A FUTURES MARKET SIMULATION WITH NON-RATIONAL PARTICIPANTS

Michael de la Maza and **Deniz Yuret**
Numinous Noetics Group
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
mdlm@ai.mit.edu, deniz@ai.mit.edu

Abstract

This paper describes a set of experiments performed with an artificial futures market simulation. The non-rational market participants, which evolve simple strategies using genetic algorithms, compete against each other to make profits by buying and selling futures contracts. The dynamic and equilibrium behavior of the participants is studied under a variety of conditions. The results suggest that in simple markets with non-homogenous participants opportunities for making consistent profits over extended periods of time exist.

Introduction

Although futures contracts are a relatively recent addition to the cornucopia of financial products available to market participants, their rate of growth has positioned them as one of the most traded financial vehicles in the world. Scalpers, day traders, institutions, and hedgers all compete to exchange risk and return in ways that meet their objectives. All participants in the market have data about the cumulative actions of all of the other participants and they adjust their behavior to exploit perceived opportunities.

The dynamic and equilibrium behavior of such a complex adaptive system cannot be easily understood by straightforward application of traditional econometric analysis. In this paper we present a program that simulates an artificial futures market and describe several experiments designed to illuminate how differences among market participants can lead to consistent profits for some of the participants.

The most important feature of this simulation is that the market participants are not rational. Curiously, when the rationality of participants in a market is bounded, as it is in this work, formal analysis becomes more difficult, not less difficult, because the number of parameters needed to describe each participant increases. Thus, market simulations are a tool that help us study complicated systems.

We stress, however, that our simulation of an artificial futures market does not do justice to the intricacies of real futures markets. In an attempt to capture the essentials of the futures markets in order to study them, we may have inadvertently removed the very components that make them worth studying.

Efficient markets

A market in which prices fully reflect all of the information in an information set is an efficient market. When this information set consists of price and volume data the market is said to be weakly efficient. If the information set consists of all publicly available information then the market is efficient in the semi-strong sense, and if the set contains all information, both public and private, then the market is efficient in the strong sense.

The most commonly used model for describing the behavior of market participants is called the rational expectations equilibrium model. The rational expectations model makes two assumptions [Dornbusch and Fischer, 1987, Sargent, 1986]. First, market participants are *rational* and are able to optimize an objective function. Second, the same information is available to all market participants.

In this paper both of these requirements are weakened. The first condition is not met because the participants use genetic algorithms to set the parameters of a model. Genetic algorithms are not guaranteed to find optimal solutions and participants cannot search over the space of models. Such participants have limited, or bounded, rationality [Simon, 1957, Simon, 1982]. Also, in some of the experiments, participants have different information sets, thus the second condition is not met.

Note that prices in a market in which the participants have bounded rationality and which have access to different information sets may still converge to the level predicted by rational expectations theory (for an example see [Sargent, 1993]). Thus, at a macro-level these two types of markets may be indistinguishable.

An artificial futures market

This section describes the framework we used for testing various hypotheses about the market.

The market consists of a number of participants who buy and sell futures contracts. For the purposes of this paper the primary distinction between a futures contract and a stock is that the futures contract has an expiration date. On the day of expiration the price of the futures contract converges to the price of the underlying commodity, which is called the spot (or cash) price. For most futures traded on American exchanges the expiration dates are three months apart. In this simulation, the participants buy and sell futures that expire in one day, so they are effectively attempting to predict the next day's spot price. One iteration of the simulation consists of a single round of buy and sell orders on one futures contract.

In the simulations described below participants attempt to buy futures contracts at prices below the price on the expiration day and sell futures contracts at prices above the price on the expiration day. The number of futures contracts that have been bought is always equal to the number of contracts that have been sold; thus, one participant's gain is another participant's loss. No transaction fees are charged when a participant buys or sells contracts.

The price of the futures contract on the day of expiration is a function of previous prices. Hence, the market participants are in effect setting the price by their actions.

