

Overview of CAT: A Market Design Competition

Version 2.0
Issued: 30 April 2009

Kai Cai, Enrico Gerding, Peter McBurney,
Jinzhong Niu, Simon Parsons and Steve Phelps

1 Introduction

This document presents detailed information on the Trading Agent Competition (TAC) Mechanism Design Tournament, also known as the CAT Tournament or CAT Game, being held in July 2009. CAT is an abbreviation of *catallactics*, the science of economic exchange. The document has been prepared by the following people: Kai Cai (The Graduate School and University Center, CUNY, USA), Enrico Gerding (University of Southampton, UK), Peter McBurney (University of Liverpool, UK), Jinzhong Niu (The Graduate School and University Center, CUNY, USA), Simon Parsons (Brooklyn College, CUNY, USA) and Steve Phelps (University of Essex, UK). Responsibility for the contents of the document rests with Peter McBurney, GameMaster for the CAT Tournament in 2009. The document is version 2.0 of *catoverview.pdf*, and is issued on 30 April 2009. The document should be cited as [1].

The 2009 CAT Tournament is hosted by the Trading Agent Competition, which will be held at IJCAI 2009, in Pasadena, CA, USA. The document needs to be read in conjunction with other documents and software for the game, particularly the document defining the agent API and CAT Communication Protocol, *CATP* [12]. The various CAT documents and software, including the JCAT server platform, are available for public download from the sourceforge site for project JCAT:

<http://jcat.sourceforge.net/>

JCAT has been developed by a team at the Graduate School and University Center, and Brooklyn College, CUNY, USA, comprising Kai Cai, Rayman Chan, Jinzhong Niu and Simon Parsons. JCAT is a modification of the Java Auction Simulator API (JASA), initially developed by Steve Phelps (formerly University of Liverpool, now University of Essex, UK). Further information about JASA is available from here:

<http://jasa.sourceforge.net/>

The contents of this document are as follows:

Section 1: Introduction

Section 2: Game Overview

Section 3: Assessment

Section 4: Trading Agent Strategies

Section 5: Qualification Stage

Section 6: Further Reading

Section 7: Results of Past Games.

Operational testing of the CAT tournament software, simulation analysis needed to decide tournament parameters, development of web interfaces, and preparation of the CAT Tournament technical documents, have been undertaken with assistance from the following people, for which we are very grateful:

- Kai Cai, Rayman Chan, Joel Kammet, Jinzhong Niu, Simon Parsons and Elizabeth Sklar (Brooklyn College, CUNY, New York City, NY, USA)
- Dan Cartwright, Ken Chan, Peter McBurney, Thierry Moyaux and Dave Shield (University of Liverpool, UK)
- Ronan Bouichet, Benjamin Broix, Arthur Coussy and Mathieu Lassalle (Ecole des Mines de Nantes, France, and University of Liverpool, UK)
- Enrico Gerding, Albert Mmoloke and Krishnen Vytelingum (University of Southampton, UK)
- Steve Phelps (University of Essex, UK)
- Peter Stone (University of Texas at Austin, TX, USA).

Assistance with identification of bugs and anomalies was also gratefully received from participants in the 2007 and 2008 CAT Tournaments. Current information on the 2009 CAT Tournament is available here:

<http://www.marketbasedcontrol.com/cat/>

The CAT tournament is an initiative of, and has been financially supported by, the UK EPSRC project, *Market Based Control of Complex Computational Systems* (MBC Project) (Grant number GR/T10657/01), a project undertaken jointly by the Universities of Birmingham, Liverpool and Southampton, with the support of BAE Systems, British Telecommunications and Hewlett-Packard. We are grateful for the

support from the TAC organization and from the MBC Project. The views expressed in this document are those of the authors in their personal capacities only, and are not necessarily the views of their employers, sponsors or affiliated organizations or companies. In addition, the authors and their affiliated and sponsoring organizations are not responsible for actions taken, or not taken, or the consequences of such actions taken or not taken, as a result of the information contained in this document.

This document and its contents are Copyright 2007, 2008 and 2009 to the MBC Project. All rights are reserved. The contents of this document may not be copied without prior written permission from the authors.

1.1 Change Log

- Version 1.0 (May 2007): First version.
- Version 1.1 (June 2007): Revisions:
 1. Section 1 — Minor edits and additions to text.
 2. Section 2.4.1 — Details on client-server operation of game.
 3. Section 2.4.3 — Price lists broadcasted to all specialists.
 4. Section 2.4.4 — A specialist's information about the number of traders.
 5. Section 3.2 — Information added on the assessment score algorithms.
 6. Section 5 — Information added concerning the Qualification Stage and Final Stage of the 2007 CAT Tournament.
- Version 2.0 (April 2009): Revisions:
 1. Section 1 — Minor edits and additions to text; and web site addresses updated.
 2. Section 2 — Minor edits and additions to text; footnote regarding alternatives of socket-based communication added; and citation of [12] about CATP added.
 3. Section 2.2 — Minor edits and additions to text; illustration of paying information fees corrected; and 'accepting policy' changed to 'shout accepting policy'.
 4. Section 2.4.1 — Clarification of requirements regarding IP addresses and user names for specialists; explanation of revision to specialist connections and disconnection procedures added.
 5. Section 2.4.4 — 'Specialist selection' changed to 'market selection' to comply with the terminology in JCAT; merger of the two paragraphs about how traders and specialists respectively can use information about specialists; and a footnote about a specialist subscribing for information from itself added.
 6. Section 3.1 — Revision of text to align description with current CAT Game practice.

