unit is carried over and so the price rises in the Yellow Season. Apparently noticing this increase in price, buyers carry over more units in the next period and then learn in the Period 4 the consequences of carrying over too many units—a crash at the end of the period caused by an excess of supply. Efficiency, as measured by the ratio of profits captured by the subjects relative to the maximum possible surplus that they could extract from the market, runs from 95 to 100 percent in each period except the first, when it is 92.6 percent.

Transferring this experimental design to robot traders in a manner consistent with the Gode-Sunder zero-information markets is straightforward. The only complication is that human buyers are given some discretion as to the order in they can purchase units during the Blue Period. For example, in the example given above, as long as a unit is being offered at a price below $2.00, the buyer can either buy a Blue Period unit or a Yellow Period unit first. Because zero-information agents lack the intelligence to exercise discretion, they were programmed to purchase their highest value units first. Hence, in this example, the robot trader would always buy the Yellow Period unit before purchasing a Blue Period unit during the Blue Period.

As in the Gode-Sunder robot markets, buyers and sellers randomly generate their bids and offers. Using a random-number generator, the computer first chooses a side of the market (buyer or seller) with equal probability and then picks a specific buyer or seller at random, also with equal probability. A bid price or offer price is then chosen from a uniform distribution from the whole-cent prices along the interval [$0, $4]. The bid or offer is immediately discarded if the agent chosen has used up its allotment of units or if all its units cannot be traded without incurring a loss at the randomly chosen price. (Breakeven trades are allowed.) If a buyer's bid price equals or is greater than the going offer price or if the seller's offer price is equal to or less than the going bid price, a transaction is consummated at the midpoint of these two prices. (With human subjects, buyers or sellers would simply "accept" an outstanding bid or offer.) If a trade is not possible, then if the buyer's bid is higher than the outstanding bid or the seller's offer is lower than the outstanding offer, that order is placed on the books. As in the Miller-Plott-Smith experiments and previous robot-trader experiments, the order book is wiped clean after each trade.

The robot market is organized in the same manner as the Miller-Plott-Smith human experiments except that we follow the Gode-Sunder rule of not allowing buyers or sellers to accept outstanding bids or offers, but rather use the midpoint rule to "cross" trades. (For example, if the outstanding offer is $2.10 and a bid is placed at $2.27, a trade is consummated at $2.185.) The other difference from the Miller-Plott-Smith experiments is that the five-cent commission for each trade is not paid to the robots because it is unnecessary. The robots are programmed to seek out all trades and so the commission would have no effect on the results because the robots are willing to enter into trades that earn them nothing. All results presented here were programmed in *Mathematica* and use its built-in random-number generator to choose among buyers and sellers and to generate their bids and offers.

Robots fail to replicate human results

Typical results from running the zero-information robot market for a single two-season period are presented on the left side of Figure 2. This trial, like all the other trials reported in this paper, was run using a total of 1,000 randomly generated bids and offers for each season. (This number was chosen to be sufficiently large that all possible trades would almost certainly be executed long before the season ended.) Notice that in the Blue Period every trade is executed at a price above the intertemporal competitive equilibrium price of $1.70 and many units trade above $2.00. In the Yellow Season prices crash to below the intertemporal competitive equilibrium and remain there. The total surplus extracted by the agents in this example is $13.05 or 87 percent of the maximum possible surplus of $15.00, which is the surplus available at equilibrium. By comparison, double-auction markets with human subjects rarely drop below 90 percent efficiency in a single period and typically average at least 95 percent over the course of an experiment.

As with single-season market, the prices viewed over a single market period tend to be volatile and the failure of the zero-information agents to learn implies that the volatility is not reduced in future periods. Some of the noise from the results, however, can be removed by averaging many periods (2,000 is more than sufficient) together. The average allocative efficiency (surplus relative to the maximum possible surplus) over 2,000 runs is 87.727 percent, slightly more than in the sample period shown in Figure 2.