Each participant in these simulations learns a strategy that calculates a "fair" value for the future contract. During each iteration of the simulation this fair value is revealed and an equilibrium price is computed which balances the number of buyers and sellers. If there is a range of possible equilibrium prices then the equilibrium price is set to be the average of the extremes of this range. If the fair value computed by a participant is above the equilibrium price, then the participant buys at the equilibrium price. If the fair value is below the equilibrium price then the participant sells at the equilibrium price. Note that in this simulation participants do not have the option of remaining on the sidelines and observing the price: they must buy or sell during each iteration.

Each strategy is a set of coefficients that is multiplied by the same set of variables used to determine the spot price. For most of the experiments described in the next section three variables are available to each of the participants: a ten day moving average of the equilibrium prices, the last equilibrium price, and the sign of the last change in the equilibrium prices. So the function used to compute the fair value is:

$$f(a, b, c, d) = a * x + b * y + c * z + d$$

where x is the ten day moving average of the equilibrium prices; y is the last equilibrium price; z is -1 if the last change in equilibrium prices is negative, 1 if it is positive, and 0 otherwise; and a, b, c, and d are the coefficients supplied by the strategy.

The coefficients used to determine the spot price for most of the experiments described below are: $a = 0.5$, $b = 0.5$, $c = 1$, and $d = 0$.

Each participant has a visible strategy which is used to compute the fair value of the futures price. In addition, each participant has a set of strategies that are not visible to other participants. Each of these non-visible strategies also computes a fair value for the futures contract. At each iteration of the simulation the best strategy (i.e., the one whose fair value is closest to the spot price) in the previous iteration becomes the visible strategy.

The strategies maintained by each participant are modified and improved by a genetic algorithm [Holland, 1975, Goldberg, 1989] with traditional mutation and crossover operators. The mutation operator adds a number uniformly distributed between -.1 and .1 to a coefficient and the crossover operator is traditional single point crossover. The best strategies are selected and copied using standard proportional selection with a constant offset that ensures that all fitness values are positive [Michalewicz, 1992]. The fitness values are calculated by subtracting the fair value computed by a strategy from the spot price on the next day.

Experiments

This section describes some experiments that suggest that when participants are asymmetric opportunities for making consistent profits exist.

In experimenting with this futures market simulation program we have found that small perturbations sometimes drastically change the behavior of the participants. As a result we have proceeded cautiously, changing only one parameter at a time in order to increase our understanding and intuition.

Preliminary experiments were used to fix several parameters of the simulations: the mutation rate, the number of participants, and the number of strategies per participant. The first experiment discussed below is an example of a preliminary experiment. It tests to what extent participants can ignore noise.

Experiment 2 serves as the gold standard against which experiments 3-7, which introduce sources of potential systematic profits, should be compared.

Experiment 1: Impact of noise

This experiment provides evidence that the participants can withstand noise. The three variables available to the participants, the ten day moving average, the last equilibrium price, and the sign of the last change, are all replaced with random numbers chosen uniformly between -5 and 5. The spot price is fixed at 1. So the participants should set the first three coefficients of their strategies to 0 and the fourth coefficient to 1. The participants' strategies are all initialized with small random numbers near 0.

Figure 1 shows how the equilibrium price changes as a function of time. Each day the simulation performs one iteration of the algorithm. This involves calculating an equilibrium price from the fair values submitted by the participants and updating the participants' strategies with the genetic algorithm.

Experiment 2: Comparison point

The second experiment provides a point of comparison for the rest of the results because all of the participants are identical. The spot price is a function of the moving average, the last equilibrium price, and the sign of the last change in the equilibrium price: $x/2 + y/2 + z$. Because all four participants are homogenous and have access to all relevant variables, the profit collected by the most profitable participant should approach zero over time.

Figure 2 shows the profit received by the most profitable participant as a function of time. The profit first increases sharply and then falls to less than 0.01. These profits are averaged over ten runs. Figure 3 displays the spot price, equilibrium price, and the fair value of the futures contract predicted by each of the four participants for one of the ten runs. Note that the coefficients of the strategies are initialized to values that cause the participants to overestimate the spot price during the beginning of the simulation.