7. Section 3.2 — Minor edits and additions to text.
8. Section 3.3 — New subsection added on display of tournament results.
9. Section 4.1.2 — Typos corrected.
10. Section 4.2 — Addition of high-level information on methods of trader selection of specialists.
11. Section 5 — Clarification of purposes and operation of Qualification stage, and update for 2009 procedures.
12. Section 6 — New section added with pointers to further reading.
13. Section 7 — New section added with entrants and results of 2007 and 2008 CAT Tournaments.

2 Game Overview

A CAT game consists of *trading agents* (or simply *traders*), i.e., buyers and sellers, and *specialists*. Each specialist operates and sets the rules for a single exchange market, and traders buy and sell goods in one of the available markets. In the CAT tournaments, the trading agents are provided by the organizers, whereas specialists (and the rules of the markets) are designed by the entrants. Each entrant is limited to operate a single market.

During a CAT game, trading agents or specialists (markets), each as a client, connect to a central CAT server through socket.¹ CAT clients do not talk to each other directly; instead they send messages to the CAT server and the server responds to these messages and forwards information to other clients as appropriate. The communication between the server and the clients is regulated by the CAT Protocol (CATP), which is similar to HTTP for web communication. The detailed specification of CATP protocol is given in [12].

2.1 Trading Agents

Each trading agent is endowed with a *trading strategy* and a *market selection strategy* and assigned private values for the goods being traded. The traded goods are assumed to be homogenous and non-divisible. Furthermore, each buyer and seller has a certain demand and supply. Private values and the demand and supply of the markets are allocated by the CAT server, and are reset at the start of each day. Although private values remain constant during a day, these values may change from day to day.

The trading strategy is used to generate bids and asks (both also called *shouts*), whereas the market selection strategy is used to select a market or specialist. More detailed information about the trading agents' strategies is provided in Section 4.

Each trader is furthermore endowed with a limited budget that they can spend within a trading day. This budget prevents a trader from paying excessively high fees. Budget sizes are unknown to the specialist.

¹Socket communication is necessary for clients running locally around the world to participate in the CAT Competitions. However JCAT provides alternatives to run the server and multiple clients locally in a single process, which enables entrants to test their designs more quickly and in a deterministic manner.

2.2 Specialists

Specialists facilitate trade by matching bids and asks and determining the trading price in an exchange market. Each specialist operates its own exchange market. Each entrant in the game is required to design a single specialist or market, which is achieved by implementing the following policies:

- *Charging Policy*. This policy sets the fees (also referred to as *price list* in this document) which are charged to traders and other specialists who wish to use the services provided by the specialist. Each specialist is free to set the level of the charges (from zero up). These are the following:
 - *Registration fees*. Fees charged for registering with a specialist.
 - *Information fees*. Fees for receiving market information from a specialist.
 - *Shout fees*. Fees for successfully placing bids and asks.
 - *Transaction fees*. A flat charge for each successful transaction.
 - *Profit fees*. A share of the profit made by traders, where a trader's profit is calculated as the difference between the shout and transaction price.

The first 4 types of fees are each a flat charge, and the last one is a percentage charged on the profit made by a trader. A trader pays the registration fee to the only specialist it registers with once every trading day. A trader or specialist pays (other) specialists their information fees so as to receive markets information from them once every trading day. An example of how the last three fees are applied is given in Section 2.3.

- *Shout Accepting Policy*. This policy determines which shouts are accepted. A specialist has the option to reject shouts which do not conform to the specialist's policy. For example, a "beat the quote" policy requires a shout to beat the market quote. If the received shout violates this, it can be immediately rejected and will not be considered for a transaction, allowing the trader to submit a new shout. The policy, for instance, may reject any shout that is unlikely to trade so as to increase the specialist's transaction success rate, which is one of the assessment criteria for CAT Tournaments (see Section 3).
- *Clearing Policy*. This policy determines the way in which bids and asks are matched. For example, in the well-known continuous double auction, the market is cleared as soon as new shouts arrive.
- *Pricing Policy*. This policy determines the transaction price of a matched bid and ask. The most common is to set the price half way between the bid and ask.

Specialists are permitted and even encouraged to have adaptive strategies such that the policies change during the course of a game in response to market conditions.

