The effects of periodic and continuous market environments on the performance of trading agents

Satpal Singh Chaggar[1], Jason Noble[1] and Dave Cliff[2]
[1]School of Electronics and Computer Science, University of Southampton
[2]Department of Computer Science, University of Bristol
jn2@ecs.soton.ac.uk

Abstract

Simulation experiments are conducted on simple continuous double auction (CDA) markets based on the experimental economics work of Vernon Smith. CDA models within experimental economics usually consist of a sequence of discrete trading periods or "days", with allocations of stock and currency replenished at the start of each day, a situation we call "periodic" replenishment. In our experiments we look at both periodic and continuous-replenishment versions of the CDA. In this we build on the work of Cliff and Preist (2001) with human subjects, but we replace human traders with Zero Intelligence Plus (ZIP) trading agents, a minimal algorithm that can produce equilibrating market behaviour in CDA models. Our results indicate that continuous-replenishment (CR) CDA markets are similar to conventional periodic CDA markets in their ability to show equilibration dynamics. Secondly we show that although both models produce the same behaviour of price formation, they are different playing fields, as periodic markets are more efficient over time than their continuous counterparts. We also find, however, that the volume of trade in periodic CDA markets is concentrated in the early period of each trading day, and the market is in this sense inefficient. We look at whether ZIP agents require different parameters for optimal behaviour in each market type, and find that this is indeed the case. Overall, our conclusions mirror earlier findings on the robustness of the CDA, but we stress that a CR-CDA marketplace equilibrates in a different way to a periodic one.

Introduction

The Continuous Double Auction (CDA) is a market institution that plays a fundamental role in the world economy. It is the principal trading format for commodity markets, equity exchanges, foreign exchange, and derivatives markets. Real-world examples of CDA-based markets include the NYSE and the Chicago mercantile exchange. Although we have a great deal of observational data on these markets, it would be both difficult and illegal to manipulate them experimentally. Our understanding of how CDAs work has therefore been greatly enriched first by the discipline of experimental economics (Smith, 1962), in which human subjects participate in economic games in the laboratory. More recently CDAs have been studied using the methods of artificial life: in agent-based computational economics (see Tesfatsion, 2002, for a review) the behaviour of a simulated market emerges from the interactions of many relatively simple trading agents.

Our particular interest is in how the temporal structure of a CDA can affect both overall market performance and the optimal strategies for agents participating in that market. We look at two variant CDAs: one is an explicitly periodic market in which there is a discrete trading period with daily opening and closing points; we refer to this as the day-based or periodic-replenishment (PR) market. The second variant involves a non-periodic or continuous-replenishment (CR) market which allows for trading without interruption. We refer to the continuous-replenishment variant of the CDA as the CR market. These two types of CDA have important real-world exemplars: most stock exchanges are day-based, for instance, whereas the global foreign exchange markets are continuous-time. Intuition suggests that these markets are significantly different playing fields. Our goal is to use an agent-based model to find out how different these two CDA variants really are.

Experimental economics

The motivation of experimental economics is to model economic phenomena using human participants in controlled laboratory situations. Smith (1962) conducted pioneering studies in which a small number of inexperienced human traders participated in a CDA and were able to reach a competitive equilibrium price and equilibrium quantity of a traded commodity. Smith derived a qualitative indication of the relationship of supply and demand curves in producing equilibrating transaction prices and presented results suggesting the replication of classical microeconomic theory, all from a surprisingly simple model.

Smith's studies are recognized as the standard modelling framework for CDAs and the simplicity of Smith's concept has been integral to its success. Recent research has focused on establishing the robustness of Smith's general findings and examining the fidelity with which these experiments reproduce phenomena from real CDA markets. The reproducibility of economic phenomena is important as it means that on the one hand market makers (e.g., a regulatory agency setting up a new marketplace) can use these experiments to develop fairer and more robust market mechanisms. On the other hand, traders (and the operators or regulators of financial institutions) can use results from experimental economics to identify and exploit strategic niches in their existing marketplaces.