Generating an average time series for prices is complicated by the fact that the quantity traded in each season varies from period to period. To ensure comparability across runs, the 720 (out of 2,000) runs with 9 trades in the Blue Season followed by 3 trades in the Yellow Season are averaged together to provide a smoother picture of the dynamics of a typical period as shown in right side of Figure 2. Trades in the Blue Season tended to be at prices slightly above $2.00 and there was only a slight upward trend in those prices within

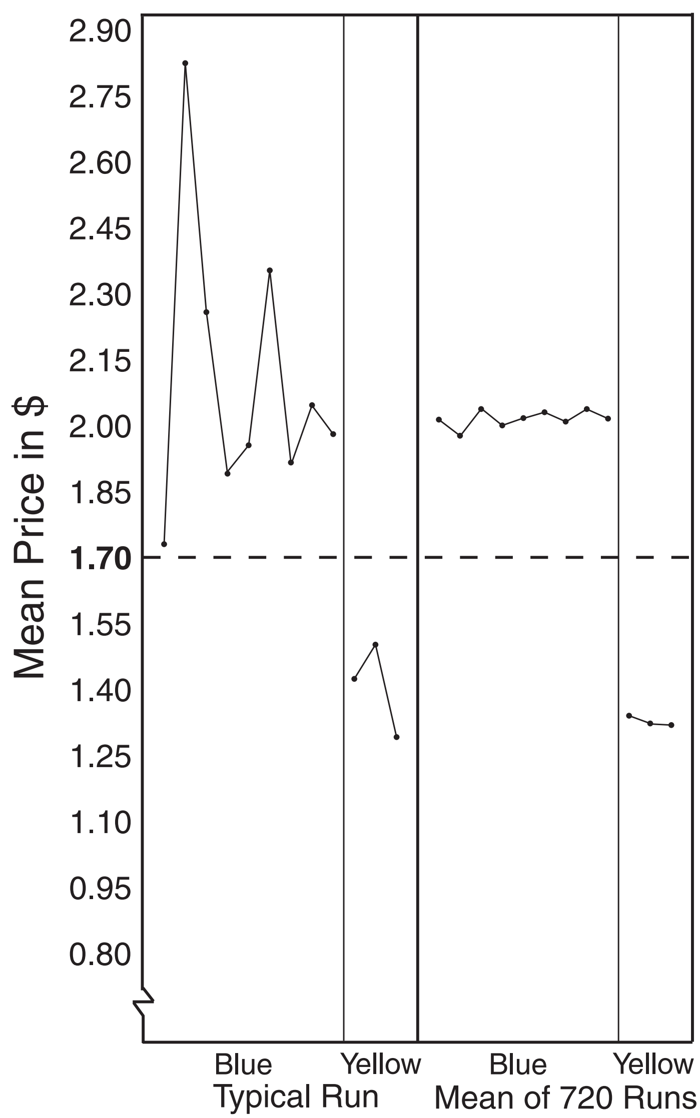

Figure 2. Results with zero-information agents in the Miller-Plott-Smith esxperiment.

the season. Prices then plunged to an average of \$1.34 for the first trade of the Yellow Season and continued down so that the third and final trade averaged less than \$1.32. Qualitatively similar results are obtained when averaging prices for periods with other quantities of units traded in the two seasons.

It should also be noted that changing the random-number distributions has only a minimal effect on prices and efficiency. Picking from prices uniformly distributed from \$0 to \$3.21 (just above the highest redemption value for a buyer), somewhat reduces prices in the Blue Season to an overall average of \$2.01 and increases them by less than a cent in the Yellow Season. It is likely than under any reasonable distribution of random bids and offers the basic result that prices remain well above equilibrium in the Blue Season and then crash below it in the Yellow Season will continue to obtain.