Experiment 3: Forced liquidations

Experiment 3 simulates the effect of forced liquidations on the futures price and on profitability. With probability .2, a participant is forced to post a fair value that is one half of that computed by its visible strategy. These occasional shocks are designed to mimic unforeseen events, such as bankruptcies, that are sometimes experienced by market participants.

For this experiment the profit does not tend toward zero, as Figure 4 shows. Large profits continue to be made by those participants who are fortunate enough not to be forced to liquidate. Thus, participants with stable bankrolls appear to have an advantage over those who do not. To the extent that this is true of real world markets, they are not weakly efficient.

Figure 5 shows the spot, equilibrium, and fair value prices for one run. The equilibrium prices typically overshoot the spot price, thus providing some market participants with the opportunity to make profits.

Experiment 4: Unknown variable

In this experiment the spot price is partially determined by a fourth variable, the moving average over the last five days, that is not available to the participants. Thus, even the "optimal" strategy cannot correctly predict the

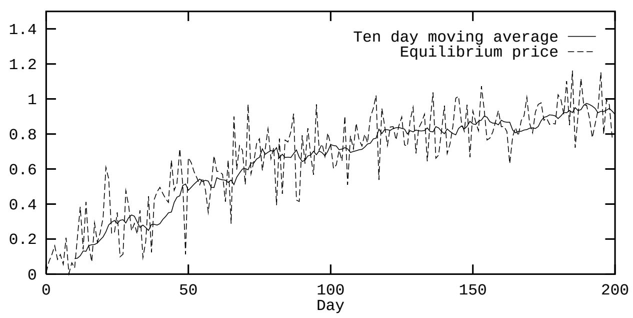

Figure 1: Equilibrium price as a function time for one run of experiment 1. In this experiment, the first three variables are randomly and uniformly distributed between -5 and 5 and the spot price is fixed at 1. This data is averaged over ten runs each of which lasted for 200 days.

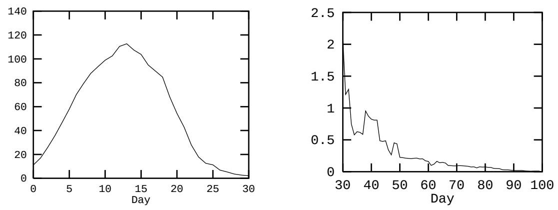

Figure 2: Profit of the most profitable participant as a function of time for experiment 2. As expected, the profit approaches zero as time passes.

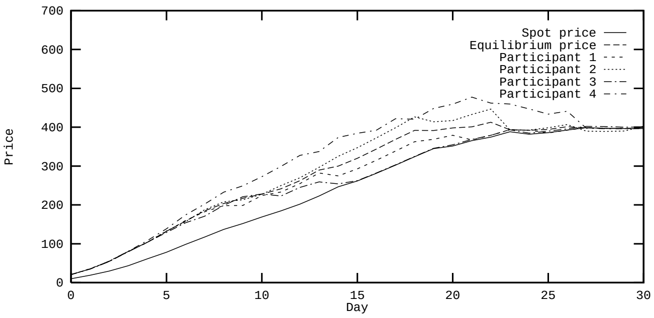

Figure 3: The spot, equilibrium, and fair value prices as a function of time for one run of the simulations performed for experiment 2.

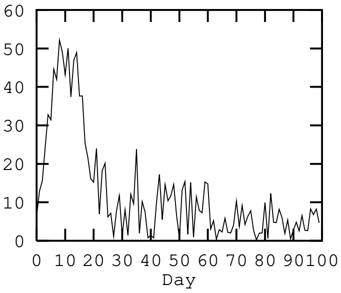

Figure 4: Profit of most profitable individual as a function of time in experiment 3. Unlike experiment 2, the profit does not tend to 0. This data is averaged over ten runs.

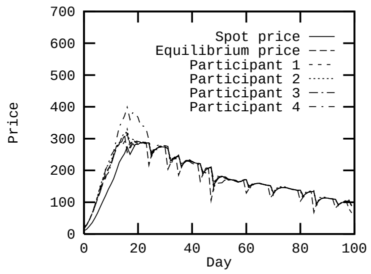

Figure 5: Spot, equilibrium, and fair value prices for one simulation of experiment 3. The dips in the price occur as a result of forced liquidations.