Computational economics

If we take the human traders of the experimental economics paradigm and replace them with programs representing different trading strategies, we get agent-based computational economics (ACE) (Gode and Sunder, 1993; Cliff, 1997; Tesfatsion, 2002). An important aspect of this research has been finding the simplest algorithm capable of producing equilibrating market dynamics in a similar fashion to human participants. Cliff (1997) introduced the Zero-Intelligence Plus trading agent (ZIP) as an algorithm with minimal intelligence that nevertheless produced market behaviour that was very close to that of human traders. ZIP trading agents are a modified version of an earlier agent known as ZI (Zero Intelligence), created by Gode and Sunder (1993). ZI traders are simply stochastic agents that announce random prices for bids and offers. ZIP is able to model CDA price formation based on an intuitive heuristic "decision tree" algorithm coupled with elementary machine learning techniques (Cliff, 1997).

Computationally lightweight autonomously adaptive ("intelligent") trading agents (such as ZIP) are extremely significant given the emergence of virtual market-places. On the side of the market designer, iterative economic simulations using ZIP allow experiments to be conducted faster and yield significant results insofar as the ZIP trader can be seen as a realistic model. On the side of financial institutions that act within the market there is an incentive to replace human traders with automated trading agents. A fair chunk of work in ACE modelling to date concerns the use of agents inspired by the ZIP architecture in CDA markets. Studies have concentrated on evolving more robust agents and trading strategies. A basic ZIP agent acting in a periodic-replenishment (PR) CDA market with fixed supply and demand curves (as in the classic Smith experiment) has been used by a number of authors as the *de facto* benchmark for demonstrating equilibrating price formation with artificial agents.

The impact of replenishment in markets

Past work using intelligent agents in CDA markets has rarely explored the importance of the replenishment schedule within the market model. Round-the-clock 365-days-per-year environments are emerging at a fast rate in the real world, and yet continuous-replenishment modelas are perhaps one of the least discussed CDA variants (Cliff and Preist, 2001) in experimental economics. The standard Smith CDA model is conducted over discrete intervals known as trading days, and the dynamics of the market are centered around this day-trading structure. As not all real CDA markets are periodic the applicability of a day-based model to these variants is dubious. In what we believe to be the first human-based experimental economics studies to address this issue, Cliff and Preist (2001) explored the effect of removing periodicity from the standard CDA model by allowing continuous trading — i.e., switched from PR to CR CDA models. Cliff and Preist's general conclusion was that the ability of a CDA market to reach an equilibrium price did not seem to be affected by the switch from PR to CR. However, due to the inherent difficulties in human experimentation, the sample size in these experiments is really rather small.

Experimental aim

Our goal is to look at whether PR and CR markets produce different trading dynamics, and ultimately we would like to examine optimal trading behaviour across a wide range of different replenishment structures of the marketplace. In this paper we directly extend the work of Cliff and Preist (2001) by developing both continuous- and periodic-replenishment markets with ZIP traders instead of humans. We are especially interested in potential differences between the two market types that may have been too subtle to be detected given Cliff and Preist's limited sample size.

Method

We wrote computer simulations recreating the methods of Cliff (1997) and Cliff and Preist (2001), which are both adaptations of the experimental economics methods of Smith and Williams (1983). The method is a static model of a continuous double auction: i.e., supply and demand curves are fixed, and market participants ("traders") each privately know how many units they are willing to trade and the cost or value of each of their units, but not the allocations of any other traders.

There are 22 trader-agents in the simulated market: 11 buyers and 11 sellers. Each individual agent is allocated a private fixed *limit price*. The limit price specifies, for sellers, the minimum price at which they can sell, and for buyers, the maximum price at which they can buy. The difference between an agent's limit price and the actual transaction price they may achieve for the commodity is their utility — "profit" for sellers, "savings" for buyers. Limit prices for each of the agents are different, i.e., the agents vary in how much the commodity is worth to them. Limit prices range between \$0.75 and \$3.25 as shown in figure 1.

At the start of the experiment the 11 buyers and 11 sellers enter the market, with the sellers each in possession of one unit of the commodity, and the buyers each seeking to purchase one unit. We refer to these units as the agents' *entitlements* to buy or to sell. A single experiment — in the standard, periodic-replenishment case — consists here of a