Modifications and extensions of the robot market

The results described above clearly demonstrate that the zero-information agents achieve an allocation that is both different and less efficient than competitive theory suggests and that human subjects are able to achieve under identical conditions. The "eagerness" of zero-information agents to trade appears not only to undermine the efficiency of the market (unlike in static settings where it promotes efficiency) but leads to a rudimentary market bubble and crash.

These results become even more dramatic when the market reorganized so that it is more like previous experiments with zero-information agents and less like the experiment with human subjects. This requires two changes. First, demand is revised so that rather than increasing from the Blue Season to the Yellow Season, the average demand over the two seasons is used for each season. In this way, we can eliminate any overbidding in the Blue Season that might be attributed to the projected increase in demand. Second, the number of agents is doubled so that each agent has a single unit to trade in each season. This prevents one of an agent's units from "casting a shadow" over the other unit and eliminates a choice for the agents. Now buyers are simply programmed with the obvious priority, they must buy a unit for redemption in the Blue Season before they can purchase one for redemption in the Yellow Season.

The supply and demand curves are given on the left side of Figure 3. Before these experiments were run, a set of single-season control experiments were run to assure that the results from these robot markets are similar to the single-season Gode-Sunder markets. Because the supply-and-demand configuration used in the Miller-Plott-Smith experiments is more widely dispersed and has a greater proportion of extramarginal units, the single-period market based on these parameters generates a greater dispersion of prices and a somewhat less efficient market than is typical for the Gode-Sunder markets. Averaging over 2,000 single-season trials, these markets are 95.61 percent efficient, comparable to the

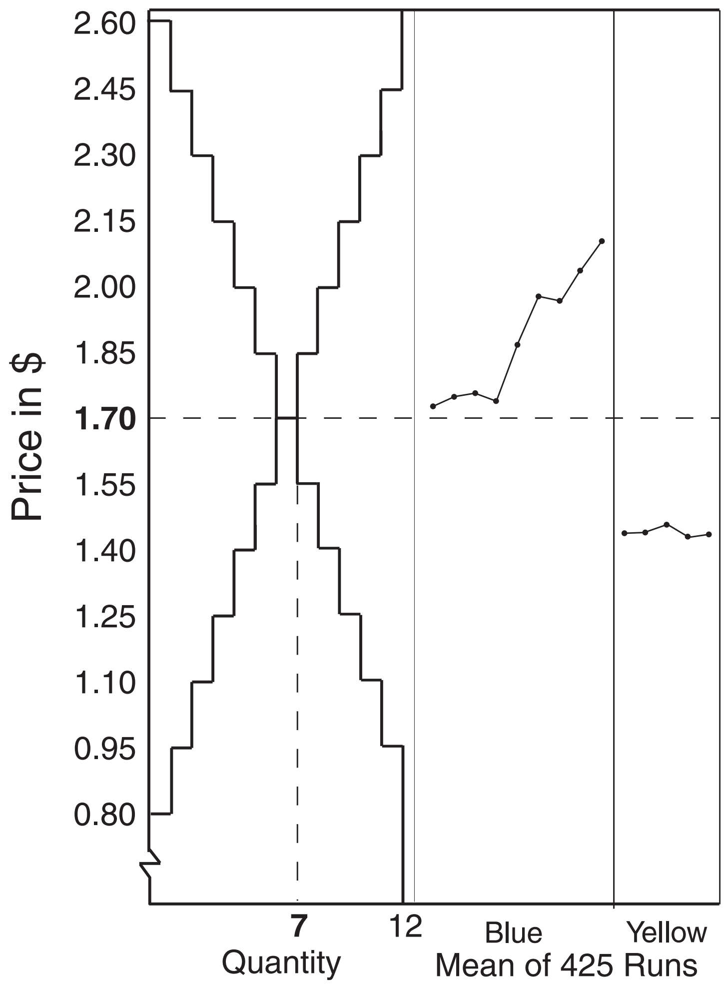

Figure 3. Supply and demand reparameterized for constant demand across seasons along with results.

low end of the results for human subjects in similar markets, and generate a mean price of \$1.7072, slightly above the competitive equilibrium of \$1.70. In addition, while prices can appear to be trending up or down within an individual trial season, on average prices exhibit no consistent trend.