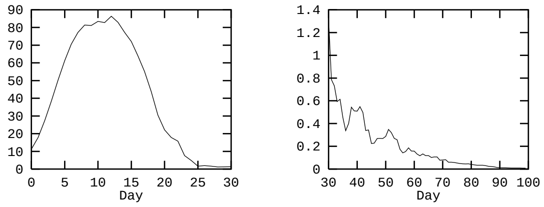

Figure 6: Profit of the most profitable participant as a function of time in experiment 4. Surprisingly, the profit achieved by the most profitable participant is less than in experiment 4.

spot price in all cases. The new spot price is: $x/4+y/4+4z/5+w/5$ where w is the five day moving average of the equilibrium prices and all other variables have been defined previously.

We expected profits to persist longer in this experiment than in experiment 2. To our surprise, as Figure 6 shows, profits are uniformly lower. Why is this the case? At first we thought that the addition of the five day moving average might reduce the initial sharp rise in the price, thus allowing the strategies to better estimate the spot price. However, Figure 7 demonstrates, for one representative example, that this is not the case.

The lower profits appear to be due to a reduction in the impact that the sign variable, z, has on the price. Because it is the most volatile variable and the only non-linear one, reducing its contribution increases the predictability of the spot price, therefore reducing profits.

Experiment 5: Protected participant

This experiment, in which only three of the four participants are forced to undergo liquidations, is to be compared with experiment 3. Over time, the protected participant is expected to accumulate profits that dwarf the profits of the other three participants.

Over ten runs, the protected participant accumulated a total profit of 4162.5 points, while the unprotected participants had profits of -56.2, -4128.7, and 22.4. The

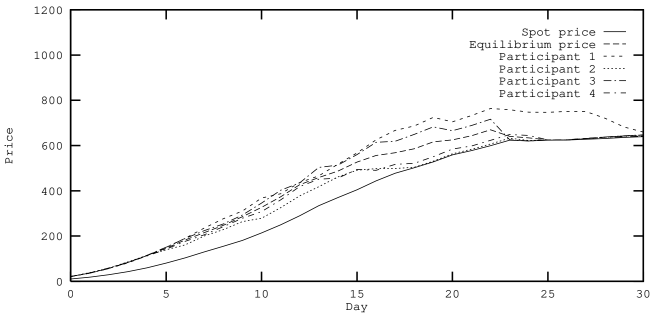

Figure 7: Spot, equilibrium, and fair value prices for one simulation of experiment 4. The spot price rises faster than in experiment 2, thus contradicting the hypothesis that the lower profits accumulated by the most profitable individual are due to increased price stability.

protected participant does not, however, accumulate the most profits in all of the runs. Figure 8 shows one simulation in which the protected participant is only slightly more profitable than one of the unprotected participants.

Experiment 6: Sharing strategies

Experiment 6 shows the effect of sharing strategies on the profitability of the participants. At the end of each iteration each participant incorporates the visible strategies of the other participants into its strategy set. We expected this sharing to increase the homogeneity of the strategies across the participants and, therefore, to reduce the profits accumulated by the most profitable participant.

These expectations were borne out by the results, which are summarized in Figure 9 and Figure 10. The profits are significantly less than in experiment 2 and the fair values computed by the visible strategies are all but indistinguishable.

This experiment helps to explain why most of the papers published in the academic literature about making money in the market report negative results.

Experiment 7: Asymmetric information sets

This final experiment explores how different information sets affect the profitability of participants. As before, one participant has access to all three variables, but the other three participants do not have access to the ten day moving average of the equilibrium prices.

Over ten runs the participant with the larger information set accumulates a profit of 666.67 units, while the other three participants have profits of -235.97, -173.81, and -256.89. The cumulative profits for one simulation run are shown in Figure 11. These results are in accord with our intuitions: the participant with the larger information set makes greater profits. Note, however, that this result may not hold as the information set increases in size because the search space becomes correspondingly larger.