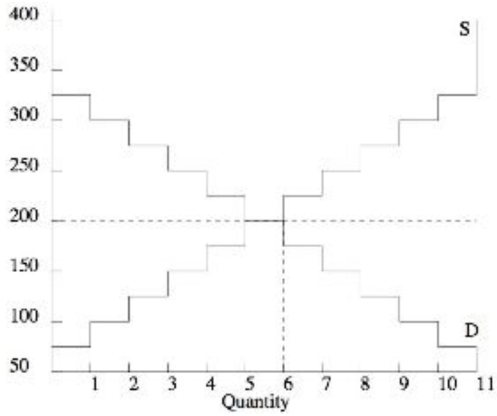

Figure 1: Stepped market Supply curve S and demand curve D, for 11 buyers and 11 sellers. Vertical axis is price in cents ($0.50 to $4.00); equilibrium price P_0 = $2.00. Supply and demand curves are fixed and symmetrical for all experiments. Figure is reproduced from Cliff (1997).

sequence of 20 trading periods, referred to as days. Each day is separated into 120 trading intervals (referred to as ticks). A tick is a discrete boundary of time at which a complete trading interaction can be executed (i.e., up to 120 attempted trades can take place during a day). Buyers and sellers normally have their entitlements reset to at the start of each trading day, by replenishing money to buyers and stock to sellers.

The arrangement of buyer and seller limit prices creates a stepped supply and demand curve for the imaginary commodity with a theoretical equilibrium price (P_0 = $2.00) and theoretical equilibrium quantity ($Q_0 = 6$) of units traded. Economic theory suggests that for rational agents participating in such a market, trading dynamics will show the competitive equilibration colloquially known as "the laws of supply and demand". In excess demand (trading taking place below the equilibrium) there is an incentive for the buyers to raise their bids to ensure they make a trade, and in excess supply (trading taking place above equilibrium) there is an incentive for sellers to lower offers to ensure a successful trade with a buyer (Cliff, 1997).

Trading process

With the market set up as described, buyers and sellers then engage in a CDA, in which they are free to announce and accept bids and offers for the commodity. The auction procedure is the same as that used by Cliff (1997).

1. At each tick a randomly selected agent quotes a price. This will be a *bid* if the agent is a buyer or an *offer* if the agent is a seller. The quoted price is made public to all agents from both communities and is the future transaction price for the trade. The agent's choice of price to quote is a function of its strategy.
2. Agents of the "contraside" (i.e., buyers responding to an offer, or sellers responding to a bid) make an assessment on whether dealing at the quoted price would be profitable for them. Again, this decision is a function of the agent's strategy. For ZIP agents, the decision will be influenced by their limit price but also by their current estimated valuation which is based on the recent history of successful trades in the marketplace.
3. If no willing agents are present in the market, i.e., the quoted bid is too low or the offer is too high, that tick-step is designated as a failed trade, and the market progresses onto the next tick.
4. If an agent decides that the shouted price is acceptable, it designates itself as a willing agent.
5. Prices of willing agents are arranged into a queue similar to NYSE rules (i.e., a trader makes a bid or offer at any time, but once made it is persistent until the trader alters it for a better price or it is accepted).
6. An agent is chosen from the queue, and the *quoted* price is the transaction price for the trade. The entitlements of both agents decrease by one and the profit and bank balances of the agents are adjusted according to the transaction price.
7. Finally, agents are assessed on their market activity state. Agents with no remaining entitlements to trade drop out of the market (although entitlements may later be reset, e.g., at the beginning of the next trading day).

A day's trading can be terminated prematurely if there are no active agents remaining in the market. Otherwise the market is open for 120 ticks, the duration designated for a trading day. For our markets we arbitrarily set the number of trading days to 20, to measure market performance over a reasonable period of time.

The periodic CDA

The replenishment schedule in a CDA market model effectively determines how and when the buying and selling entitlements of traders are reset. The periodic-replenishment (PR) variant is the default condition that has been described above; this is a replication of the Smith and Williams (1983) and Cliff (1997) models. The PR market forces the simultaneous and uniform renewal of all trading entitlements at the start of each day.

The continuous-replenishment CDA

For the continuous-replenishment (CR) CDA we recreate the market model of Cliff and Preist (2001) where there is no division of time into trading days. Once opened, the market continues for 2400 ticks until the end of the experiment. Every 120 ticks (the equivalent time frame for a day in periodic market) the entitlements for each agent are updated independently and with staggered phases. In short, the market is always open, and although agents temporarily drop out of trading after successfully buying or selling their single unit, they will return to trading at a randomly determined point in the future.

We have implemented two variations of the staggered renewal of agent entitlements, one referred to as *periodic continuity* or continuous(P) and the second referred to

Figure 4: Modelization of an example of the Gene Regulatory Network. A, B, C and D are 4 actions with their efficiency coefficient. The transfer coefficients are given by the arrows.

experiment consists in developing a system able to move substrates in the environment whereas the second one creates simple shapes like starfish or jellyfish.