The results for two-season markets are similar to those that were obtained by placing zero-information agents in a market with the Miller-Plott-Smith parameters. For the 2,000 independent two-season runs, average allocation efficiency increases to 88.83 percent. This is still below the efficiency that is achieved with human subjects. Even without an increase in demand in the second season, prices still are well above the competitive equilibrium price of \$1.70 in the Blue Season and well below it in the Yellow Season. On the right side of Figure 3, we see the results averaged over the 425 trials where 9 units are traded in the Blue Season and 5 are traded in the Yellow Season. Now with all units exposed to the market for the beginning of each season, there is a pronounced tendency for prices to rise within the Blue Season before crashing at the start of the Yellow Season. Furthermore, once prices have crashed, they tend to be nearly constant (on average) for the rest of the Yellow Season. In general, the more units that are traded in the Blue Season, the higher prices rise before plunging at the beginning of the Yellow Season.

Extending the number of seasons from two to an arbitrarily large number changes the path that prices take. Figure 4 shows the period-by-period average prices for a typical

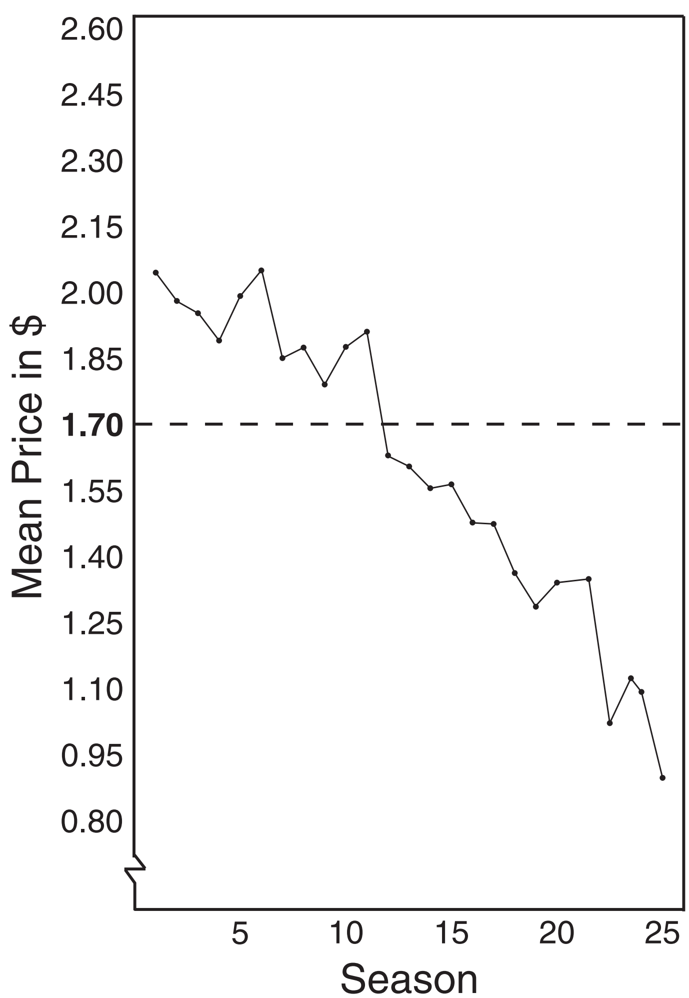

Figure 4. Mean prices in each season of a 25-season experiment.