A specialist will compete with the other specialists in the market on two levels. First of all, specialists compete for traders; at the beginning of each day of the game, a trader will choose a single specialist to register with and trade for that day (more details

are provided below). A specialist therefore has to design the market rules and fees such that the specialist attracts both buyers and sellers. Secondly, specialists compete to obtain the highest score in the tournament. This score each day is determined not only by the profit made by the specialist that day, but also by the market share (the proportion of traders registered with the specialist compared with other specialists), and the proportion of bids and asks which resulted in a successful transaction with that specialist on that day. These factors together are aggregated to determine the winner of the competition. A detailed description of the scoring rules is provided in Section 3.

2.3 Transaction Example

We present a simple example to illustrate the principles involved in a specialist charging fees to traders. We ignore the *registration* and *information* fees in this example and consider interaction between a single buyer, seller, and specialist.

Suppose a specialist has the following charging policy:

- Shout fees = \$2
- Transaction fees = \$5
- Profit fees = 10%

The *buyer* has private value of \$90 for each unit of the good and the *seller* has a private value of \$80; these values are unknown to the specialist. Suppose that traders shout their true values: thus, the buyer places a bid of \$90 and the seller submits an ask of \$80, and the specialist sets the transaction price to be \$86.

The fees are calculated as follows. The transaction and shout fees amount to \$7 for each trader. A trader's *profit* is calculated as the difference between the shout and the transaction price. In this case the difference between the ask and the transaction price is $90 - 86 = 4$, and the profit fee for the buyer is $4 \times 0.10 = \$0.40$. The seller's profit is equal to $86 - 80 = \$6$ and the seller pays a profit fee of $6 \times 0.10 = \$0.60$. The total fees for the buyer and seller are \$7.40 and \$7.60 respectively.

2.4 Game Procedure

In this Section we provide a detailed overview of the course of a CAT game. Each game consists of several *trading days*,² and each day consists of several rounds. We now explain what occurs at various stages in the game.

2.4.1 Before the Game Starts

As in previous years, the 2009 CAT Tournament will be operated from a server machine hosted in the Department of Computer Science at the University of Liverpool, UK. Entrants to the Tournament write programs to operate as specialists in the game. Participation in the Tournament will require a specialist to operate on a client machine,

²Note that a trading day is a virtual day and is different from an actual tournament day in real time.

hosted locally by the entrant concerned, which is able to establish and maintain continuous or nearly-continuous connection to the Liverpool client over the Internet. Prior to commencement of the Qualifying Stage Game and of the Final Stage Game (see Section 5), entrants who have registered with the TAC site will receive instructions to connect to the Liverpool server. Entrants need to provide an IP address and a short name for their specialist; for security reasons, subsequent connections to the Liverpool server and participation in the CAT games must be made from the given IP address and using the short name.

In normal operation of the game, all traders and specialists are required to *check in*, i.e., make their presence known to the server, and this is normally done just prior to the start of each game (i.e., before the first Trading Day). Each trader and specialist is then assigned a unique ID at this point, and this ID will remain unchanged throughout the course of a game. Typically, the ID is the same as the short name suggested by the entrant; the server will assign the suggested short name as the ID, provided that this short name does not conflict with existing IDs in the system.

After all the specialists and traders have checked in (or until a time out), the server announces that the game is starting and informs the clients of the number of rounds per day, and the length of each round (in milliseconds).

Modifications made to the JCAT server platform before the 2008 Tournament now permit specialists that have passed the Qualification Stage to arrive after the start of a game. In other words, there is no longer a technical requirement for specialists to check-in to the Liverpool server at the very start of each game. However, we advise entrants not to rely on late arrival of their specialist, as doing so may harm the competitive position of the specialist in the game. In case a specialist loses connection with the server at any time during the game, it is the responsibility of the specialist to recognize this event and to reconnect again with the server. When reconnecting, the specialist will log-in again, using the ID previously assigned by the server, and re-enter the ongoing competition from the next trading day. Once a game starts, the server will only allow existing clients to re-connect or new clients to connect if these clients arrive from machines with a registered IP address and using an approved ID (i.e., an IP address and ID previously approved by the Tournament organizers).

2.4.2 At the Start of a Game

After the start of the game, but before the first day, the server sends a list of all traders and specialists in the game to all traders and specialists. In particular, at this point the specialist will learn the number of traders in the game, and the number of competing specialists. Note that the total number of traders remains constant for the duration of a game.³ Although the number of traders is known, the bidding and market selection strategies used by traders are not made public and these may change over time.

³In the case of an unexpected and catastrophic failure — such as a loss of power to the Liverpool network, as happened in the 2007 CAT Tournament — the organizers of the Tournament may or may not restart the game. The decision as to whether to restart a running game will partly depend on how many trading days have passed, and how many specialists are currently connected.

2.4.3 Before a Day Starts

Before the start of *each day*, the specialists are required to announce their price lists. This information will be broadcasted to all traders and specialists. If a specialist fails to announce its fees before the actual day begins, the specialist will be excluded from trade for that day (but can participate again in the next day).