A new description of participants

As a result of the experiments in this paper we have been led to a new model for characterizing market participants. Although this model contains no more informa-

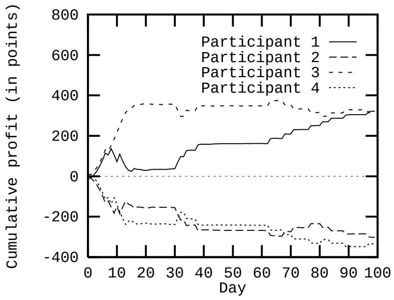

Figure 8: Cumulative profit as a function of time for experiment 5. Participant 1 is not subject to forced liquidations while the other three are.

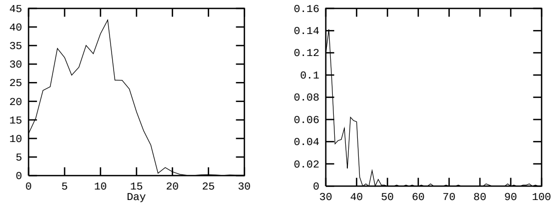

Figure 9: Profit of the most profitable individual as a function of time in experiment 6. As expected, the profits decrease more rapidly than in experiment 2. This data is averaged over ten simulations.

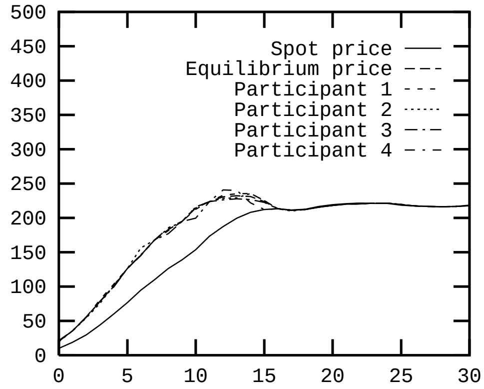

Figure 10: Spot, equilibrium, and fair values for one run of experiment 6. Note that the sharing of visible strategies causes the fair values to be all but inseparable.

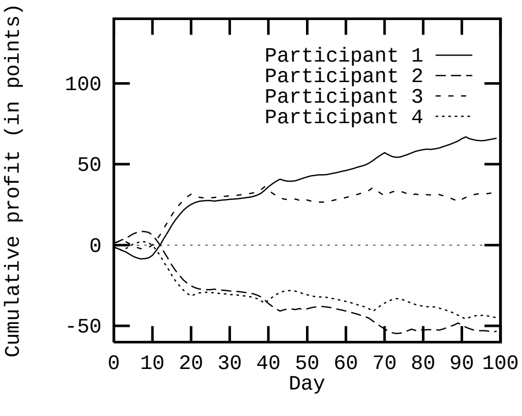

Figure 11: Spot, equilibrium, and fair values. Participant 1 has access to the ten day moving average while the other participants do not.

tion than the rational expectations equilibrium model, we feel that this description language better captures our intuitions about what differentiates market participants. Each participant is parameterized by four variables:

- Information set. What information does the participant have access to? Prices? Volume? Macroeconomic indicators?
- Algorithm set. What optimization algorithms and learning algorithms can the participant employ? Genetic algorithms? Linear regression? Neural networks?
- Model set. What kinds of models can the participant use to describe the data and predict future prices? Linear? Non-linear? Turing equivalent?
- Constraint set. What kinds of restrictions is the participant subject to? High transaction costs? Low initial capital? Sudden liquidations?

Some of this participant space has been explored in detail by economists. Linear models, for example, are much studied in econometrics and the effect of certain constraints on arbitrage opportunities is also well studied (see, e.g., [Kolb, 1991]). However, very little is known either about the dynamic or equilibrium behavior of systems in which the participants employ learning algorithms from artificial intelligence or in which the participants have grossly asymmetric information sets. In this paper we have kept the algorithm set and model sets fixed and changed the information sets and constraint sets of the participants. Exploring this space in a more complete fashion is a topic for future work.

Related work

Nottola, Leroy, and Davalo [Nottola *et al.*, 1992] describe an artificial market system in which agents (called participants in this paper) learn rules that predict movements in price. Much of their paper describes different approaches to modeling markets, but they do discuss one experiment which compares the performance of agents with different learning strategies and they conclude that adaptive agents out-perform non-adapting ones.