To find the creature the most adapted to a specific problem, we use a genetic algorithm. Each creature is coded with a genome composed of three different chromosomes:

- The list of available actions, a subset of the environment possible actions. This list allows the cell to activate or inhibit some actions.
- The action selection system that contains a list of rule to apply actions.
- The gene regulation network that allows cell specification during duplication.

The creature is tested in its environment that returns the score at the end of the simulation. To increase the genetic algorithm power, we use a computational grid parallelized genetic algorithm. This parallelization allows the computation of hundreds of creatures at the same time.

Experiments

Developing a transfer system

The first experimentation consists in developing a simple organ : a transfer system. In other words, the cell structure must be able to transport substrate from one point to another. To do that, we imagine an environment composed of 2 substrates:

- The red is the substrate that must be moved by the organism. This substrate has the specificity not to spread in the environment, in order not to impact on the organism work.
- A gray that will be used by the cell as fuel and duplication material.

The cell can perform the following actions:

- duplicate (needs one gray substrate and vital energy),
- absorb or reject substrate (consume vital energy),
- transform one gray substrate in vital energy.

We place 10 red substrate units into a specific cross of the grid (at the top left of the environment) and diffuse gray substrate all over the environment. The creature's score is given by the squared sum of the red substrate distance to the goal point (at the bottom right of the environment). The parameters of the genetic algorithm are:

- selection: 7 tournament competition with elitism,
- mutation rate: 5%; crossover rate: 65%,
- substitution: worst individuals,
- population size: 500 individuals,

Figure 6 shows the convergence curve of the genetic algorithm. It shows the variation of the minimum, the average and the maximum fitness of the population for each generation. The genetic algorithm's aim is to maximize fitness, which is the creature score. A relevant organism appears quickly. After 3 generations, the organism is able to move the red substrates but not in the right direction. After 10 generations, it is able to move closer to the goal point. The genetic algorithm converges after 22 generations (the average fitness is close to the best).

Figure 5 shows the development of the best organism[1]. We can see that only the cells on the way from the initial point to the end point are created. Moreover, the organism uses absorption and rejection actions to transfer the substrate gradually. Cells that overtake the final point die quickly so as not to interact in the transfer. During the convergence of the genetic algorithm, it is interesting to observe the evolution of the organism strategy towards the best solution. The first step is to learn to survive in the environment, absorbing gray substrate and transforming it in vital energy. The next step is to learn to duplicate in the right direction. Intermediate solution organisms are able to transport the red substrate

[1]Videos of all presented creatures in this paper are available on the website http://www.irit.fr/∼Sylvain.Cussat-Blanc