Meanwhile, the CAT server assigns to each trader a private value and the number of goods the trader needs to trade on the day. Different traders have different private values, each drawn from a random distribution. Currently, for any one trader, all the goods that the trader needs to trade have the same private value. These various private values and entitlements form the supply and demand schedules of the global market.

2.4.4 During a Day

Market Selection Once a day has started, traders can register with one of the markets (and only one market). Their choice of market is determined by their *market selection strategy* and depends on both the announced fees for that day, such as the registration fee, but also on the profits obtained in previous days. Different traders can have different market selection strategies (more details of the various market selection strategies in the game are given in Section 4.2), and the actual strategy used by a particular trader and the parameter settings remain hidden. The strategies used by the traders may also change over time. Whatever their strategy, however, obviously traders tend to go to specialists where they expect the highest profits.

Information about Specialists As mentioned above, before the actual start of the day, traders and specialists are automatically informed about the price lists set by all the specialists in the game. Alternatively, a trader or a specialist may freely request this information at any time once a day has started. Furthermore, at the end of each day, the traders and the specialists are automatically informed about both the profits obtained by all the specialists on that day (including their own profits) and the number of traders that have registered with each specialist.

In addition to the above free information, traders or specialists may subscribe to receive market information from one or more specialists. This information may not be free, and the price for a subscription (a fixed fee) is set by each specialist individually. The information that subscribers receive from a specialist is the following: (1) all of the successfully submitted shouts (i.e., bids and asks) with that specialist,⁴ and (2) all successful transactions.

Traders can use this information to adapt their *trading strategy* which is used to generate the bids and asks; a detailed description of the set of trading strategies used by bidders is given in Section 4. Note that a trader can subscribe to any number of specialists, and such a subscription is independent of the specialist that the trader registers with.

Specialists can use this information, for example, to track the decisions made by individual traders (such as their choice of specialist) using the unique trader ID which is

⁴Shouts can be rejected by a specialist, in which case they are not passed on as information to subscribers. See below for details.

provided with each shout. Also, the transaction success rate can be inferred from this information. Together with the information received at the end of the day regarding the profit and the number of traders, this can be used to calculate the score of other specialists (see section 3.2 for details on how this score is calculated).⁵

Trading We now discuss the actual trading. A day is divided into a number of rounds, during which traders submit shouts to the specialist they are registered with. Each shout specifies the limit price for a *single unit* of the traded good.

A specialist has the option to either accept or reject a shout. This can be used to enforce a certain bidding rule. For example, the rule “beat the quote” requires a shout to beat the current market quote. A specialist is not required to follow any explicit protocol, however, and can choose to reject or accept any shout. Once a shout is accepted, it becomes *active* and remains active either until a transaction is successfully completed or when the trading day ends.

For each trader only a single shout (for a single quantity of the good) can be active at any time. This means that a trader has to wait until a transaction is completed (or until the next trading day) before placing another shout. The reason for this restriction is to encourage early transactions and to avoid delays. Although a trader cannot submit a new shout when another shout is still active, the active shout can be revised by the trader by submitting a new shout. If the new shout is accepted by the specialist, this replaces the active shout. If the shout is not accepted, nothing changes. Note that traders do not pay any fees for revising their shouts. Also, note that revisions are not counted as new shouts and therefore do not directly affect the transaction success rate (see also section 3.2).

A specialist’s task is to match bids and asks and to determine the transaction price. Matching bids and asks can be done at any time during a round, provided the transaction is valid. For example, a continuous clearing mechanism will try and match any incoming bids as soon as they arrive. Alternatively, for example, clearing can occur at the ending of each round or only at the last round of the day. Whatever the timing, a transaction needs to be valid. This requires the bid price to match the ask price, and the transaction price to be set at some point in-between the bid and the ask price.

Note that active shouts persist across rounds, but expire as soon as a day ends.

3 Assessment

This section presents the process for assessment of entries to the CAT game. The core idea is that entrants will be assessed on multiple criteria (and not just the profits), which will be evaluated on a number of trading days.⁶ The assessment criteria to be used are described in Section 3.2 below. In order to avoid effects arising from the fact that

⁵A specialist can always subscribe for information from itself for free since it pays the information fee to itself. Technically speaking, a specialist can benefit from doing this due to the fact that it otherwise has no way to know the identities of traders who placed shouts to the specialist itself.

⁶Multiple criteria have been used for this tournament in order to better reflect competition between real-world exchanges, where the need by an exchange to enable specialists (or their equivalents) to make profits must be balanced with the need to attract traders to the exchange to ensure liquidity.

the tournament has a start-day and an end-day, not all the trading days will be used for assessment purposes. The process by which days will be selected is described in Section 3.1 below.