LeBaron [LeBaron, 1994] has developed a stock market simulation in which participants learn strategies using a genetic algorithm, as they do in this paper. The strategies of his market participants are encoded as bitstrings in which each bit corresponds to a boolean variable that is precomputed. His system features a mechanism for easily changing the risk aversion of individual participants. LeBaron notes [Freedman, 1993] that even if there are regularities in real market data that can be exploited, the data might contain so much noise that the regularities are barely perceptible.

There is substantial evidence that the currency futures market is not efficient [Thomas, 1986, Cavanaugh, 1987, Glassman, 1987, Harpaz *et al.*, 1990]. As empirical evidents mounts against the view that markets are efficient, theoreticians will be forced to provide alternative explanations of the market's behavior. One of the most likely candidates for this is Simon's bounded rationality theory [Simon, 1957, Simon, 1982]. Sargent [Sargent, 1993, page 4] explores markets in which the participants have bounded rationality and writes that "the bounded rationality program wants to make the agents in our models more like the econometricians who estimate and use them."

Conclusions and future work

As researchers we are interested in understanding whether or not certain participants in a market can consistently make profits. As money managers we are interested in uncovering opportunities for making profits. In this paper we have taken a small step towards addressing parts of both of these issues. Participants with stable bankrolls and access to more relevant information appear to have an advantage over other participants in our admittedly limited simulations. Some of the results that stem from this research have been used to manage a small options account since September, 1993. We participated in the options division of the 1993 U.S. Investing Championships using this account and finished fifth with a 43.9% return.

In the future we expect to create more accurate participant models of the market. We plan to circumscribe the scope and extent of the participant model described in section by running simulations with different settings for the four parameters. In particular, our current research focuses on expanding the participants' model sets so that they include explicit expectations about the model sets of other participants.

Acknowledgments

This paper describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research Contract N00014-91-J-4038. The first author is an NSF graduate fellow.

References

[Cavanaugh, 1987] K. Cavanaugh. Price dynamics in foreign currency futures markets. *Journal of International Money and Finance*, 6(3):295–314, 1987.

[Dornbusch and Fischer, 1987] R. Dornbusch and S. Fischer. *Macroeconomics*. McGraw-Hill Book Company, New York, 1987.

[Freedman, 1993] D. Freedman. Enter the market merlins. *Forbes ASAP*, 1993.

[Glassman, 1987] D. Glassman. The efficiency of foreign exchange futures markets in turbulent and non-turbulent periods. *The Journal of Futures Markets*, 7(3):245–267, 1987.

[Goldberg, 1989] D. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, 1989.

[Harpaz *et al.*, 1990] G. Harpaz, S. Krull, and J. Yagil. The efficiency of the U.S. dollar index futures market. *The Journal of Futures Markets*, 10(5):469–479, 1990.

[Holland, 1975] J.H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.

[Kolb, 1991] R. Kolb. *Understanding Futures Markets*. Kolb Publishing Company, Miami, 1991.

[LeBaron, 1994] B. LeBaron. An artificial stock market. MIT AI vision seminar colloquium, 1994.

[Michalewicz, 1992] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, Berlin, 1992.

[Nottola *et al.*, 1992] C. Nottola, F. Leroy, and F. Davalo. Dynamics of artificial markets: Speculative markets and emerging 'common sense' knowledge. In *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 185–194, 1992.

[Sargent, 1986] T. J. Sargent. *Rational Expectations and Inflation*. Harper & Row, New York, 1986.

[Sargent, 1993] T. J. Sargent. *Bounded Rationality in Macroeconomics*. Clarendon Press, Oxford, 1993.

[Simon, 1957] H. Simon. *Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in Society Setting*. John Wiley, New York, 1957.

[Simon, 1982] H. Simon. *Models of Bounded Rationality*. MIT Press, Cambridge, MA, 1982.

[Thomas, 1986] L. Thomas. Random walk profits in currency futures trading. *The Journal of Futures Markets*, 6(1):109–125, 1986.