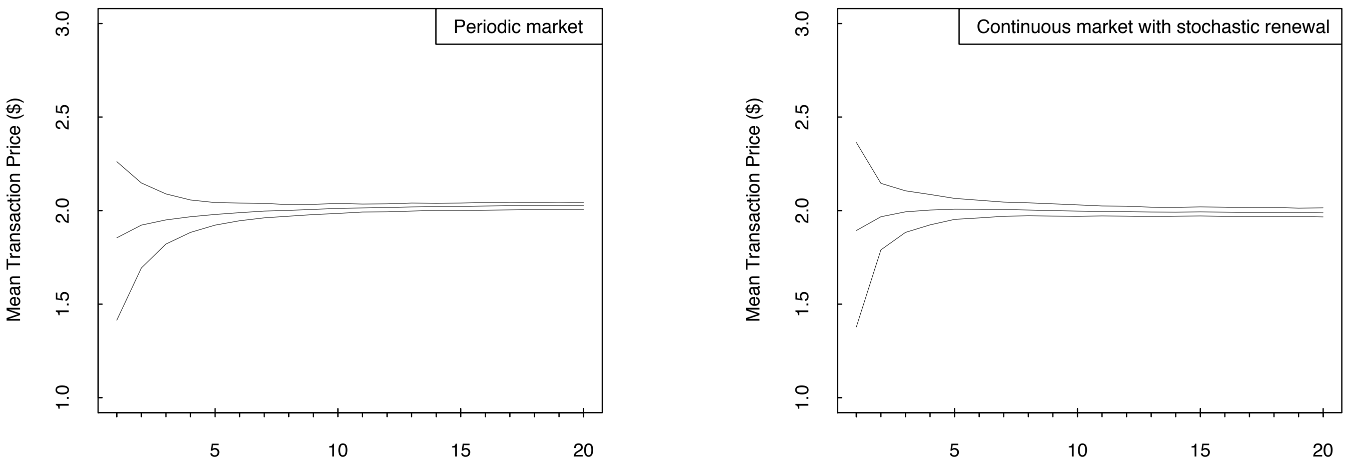

Figure 2: Mean transaction price over time (days or pseudo-days) for 500 ZIP experiments, for symmetrical supply and demand (P_0 = \$2.00) in a PR CDA (left), and a CR CDA with stochastic renewal (right). Dashed lines indicate the mean upper and lower transaction price boundaries at each day.

to occur in a range of prices around the equilibrium P_0 rather than convergence on the theoretical optimum price. Cliff and Preist (2001), in their continuous-time markets with human participants, found impressively low values of α that were below 0.1 within 600 seconds of the start of the experiment. Our overall data shows a failure to reach average α-values as low as 0.1, although we occasionally see single experimental runs with these low values.

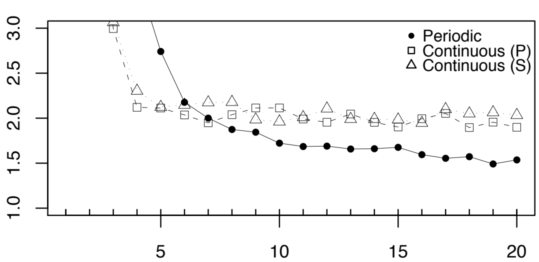

Figure 3: Mean α values over time (days or pseudo-days) for 500 ZIP market experiments: one periodic- and two continuous-replenishment CDA variants shown. Note that α values are very high (up to 20) in the earliest days of the experiments.

Hypothesis 3: CR markets exhibit greater experimental late-phase stability than PR markets. To investigate this question we split our data sets and looked only at results from the second half of each experimental run, i.e., days 11–20. This means we can look at market equilibration — effectively, long-term market efficiency — without the initial transients distorting the picture. We measure efficiency using both Smith's α and another measure, "profit dispersion". Gode and Sunder (1993) describe profit dispersion as the cross-sectional root mean squared difference between actual profits and equilibrium profits of an individual trader. For a group of n traders profit dispersion is given by $\sqrt{\frac{1}{n}\Sigma(a_i - p_i)^2}$ where a_i is the actual profit earned by trader i and p_i is p_0 for that trader. The more efficient the market, the lower the profit dispersion.

Figure 4 shows both mean α and mean profit dispersion for late-phase markets. Periodic-replenishment markets are consistently more efficient according to Smith's α, which means that transactions occur at prices closer to P_0 than in continous markets. There is also a very low variance in α-values for periodic markets. In terms of α performance the continuous markets with periodic renewal perform marginally better than the stochastic renewal version. In contrast, profit dispersion levels for all three market variants are approximately equal. This indicates that individual traders are not any more or less likely to trade at prices further from their personal equilibrium price in one type of market or another.

Hypothesis 4: Price formation in periodic markets is distributed around the opening of the market. We defined the "morning" period as being the first 25% of each trading day or pseudo-day, i.e., ticks 1–30. The trading volume during the morning period was approximately 3.5 times higher in periodic markets compared to continuous ones. This is not unexpected, as in the periodic CDA the entitlements of all traders are reset simultaneously as the market opens. This leads to an opportunity for many deals to be done immediately. More interestingly, despite the influx of entitlements to a morning market the transaction prices for periodic markets have a mean of 2.0147 (σ = 0.037). The transaction prices for both continuous markets in the morning period