3.1 Sampling Method

- **Step 1:** For each Tournament game, we choose a random starting day, and a random ending day. Both selections are made prior to the commencement of the game, and remain secret from game entrants. The starting day will only be revealed at some point after it occurs. The ending day will only be revealed after that day has passed, which is normally near to the completion of the game.
- **Step 2:** At the end of each Assessment Day following revelation that the assessment process has commenced, the scores of each specialist against each criterion, and their total score, will be made public. For details of the process of publication of assessment results, see Section 3.3 below.

3.2 Scoring

Each specialist is assessed on three criteria on each Assessment Day, and these criteria are then combined into a single score for that day. These three criteria are as follows:

- **Profit Share:** The profit share score of a specialist on a particular day is given by the total profit obtained by that specialist on that day as a proportion of the total profit obtained by all specialists on that same day. The profit share score is thus a number between 0 and 1 (both inclusive) for each specialist for each day, and the profit share scores for all specialists on the same day will sum to 1.
- **Market Share:** Of those traders who have registered with a specialist on a particular day, the market share score of a specialist on that day is the proportion of traders that have registered with that specialist on that day. Again, the market share score is a number between 0 and 1 (both inclusive) for each specialist for each day, and the market share scores for all specialists on the same day will sum to 1.
- **Transaction Success Rate:** The transaction success rate score for a specialist on a given day is the proportion of bids and asks successfully placed with that specialist on that day which that specialist is able to match eventually. Let us denote:
 - N_b for the number of new bids successfully placed with a specialist on a given day,
 - N_a for the number of new asks successfully placed with that same specialist on the same day, and
 - N_m for the number of successful matches executed by that specialist from the bids and asks the specialist received that day.

Then, the transaction success rate score for that specialist on that day is calculated as:

$$\frac{2 \cdot N_m}{N_b + N_a}$$

In the case where both N_b and N_a are zero, the transaction success rate score is calculated as zero. Thus, the transaction success rate score for any specialist on any day is also a number between 0 and 1 (both inclusive). Note that *revised* shouts (see also section 2.4.4) are not counted as new bids or asks, and therefore do not affect the values of N_b and N_a .

Each of these three criteria results in a value for each specialist for each day between 0 and 1. The three criteria are then weighted *equally* (i.e., weighted one-third each) and added together to produce a combined score for each specialist for each Assessment Day. Scores are then summed across all Assessment Days to produce a final game score for each specialist. If multiple games are run during a CAT Tournament, which is typical in prior tournaments, the final score of a specialist will be the sum of its final scores of all the games. The specialist with the highest final score will be declared the winner of the tournament.

3.3 Display of Tournament Game Results

The JCAT platform includes a system output console to display the results of a CAT game on the server, updated as each trading day proceeds. Because not all trading days are necessarily used for assessment in Tournament games (as mentioned above), the console results are not made publicly available. Instead, information from the game operation is filtered according to whether the relevant Trading Day is an Assessment Day or not, and then, if so, presented on a dedicated webpage showing the progress of each Tournament game.

This webpage can be found here:

<http://cgi.csc.liv.ac.uk/~mbc/viewer/>

This webpage and its follow-on page are automatically refreshed (i.e., reloaded), approximately once per minute. As Trading Days which are also Assessment Days are completed in a game, the content presented on these pages is also updated. Entrants developing adaptive strategies for their specialists may use the data displayed on these webpages to update the knowledge base of their specialist and thus inform its strategy.

The software for the CAT Tournament web-viewer was written in PHP and HTML, and is available for download from the JCAT sourceforge pages, here:

<http://jcat.sourceforge.net/>

Note that use of the JCAT web viewer is not needed for the operation of the JCAT platform, nor for creating or running an entry to the CAT Tournament.

4 Trading Agent Strategies

This section describes the different strategies that are used by the trading agents. Each trader is equipped with both a bidding strategy for generating bids and asks, and a market selection strategy for selecting a specialist each day.

4.1 Bidding Strategies

Each trader uses one of the following four strategies: GD (Gjerstad Dickhaut), ZI (Zero Intelligence), ZIP (Zero Intelligence Plus), and RE (Roth and Erev). These strategies are selected since they have been extensively researched in the literature and some of these have shown to work well in practice. Note that the strategy used by a trader is private knowledge, and is not revealed to the specialists. Also, the proportion of traders executing each of the four strategies will not be revealed to specialists; this proportion is likely to change over time as the game proceeds. Finally, most of these strategies are adaptive and optimize their bidding behaviour.

We now explain the strategies in more detail.⁷

4.1.1 Zero-Intelligence

One of the most prominent strategies that have been developed over the years is Gode and Sunder's *Zero-Intelligence* (ZI) trading strategy [7]. The ZI agent is not motivated to seek trading profits and ignores all market conditions when forming a bid or ask by selecting a bid or an ask drawn from a uniform distribution over a given range (hence the term zero-intelligence). In terms of our categorisation, ZI is non-history-based and non-reactive as it does not consider the market condition in its decision-making process.