25-season experiment with zero-information traders. Prices start at \$2.04 in Period 1, which is well above equilibrium, and then ratchet down several times during the remainder of the experiment, including a mini-crash from \$1.88 in Period 11 to \$1.63 in Period 12, until they are well below the equilibrium price at the end, with an average price of \$0.895 in the Period 25. The discontinuous nature of the decline in prices comes from the \$0.15 increments in demand—buyers tend to buy units with higher redemption values first and so as these units disappear from market over a series of seasons, prices must drop by at least \$0.15 to bring new buyers into the market. Also, within a period there is a tendency for prices to plunge at the end of a period when high-cost sellers sell their units before low-cost sellers. The path that prices takes resembles a typical "bear market" that breaks the decline in prices down into occasional "panics" that are often followed by brief recoveries more than it resembles a precipitous market crash.

Finally, the 25-season market is significantly less efficient that the two-season market. Averaged over 100 independent runs the market yields only 83.24 percent of the possible surplus. Even the most efficient market of these runs is just 85.33 percent efficient. Although equivalent 25-season experiments have not been run on human subjects (and are not likely to be given the time and expense), this result provides further evidence that the rules of the market alone without the aid of human intelligence are not sufficient to make markets efficient.

Robot bubbles vs. human bubbles

When compared with similar experiments with human subjects, the results of this series of experiments on zero-information robot traders demonstrate that some aspect of human intelligence appears necessary for even the simplest asset market to converge to an inter-temporal competitive equilibrium. Without the ability to plan or respond to the information contained in price signals, the robot traders "behave" as if there were no tomorrow, driving prices up only to have them plunge or suffer a protracted decline depending on the number of seasons.

While it is of considerable interest that humans can so handily outperform simple robot traders in an easily replicable experiment, the failings of the robot traders may shed light on the results in other asset-market experiments involving human subjects. Beginning with the asset-market bubble experiments of Vernon Smith, Gerald Suchanek, and Arlington Williams [1988], investigators had observed the ease with which market bubbles can be generated under a variety of experimental conditions. In these experiments, the value of an item is determined not by presenting subjects with payoffs that induce the desired supply and demand schedules, but rather by endowing them with "assets" that pay dividends and "money" that can be used to purchase assets. Trading of the asset occurs because of differences in risk tolerance or for purely speculative reasons.

Most of these asset-market experiments exhibit similar patterns of prices. The price of the asset begins trading below its “intrinsic value” as determined by the expected value of its future stream of dividend payments. Over the course of the experiment, the price of the asset tends to rise at the same time its intrinsic value is falling as dividend payments are being made. The price of the asset not only moves above its intrinsic value, it tends to move explosively higher. At some point, as the pre-announced end of the experiment draws near, the price of the asset will crash, usually dropping back below its intrinsic value. A typical experiment runs for 15 “periods” (the equivalent of “seasons” in the experiments described earlier) and usually takes at least 2 hours to conduct. Additional iterations of the experiment require that subjects return at another time, unlike the Miller-Plott-Smith experiments where five iterations could be conducted in a single sitting. Pools of subjects who have experienced one asset-market bubbles are more likely to avoid one in a follow-up experiments and those who have experienced two bubbles rarely succumb to a third. In addition, subjects who have participated in other types of market experiments are far less likely to create a bubble than completely inexperienced subjects.

More recent bubble experiments by Gunduz Caginalp, David Porter, and Vernon Smith [2000] as well as mathematical models of speculative behavior developed by Caginalp and Balenovich [1999] have shown how “momentum trading” might be an important factor in bubble formation. Because prices in these asset markets tend to converge from below, traders may become accustom to increasing prices and continue to trade on this momentum, pushing prices past the competitive equilibrium in the process. The experiments on zero-information agents, however, demonstrate that while momentum trading may play a role in naturally occurring market bubbles, it is not necessary in order for bubbles to form. Indeed, given their utter simplicity, zero-information agents never take price momentum, or any other aspect of the price history, into account in placing their bids and offers.