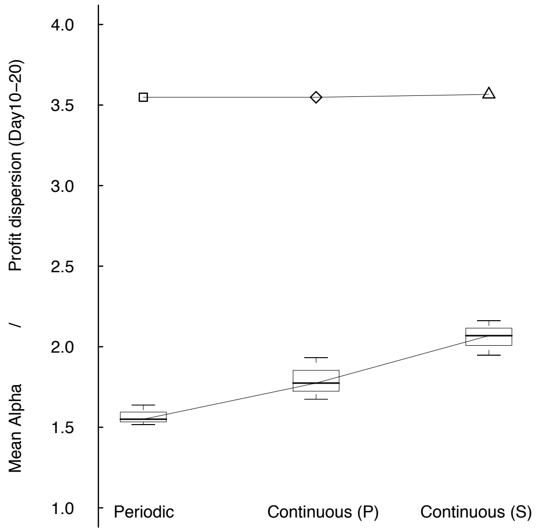

Figure 4: Mean convergence statistics for late-phase markets (days 11–20) with 500 ZIP experiments for each market model. Results for α indicate that the PR CDA is the most efficient, followed by the CR CDA with periodic renewal, and then the CR CDA with stochastic renewal. Average profit dispersion is roughly equal for all three types of replenishment.

are at a mean of 1.72 ($\sigma = 0.047$). Morning trade activity in periodic markets is very close to the equilibrium price despite the higher volume of trading. Approximately 79% of all experimental transaction occurs in the morning for a periodic market model, whereas in continuous markets the "morning" period has no particular significance and so obviously it accounts for 25% of trading.

Hypothesis 5: The optimal parameters for trading agents will take on different values depending on the market type. The behaviour of ZIP agents depends on a number of different parameters. Several different variables dictate the speed with which a ZIP trader modifies its price in the market, but the two most important are the Widrow-Hoff momentum (γ) and the agent learning rate (β). Preist (1999) demonstrates the significance of these variables. We looked at the effectiveness of different γ and β values in both periodic and continuous markets by creating surface plots of market efficiency, measured by Smith's α, for homogeneous communities of ZIP agents: see figure 5. We find that the resulting profiles of market efficiency are different for periodic and continuous CDAs. In other words, if I am a ZIP agent, the optimal settings for my core parameter values will depend on the market type I am in. In a periodic market, a value of $\beta = 0.2$ will produce the most efficient equilibrating market performance. Momentum γ, which acts to damp the oscillations for heuristic adjustments, can then vary across the range of 0.1–0.5 and this makes little difference to performance if $\beta = 0.2$ (see figure 5, left panel). In continuous-replenishment markets, the best example of market equilibration results from ZIP traders with $\beta = 0.1$. Continuous markets react more than periodic markets when γ is varied over the range 0.1–0.5. In a continuous market ZIP agents with lower momentum result in more efficient market behaviour (figure 5, right panel).

Intuitively we might expect that fast learning (a high value of β) and strong damping of adjustment oscillations (a high value of γ) would produce ZIP agents with more efficient market behaviour. Instead the trend for both markets is the opposite. Of course, we should be aware that ZIP parameters are not limited to γ and β. Our rationale for using these variables was not to find the most efficient ZIP trading strategy, but merely to illustrate that market replenishment style affects the way a ZIP trader should best operate.

It is also noteable from these results that some of our combinations of fixed γ and β ZIP variables produce markets that are almost 50% more efficient than those of the populations of ZIP agents used in the main set of experiments that featured the random assignment of parameters. This evidence is suggestive that there may be a market efficiency gain if all traders are uniform agents and consequently can be said to share the same idea of rational behaviour.

Discussion

Our experiments are, as far as we know, the first studies conducted with adaptive artificial trading agents operating in a simulation of a continuous-replenishment CDA. We have demonstrated the robustness of the CDA institution in fair price formation, by showing that groups of ZIP trading agents can consistently converge to the competitive equilibrium price and quantity governed by the supply and demand curves of the market. These results validate the observation of Cliff and Preist (2001) that both periodic and continuous markets can reach an equilibrium price. The use of simulation methods allows us to examine price formation variables more easily than in human-based experiments and we have therefore compared and contrasted the two CDA variants in more detail than was possible for Cliff & Preist in 1998.

Firstly, we found that profit dispersion between markets is almost identical in the later phase of the market for all three of our CDA variants. Secondly, we examined the α statistic over time, which calculates the divergence of market activity from the competitive equilibrium price. A comparison between the α values of periodic and continuous markets over time suggests that periodic markets equilibrate more efficiently over the long run than do continuous-replenishment markets. Comparing markets in late-phase allows measurements that are free from the effects of initial market turbulence, and thus facilitates a fair comparison between peri-