Specifically, Gode and Sunder considered the performance of both constrained and unconstrained ZI traders, which they respectively term ZI-C and ZI-U traders. The former are subject to budget constraints and are not allowed to trade at loss. The latter, however, are allowed to enter loss-making transactions (i.e. they are allowed to submit a bid (ask) that is higher (lower) than their limit (cost) prices). In more detail, the ZI-C buyer draws a bid from a uniform distribution between the minimum allowed bid and its limit price. The ZI-C seller forms an ask from a value drawn from a uniform distribution between its cost price and p_{max} , beyond which we assume no transaction can take place. For the ZI-U traders, the shout price is drawn from a uniform distribution between 0 and p_{max} . In the CAT game, we only consider ZI-C traders.

4.1.2 Zero-Intelligence Plus

The *Zero-Intelligence Plus* (or ZIP) strategy was first designed by [3] to show that more than zero-intelligence is required to achieve efficiency close to that of markets with human traders. ZIP has subsequently been used in a number of works as a benchmark for strategy evaluation (e.g. [4, 17, 20]).

⁷The descriptions of the trading strategies are taken from [18] and [14] with permission of the authors.

Specifically, while the ZI strategy ignores the state of the market and past experience, the ZIP strategy uses a history of market information and is of the predictive class, adapting the agent's profit margin to the future market conditions. That is, the agent increases or decreases its profit margin to remain competitive in the market. In this context, the profit margin determines the difference between the agent's limit price and the shout price.

At the beginning of the trading day, the agent has an arbitrarily low profit margin which it increases or decreases, depending on the different market events (submitted offers by buyers and sellers and any successful transactions) during the trading period. The ZIP buyer increases its profit margin whenever events in the market indicate that it could acquire a unit at a lower price than its current shout price, given by its profit margin. For a ZIP seller, if its last shout resulted in a transaction and its shout price was less than the transaction price, this indicates that it could transact at a higher price which would necessarily increase its profit margin.

Conversely, ZIP buyers and sellers reduce their profit margin when this margin is too high to remain competitive. In this case, the buyer would have market power (to influence the trend of ask prices) if the seller were to decrease its profit margin whenever an unsuccessful bid is submitted (since a series of unsuccessful bids would unnecessarily decrease the seller's margin). Similarly, the buyer should not lower its profit margin after each unsuccessful ask submitted in the market. Thus, a buyer lowers its profit margin only after a submitted bid is rejected, while a seller lowers its profit margin only after an unsuccessful ask. It is also necessary to consider the trader's strategy after a successful shout. If the last bid was successful, the unsuccessful seller lowers its profit margin so as not to be undercut by the competing sellers. The unsuccessful buyer lowers its profit margin after an ask was accepted by a competing buyer. We summarise the bidding behaviour of the ZIP trader in Figure 1, where $p_i^b(t)$ and $p_j^s(t)$ are the most profitable offer to buy or sell of ZIP buyer i and seller j respectively, at any time during the trading period and $s(t)$ denotes the price of the most recent shout. In particular, we have 6 different rules that specify when to increase or decrease the profit margin.

In this case, the profit margin is modified using a simple adaptive mechanism based on the Widrow-Hoff algorithm [21]. This is a continuous-space learning mechanism that back-projects the error between the current value and some desired value onto that current value. At any given time t , the ZIP trader i calculates the shout-price, $p_i(t)$, given its limit price, ℓ_i , and trader's profit margin, $\mu_i(t)$, according to the following equation:

$$p_i(t) = \ell_i(1 + \mu_i(t)) \quad (1)$$

The ZIP seller's margin is raised by increasing $\mu_i(t)$ and is lowered by decreasing $\mu_i(t)$, $\mu_i(t) \in [0, \infty)$. Conversely, the ZIP buyer raises and lowers its profit margin by decreasing and increasing $\mu_i(t)$ respectively. The initial profit margin, $\mu_i(0)$, is drawn from a uniform distribution over the range $[0.1, 0.5]$ at the beginning of the simulation. The aim of dynamically modifying $\mu_i(t)$ is for the trader's shout price to remain competitive against that of other participants. The learning mechanism of the

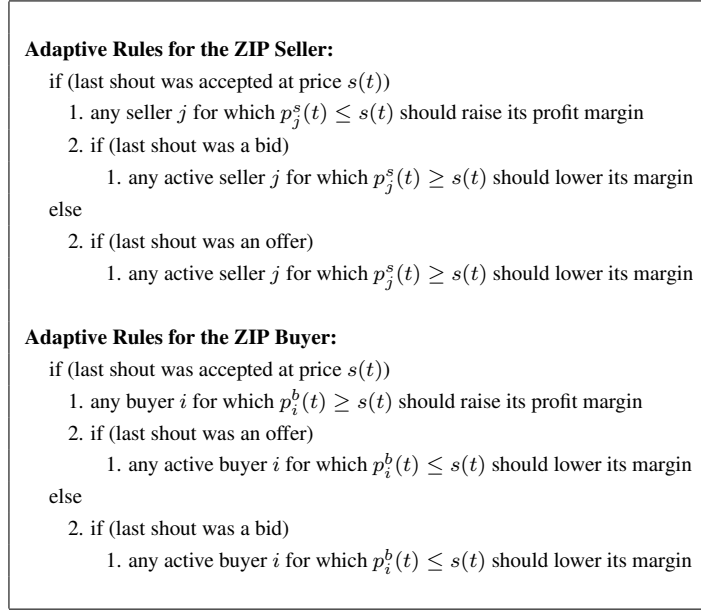


Figure 1: The ZIP Trading Strategy.