While momentum trading may very well contribute to bubble formation both in the laboratory and in naturally occurring asset markets—for example, investment bankers have an incentive to underprice their initial public offerings of companies not only to provide perks to their most-valued clients but also to create positive price momentum—the experiments with zero-information agents provide an alternative, and possibly more elegant, economical explanation of bubble formation. The bubble in Internet and other technology stocks that formed at the end of the 1990s could have been partially rooted in investors’ inability to properly anticipate the future supply of stock in Internet companies. Shares in the first Internet companies brought to market soared in part because investor demand outstripped supply. But as one would expect, Silicon Valley and Wall Street quickly geared up to meet this demand, flooding the market with new shares on a daily basis at the height of the boom. Just like the zero-intelligence robots, many investors acted without regard to the effects of future supply, especially the abundance of supply that high share prices in technology companies would stimulate.

A key difference between the bubble experiments described above and the experiments on zero-information agents is while it is easy to assign a human the role of speculator whose goal is to buy low and sell high, it is not a simple to program an agent who can speculate effectively. Despite repeated efforts by artificial intelligence research to create robot traders that can compete with humans in financial markets, it is widely believed that the intelligence required of a successful speculator vastly exceeds what is possible with the current technologies used in the field of artificial intelligence. To the extent that computer-based trading systems have been successful, that success has often been short-lived. The problem is that the appearance of profitable trading opportunities is quickly detected by competing speculators and so successful speculation requires continuous adaptation. Given that even simple learning is difficult for artificial agents in a market setting, the amount of learning required for successful speculation for more than a limited time becomes truly daunting.

The possibility that markets can form bubbles even without subjects or robots assuming the role of speculators is not without precedent in the experimental literature. Vivian Lei, Charles Noussair, and Charles Plott [2001] have produced bubbles in markets with human subjects none of whom have the ability to speculate in the market, that is, they cannot purchase items simply for future resale in expectation of making a profit. They attribute these bubbles to what they call the *Active Participation Hypothesis*—the belief that human subjects will trade in experiments even when there are reasonable expectations that sitting and doing nothing is a more profitable course of action. Hence, the eagerness of zero-information robots to trade parallels a hypothesized human behavior.

This leads to the reasonable conjecture that bubbles are more likely to form in markets where the behavior of subjects is most similar to that of zero-information agents. Not only does this conjecture explain bubble formation in nonspeculative markets, but it also explains why bubbles disappear as human subjects gain experience in markets. While it is inconceivable that any human subject follows the exact random-bidding strategy employed by zero-information agents, the actions of a subject who is overwhelmed by the novelty of the experimental environment could well mimic random behavior. Subjects uncertain as to how to plan a course of action in a novel and complex environment might simply act immediately without careful consideration of the future consequences. It is important to note that human subjects only fail to plan and coordinate their actions through the market mechanism in complex experimental settings. In less complex experiments with only a few seasons and predetermined supply and demand, human subjects converge on the competitive equilibrium while zero-information agents do not.

It should be noted that the bubble experiments that have been conducted on human subjects differ in a fundamental way from the standard double-oral auction supply-and-demand experiments. In bubble experiments allocative efficiency never enters the picture explicitly. In experiments where dividends are uncertain, an efficient allocation involves transferring assets from less risk-averse to more risk-averse subjects; however, nothing in the experimental set-up provides us with the information (a measure of each trader's risk aversion) necessary to determine the allocative efficiency of the market. When dividend payments are known with certainty at the beginning of the experiment, as in the experiments conducted by David Porter and Vernon Smith [1995], every allocation is equally efficient. In both cases, bubbles matter only to the extent that they transfer income from one subject to another.

In contrast, the bubbles that form in the experiments conducted on zero-information agents described above all have a marked impact on allocative efficiency. The size and nature of the bubbles that are formed, however, is quite different because the zero-information agents are prohibited from engaging in actions that could result in a loss and so are not allowed to speculate. Because of this restriction, the bubbles and crashes seen with the robot traders are much less dramatic than in the experiments conducted with human subjects. The highest redemption value for the buyer with the greatest demand serves as a ceiling for the market price and as those units are exhausted the ceiling drops.