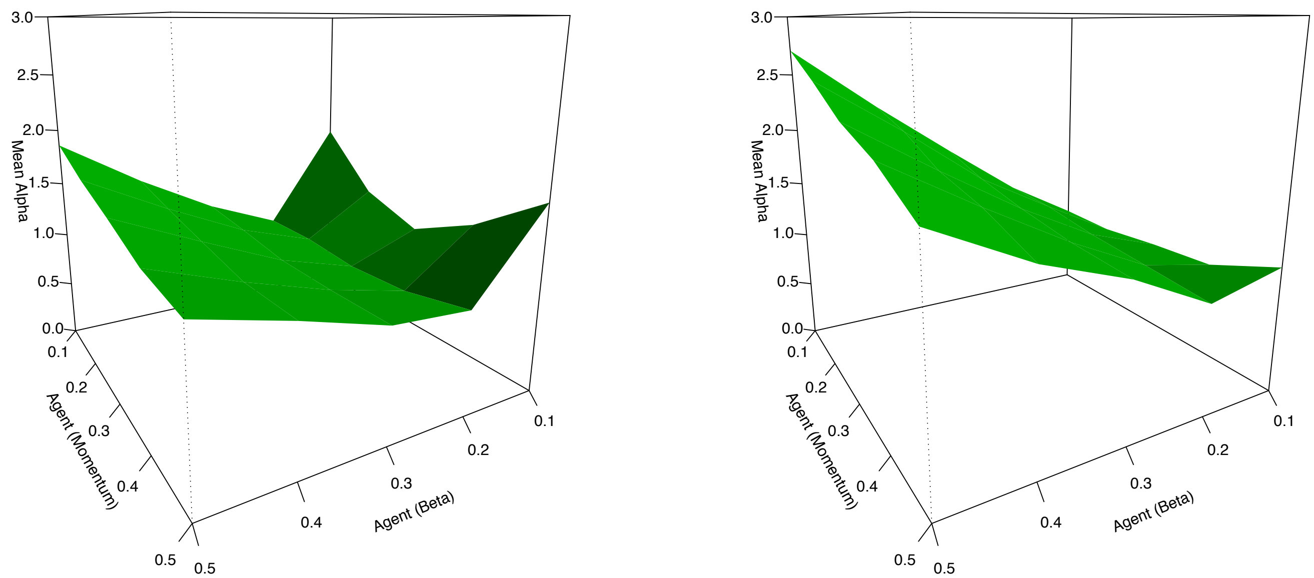

Figure 5: Surface plot of α against different values for γ (momentum) and β (learning rate) in a homogeneous ZIP agent community. The periodic market is shown on the left and the continuous market on the right. Results were generated over 500 experiments with all other agent parameters remaining at default values for ZIP version 1 (Cliff and Bruten, 1998). Note that the two surfaces are quite different, indicating that the two market types produce different optimal strategies.

odic and continuous markets. With no difference in profit dispersion across the three market types but with periodic markets achieving the most impressive (i.e., lowest) α values, this suggests that periodic markets represent a (near) Pareto-optimal solution to the problem of market design, with respect to our two measures of market efficiency.

Our original intuitions about the likely relationship between market efficiency and temporal structure were, in fact, the direct opposite of our results. We expected that the reoccuring event of an opening and closing of the market for our periodic variant would be enough to bring about a mini-fluctuation in the movement of opening prices each day and that possibly this pattern of trading would lead to oscillations around the equilibrium price at daily intervals. For CR markets, our intuition was that competitive price formation would occur early on and be maintained without such interruptions. Our original expectations can be summed up by the analogy that an engine that is continually restarted runs less smoothly than one that only starts once.

While it is not immediately clear why periodic markets over time deviate less from the competitive equilibrium price when compared to continuous markets, we can illustrate one reason for this behaviour from the perspective of the proportion of active agents within the market. The aggregate movement of price formation towards transactions at the equilibrium price only occurs if an agent is active within the market. For PR markets there is no potential delay in an agent being active for any given day, as by default all agents are deemed active at the start of each day. In a CR market, agents may in theory wait for a maximum time period equivalent to two days before being active within the market. An agent can only make meaningful contributions to the movement of the current trading price when it is active. Therefore in periodic markets, in which all agents start the day as active participants in the market, the collective action of all agents in reaching the equilibrium price will be maximally efficient. It may be that this activity being concentrated in time leads to the improved α values of the periodic market in comparison to the continuous ones.

Our average α-values for both the CR and PR variants of the market compare poorly to the reported α-values obtained by Cliff and Preist (2001) with human traders. This may well indicate relative inefficiency on the part of our ZIP agents, but it is also possible that the α-values reported by Cliff and Preist were the result of a regrettably small sample size.