profit margin is given in the following equations:

$$\begin{aligned}
\mu_i(t+1) &= (p_i(t) + \Gamma_i(t+1))/\ell_i - 1 \\
\Gamma_i(t+1) &= \gamma_i \Gamma_i(t) + (1 - \gamma_i) \Delta_i(t), \\
&\text{where } \Gamma_i(0) = 0 \forall i
\end{aligned} \tag{2}$$

$$\begin{aligned}
\Delta_i(t) &= \beta_i (\tau_i(t) - p_i(t)) \\
\tau_i(t) &= R_i(t) s(t) + A_i(t)
\end{aligned} \tag{3}$$

where the learning coefficient, $\beta_i \in [0.1, 0.5]$, determines the rate of convergence of the trader's shout price toward the target price $\tau_i(t)$. Here, R_i is a randomly generated coefficient that sets the target price *relative* to the submitted shout price $q(t)$, with $R_i \in (1, 1.05]$ to increase $\tau_i(t)$ (when increasing the profit margin) and $R_i \in [0.95, 1)$ to lower $\tau_i(t)$ (when reducing the profit margin). $A_i(t)$ is an absolute perturbation, so the target price differs by at least a few units (of the minimum increment on the outstanding bid or ask), from even relatively small shout prices. $A_i(t)$ is drawn from a uniform distribution over $[0, 0.05]$ for an absolute increase and over $[-0.05, 0]$ for an absolute decrease. If we were to set $\tau_i(t)$ to $q(t)$, the trader would never be able to submit an offer.

Furthermore, [3] improved their learning mechanism to minimise the effect of high-frequency changes in bids or asks, by considering the momentum (trend) of shout

prices. The momentum-based updates are given in Equation 2. Specifically, the *momentum coefficient*, $\gamma_i \in [0, 1]$, determines the weight of previous shout prices on the change in the profit margin. When γ_i is equal to 0, the learning mechanism is myopic and ignores past quotes, while a high γ_i gives more weight to the trend of shout prices. In their simulations, γ_i is uniformly distributed over the range $[0.2, 0.6]$. Thus, the ZIP strategy has a set of 8 different parameters (β_i, γ_i, R_i and A_i to increase, R_i and A_i to decrease the margin, $\mu_i(0)$ for the buyer and the seller) that determine how to increase or decrease the profit margin which is specified by the 6 different rules.

4.1.3 Gjerstad-Dickhaut

The GD strategy, developed by [6], is based on a belief function that an agent builds to indicate whether a particular shout is likely to be accepted in the market. In the GD strategy, buyers form beliefs that a bid will be accepted and similarly sellers form beliefs that an ask will be accepted in the market. The traders form their beliefs on the basis of the history of observed market data and, particularly, on the frequencies of submitted bids and asks and of accepted bids and asks resulting in a transaction. Given this information, the bidding strategy is to submit the shout that maximises the trader's own expected surplus, which is the product of its belief function and its risk-neutral⁸ utility function. The GD strategy also implicitly considers the notion of recency, by limiting the trader's memory length, L , to a few transactions (L was set to 5 in Gjerstad and Dickhaut's simulations). The most recent history of shouts and transactions is considered.

In Gjerstad and Dickhaut's model, the seller's belief function, $\hat{p}(a)$, is based on the following assumptions. If an ask $a' < a$ has been rejected, then an ask, a , will also be rejected. Similarly, if an ask $a' > a$ has been accepted, then an ask submitted at a will also be accepted. Furthermore, if a bid $b' > a$ is made, then an ask $a' = b'$ would have been taken, since they assume that this ask a' would be acceptable to the buyer who bid b' . Similar assumptions are made about the buyer's belief function, $\hat{q}(b)$. We now define the bid and ask frequencies $\forall d \in D$, where D is the set of all permissible shout prices in the market, used in the belief function.

Bid Frequencies: $\forall d \in D$, $B(d)$ is the total number of bid offers made at price d , $TB(d)$ is the frequency of accepted bids at d , and $RB(d)$ the frequency of rejected bids at d .

Ask Frequencies: $\forall d \in D$, $A(d)$ is the total number of ask offers made at price d , $TA(d)$ is the frequency of accepted asks at d , and $RA(d)$ the frequency of rejected asks at d .