Concluding Remarks

The point of this paper was to provide a rigorous demonstration that contrary to the previous literature on zero-information robot traders, there are simple circumstances in which intelligent agents are necessary if markets are to reach their efficient competitive equilibrium. This does not mean that the institutions that govern markets do not play a role in market performance; indeed, a simple change in market institutions can lead to an enormous boost in market performance. Such changes are only matter, however, in markets where reaching equilibrium depends on the ability of agents to learn. Miller [2002a] uses the example of A. Michael Spence's celebrated signaling model as an example of a market that depends on the ability of agents to learn for its efficiency and where a minor modification of market rules can lead to vast improvements not only in market efficiency, but also in market stability.

With the difference in market performance between humans subjects and zero-information agents firmly established, the natural next step is the determination of how much intelligence a robot traders needs in order to provide a convincing simulation of human behavior in a market setting. The design of a special-purpose agent that can trade in the simple asset markets examined in this paper as well as, if not better than, humans seems clearly within grasp. Straightforward heuristics can enable the agents to "shop around" and "plan ahead."

In a predetermined setting, zero intelligence is insufficient, but substantial intelligence is unnecessary. Indeed, the popular press has reported extensively on results by Das et. al. [2001] have show that of an enhanced version of zero-information agents over untrained human traders in double-auction experiments.

The real problem lies in the determining how much intelligence an agent requires to match human performance in a dynamic setting with constantly changing market conditions. Even more desirable, given the human propensity to be drawn into market bubbles, would be the creation of robots that can avoid these traps. Unfortunately, given the experience of market professionals who hoped to harness machine intelligence for their own enrichment, it may be a long time before robot traders can exhibit the ability to learn and adapt necessary for them to earn their keep in the marketplace.

The results of the work described in this paper are particularly interesting in light of the work of Philip Mirowski [2002], which has expressed skepticism both toward the Gode-Sunder work and toward the research program of experimental economics. Mirowski views experimental research as flawed because its very design forces the human subject to behave robotically, especially in light of research that appeared to indicate that market behavior was largely independent of whether agents were computerized robots or human subjects participated in an automated laboratory experiment. While there is no question that the laboratory environment deprives any market of some of its richness, not only do the important salient features of the market remain, but human subjects are still able to exhibit their humanity—in this instance, by using the auction mechanism to send and receive the price signals necessary to coordinate the carryover of units from one season to the next.

Whenever the topic of intelligence, whether artificial or human, is discussed, the notion of rationality (as economists use the word) must be considered. From Gode and Sunder [1993] to the recent Sunder [2002] analysis of previous results on zero-information agents, it is steadfastly maintained that zero-information agents who submit random bids and offers are not behaving rationally. Nonetheless, these agents are programmed to be rational in the limited sense that they cannot engage in trades that generate an immediate loss; however, it is possible to imagine situations where a larger strategic purpose is served by presenting the market with a "loss leader." In addition, it is possible under certain circumstances that randomly placing orders could be part of an optimal bidding strategy.

All of this begs the larger question of what exactly is the optimal bidding strategy for human or artificial agents participating in a double-oral auction. Miller [2002b] discusses this problem at length, noting specifically that the mere placing of a bid or offer has the effect of granting the market a valuable option. This creates the well-known free-rider that the optimal strategy of all agents is to wait for someone else to bid first, leading to the unfortunate situation that bidding never begins, no one trades, and no surplus is appropriated. While the random bidding of zero-information agents may not appear to be the optimal choice of a rational agent, it generates vastly superior results to a market where nothing ever happens. Indeed, a truly adaptive robot trader might well be able to recognize when it is necessary to "break the ice" and then, like the zero-information agent, behave in a sufficiently random manner in order to convey as little information as possible to the other agents.

The illusion that the intelligence of agents had little impact on how markets behave was useful in that it limited the problem of market design to the choice of the rules governing the market. Now that we have clear evidence that agent intelligence (artificial or human) matters, market design becomes both a more difficult problem and a more interesting one.