The majority of PR market transactions occur within the "morning" period (i.e., the first 25% of the trading day), whereas in CR markets the trading activity is unsurprisingly spread across the trading day as the morning has no special significance. After the rush of morning trading, the remainder of the day in a periodic market is an empty trading environment, although quotes are still continuously made. In a sense, our PR markets "waste" most of the time of their participant traders, as (in these experiments) there isn't enough market surplus to fulfill the desired shouts; and so on average our PR markets were nonuniformly — and arguably inefficiently — used over the duration of each day. In contrast, the CR market successfully facilitates continuous trading. Many of these dynamics may be attributable to the assumption that each trader makes only one trade per day. However, even if agents traded many units per day we believe that a concentration of trading volume in the morning would remain characteristic of periodic markets as opposed to continuous. Empirically testing this belief remains a topic for further work.

How does periodicity of replenishment affect the agent? Our results suggest that groups of agents with uniform trading heuristics perform differently in each market. Therefore, each market requires a different trading strategy to produce the greatest efficiency or to extract the greatest utility. From the agent perspective, these two styles of market replenishment create two different playing fields. Results show that each market is capable of reaching the equilibrium price with intelligent trading agents, but it is important to emphasize that the greatest market efficiency is achieved by different agent strategies in the different marketplaces.

Questions concerning which of PR or CR as a market model is more efficient and which model offers the fairest profit distribution are hard to clarify. Indeed, if these questions were easy to answer, we assume that all real-world CDA markets would have converged to the optimal market model. The distinction between market types exists because each possesses different practical features in their own right.

Further work

While the results presented in this report illustrate new work on the CR market model, there are still many ways in which our experiments could be extended. Firstly, we limited our ZIP agents to handling only a single trade per day. Cliff and Preist worked with traders with multiple entitlements per day, who were also able to buy or sell multiple units in one transaction. The rationale for allowing our ZIP traders multiple daily entitlements would be to look at whether more sophisticated trading takes place, based on accumulated entitlements being filled at a later time in a continuous market.

We have kept our models of agents and markets simple in the interests of clarity. However, there are numerous features of the trading agent behaviour that could be improved. ZIP agents are unable to formulate a decision process that considers waiting in the market and making full use of continuous time (i.e., they cannot make a decision as to whether waiting is better than buying now). The ZIP agents used here are the original 1997-vintage "Version 1.0", now referred to as ZIP08 (Cliff, 2008). One consideration would to implement an optimising ZIP60 agent (Cliff, 2008) based on a genetic algorithm, to properly observe how different the optimised variables would be in each market. This would be a full extension of Hypothesis 5. Additionally a ZIP agent could also be made sensitive to CR markets by receiving more informative signals on how long the market has been running, and through greater temporal awareness being able to exploit strategies such as delaying the sale of a commodity in order to exploit a shortage and higher prices later on.

We could also be more rigorous in creating a framework that is completely free from synchronous behaviour. This is obviously desirable because of the asynchronous nature of real markets. The rate at which our agents update their price information is synchronised in our models, at each tick. It is possible that experimenting with an asynchronous and varied update rate for each agent could capture the asynchronous intelligence of real-world populations of traders.

References

Cliff, D. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Tech. Report HPL-97-91, Hewlett Packard Labs, Bristol, UK.

Cliff, D. (2008). ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Transactions on Evolutionary Computing*. In press, also available from http://eprints.ecs.soton.ac.uk/14078/.

Cliff, D. and Bruten, J. (1998). Market trading interactions as collective adaptive behavior. In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors, *From Animals to Animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*, pages 417–426. MIT Press, Cambridge, MA.

Cliff, D. and Preist, C. (2001). Days without end: On the stability of experimental single-period continuous double auction markets. Tech. Report HPL-2001-325, Hewlett Packard Labs, Bristol, UK.

Gode, D. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119.

Preist, C. (1999). Commodity trading using an agent-based iterated double auction. In *Proceedings of the 3rd Annual Conference on Autonomous Agents (AA'99)*, pages 131–138.

Smith, V. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–137.

Smith, V. and Williams, A. (1983). An experimental comparison of alternative rules for competitive market exchange. In *Auction, Bidding and Contracting: Uses and Theory*, pages 307–334. New York University Press.

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82.