Acknowledgements

The author thanks Vernon Smith and Nat Wilcox for their helpful comments on an earlier draft of this paper.

References

Axelrod, Robert, 1984. *The Evolution of Cooperation*. New York: Basic Books.

Caginalp, G. and D. Balenovich,1999. "Asset Flow and Momentum: Deterministic and Stochastic Equations." *Phil. Trans. Royal Soc. London A*, Vol. 357, pp. 2119–2133.

Caginalp, Gunduz, David Porter, and Vernon L. Smith, 2000. "Momentum and Overreaction in Experimental Asset Markets." *International Journal of Industrial Organization*, Vol. 18, No. 1 (January), pp. 187–204.

Cliff, Dave and Janet Bruten, "Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments." HP Laboratories Bristol, HPL-97-91, (August).

Das, Rajarshi, James E. Hanson, Jeffrey O. Kephart and Gerald Tesauro, 2001. "Agent-Human Interactions in the Continuous Double Auction." *The 2001 Proceedings of the International Joint Conferences on Artificial Intelligence*, pp. 1169–1187.

Gode, Dhananjay K., and Shyam Sunder, 1993. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy*, Vol. 101, No. 1 (February), pp. 119–137.

Hayek, F.A., 1937. "Economics and Knowledge." *Economica*, Vol. 4 (New Series), pp. 33–54.

Hayek, F.A., 1945. "The Use of Knowledge in Society." *American Economic Review*, Vol. 35, No. 4 (September), pp. 519–530.

Miller, Ross M., 2002a. "Can Markets Learn to Avoid Bubbles?" *Journal of Psychology and Financial Markets*, Vol. 3, No. 1, pp. 44–52.

Miller, Ross M., 2002b. *Paving Wall Street: Experimental Economics and the Quest for the Perfect Market*. New York: John Wiley & Sons.

Miller, Ross M., and Charles R. Plott, 1985. “Product Quality Signaling in Experimental Markets.” *Econometrica*, Vol. 53, No. 4 (July), pp. 837–872.

Miller, Ross M., Charles R. Plott, and Vernon L. Smith, 1977. “Intertemporal Competitive Equilibrium: An Empirical Study of Speculation.” *Quarterly Journal of Economics*, Vol. 91, No. 4 (November), pp. 599–624.

Mirowski, Philip, 2002. *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge, UK: Cambridge University Press.

Porter, David P., and Vernon L. Smith, 1995. “Futures Contracting and Dividend Uncertainty in Experimental Asset Markets.” *Journal of Business*, Vol. 68, No. 4 (October), pp. 509–541.

Rust, John, John Miller, and Richard Palmer, 1992. “Behavior of Trading Automata in a Computerized Double Auction Market.” In *The Double Auction Market: Institutions, Theories, and Evidence*, edited by Friedman, Daniel, and John Rust, Reading, MA: Addison-Wesley, pp. 155–198.

Smith, Vernon L., 1962. “An Experimental Study of Competitive Market Behavior.” *Journal of Political Economy*, Vol. 70, No. 3 (June), pp. 111–137.

Smith, Vernon L., 1982. “Markets as Economizers of Information: Experimental Examination of the “Hayek Hypothesis”,” *Economic Inquiry*, 165-79 Vol. 20, No.2, pp. 165–179.

Sunder, Shyam, 2002. “Market as an Artifact: Aggregate Efficiency from Zero-Intelligence Traders.” Yale ICF Working Paper No. 02-16, 2002.

Van Boening, Mark and Nathaniel. Wilcox, 1996. “Avoidable Cost: Ride a Double Auction Roller Coaster.” *American Economic Review* Vol. 86, No. 3, pp. 461–77.

Williams, Arlington W., 1980. “Computerized Double-Auction Markets: Some Initial Experimental Results.” *Journal of Business*, Vol. 53, (July), pp. 235–258.