The *Seller's Belief Function* for each potential ask price, a , is defined as:

$$\hat{p}(a) = \frac{\sum_{d \geq a} TA(d) + \sum_{d \geq a} B(d)}{\sum_{d \geq a} TA(d) + \sum_{d \geq a} B(d) + \sum_{d \leq a} RA(d)} \quad (4)$$

⁸The risk-neutral agent uses a linear utility function and submits the price that maximises its expected profit.

The *Buyer's Belief Function* for each potential bid price, b , is defined as:

$$\hat{q}(b) = \frac{\sum_{d \leq b} TB(d) + \sum_{d \leq a} A(d)}{\sum_{d \leq a} TB(d) + \sum_{d \leq a} A(d) + \sum_{d \geq a} RB(d)} \quad (5)$$

The seller's belief function is modified to satisfy the *NYSE spread reduction rule*. Thus, for any ask that is higher than the current outstanding ask, the belief function is set to 0 (i.e. that ask cannot be accepted). Similarly for the buyer, the belief that any bid submitted that is below the outstanding bid is accepted is 0 (i.e. that bid cannot be accepted).

Furthermore, because the belief function is defined over the set of all bids and asks within the trader's memory, the belief is extended to the space of all potential bids or asks allowed in the market, constrained by the outstanding bid and ask and the step-size of the belief function. Then, *cubic spline interpolation* is used on each successive pair of data items to calculate the belief of points in between the pair of data. In particular, a cubic function, $p(a) = \alpha_3 a^3 + \alpha_2 a^2 + \alpha_1 a + \alpha_0$, that ensures each two successive pair of points is constructed with the following properties:

1. $p(a_k) = \hat{p}(a_k)$
2. $p(a_{k+1}) = \hat{p}(a_{k+1})$
3. $p'(a_k) = 0$
4. $p'(a_{k+1}) = 0$

The coefficients, $\alpha_i \forall i = \{0..3\}$, that satisfy the above properties, are then given by the solution to the following equation:

$$\begin{bmatrix} a_k^3 & a_k^2 & a_k & 1 \\ a_{k+1}^3 & a_{k+1}^2 & a_{k+1} & 1 \\ 3a_k^2 & 2a_k & 1 & 0 \\ 3a_{k+1}^2 & 2a_{k+1} & 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha_3 \\ \alpha_2 \\ \alpha_1 \\ \alpha_0 \end{bmatrix} = \begin{bmatrix} \hat{p}(a_k) \\ \hat{p}(a_{k+1}) \\ 0 \\ 0 \end{bmatrix} \quad (6)$$

Then, the buyer's belief function, $q(b)$, is constructed similarly using the pairs $(b_k, \hat{q}(b_k))$ and $(b_{k+1}, \hat{q}(b_{k+1}))$. Having defined the belief functions, Gjerstad and Dickhaut proved that the beliefs are monotonically non-increasing (see [6] for further details). Thus, the belief of an ask, $a > a'$ being accepted has to be lower than the belief of a' being accepted, and, similarly, the belief of a bid, $b < b'$, has to be lower than that of b . On this basis, we can clearly see that monotonicity of beliefs is an essential property that the belief function of the GD strategy must satisfy.

Given the belief function, the GD strategy forms an offer to buy or sell that maximises the trader's expected surplus (which is defined as the product of its belief function and its utility function, $\pi(a)$). Because Gjerstad and Dickhaut consider risk-neutral traders, the utility function is linear and equals the profit of the traders; that is, the difference between the seller's ask price and its cost price, and the difference between the buyer's bid price and its limit price. When the trader's maximum expected surplus is negative, there is no incentive to submit a bid or an ask and the trader abstains from bidding. The trader's utility function and its surplus maximisation is formulated as follows:

For a buyer i ,

$$\pi(b) = \begin{cases} \ell_i - b & \text{if } b < \ell_i \\ 0 & \text{if } b \geq \ell_i \end{cases}$$

For a seller j ,

$$\pi(a) = \begin{cases} a - c_j & \text{if } a > c_j \\ 0 & \text{if } a \leq c_j \end{cases}$$

$$b^* = \arg \max_{b \in (o_{ask}, o_{bid})} [\pi(b) \cdot q(b)] \quad (7)$$

$$a^* = \arg \max_{a \in (o_{ask}, o_{bid})} [\pi(a) \cdot p(a)] \quad (8)$$

4.1.4 Roth-Erev

The Roth-Erev algorithm is a strategy designed to mimic human game-playing behaviour in extensive form games [5]. The strategy relies on the immediate feedback from interacting with the mechanism, namely the surplus that the agent was able to obtain in the most recent round of trading. Hence, this strategy is general-purpose enough to be used with any auction-mechanism, even where the trader does not have access to market-data, for example, in repeated *sealed-bid* auctions.

Roth-Erev is one of a class of strategies in which an agent a_i chooses its markup over its valuation price v_i as follows:

$$\text{For } a_i \text{ a seller: } \zeta(i, t) = v_i + RL_{\lambda_i}(t)RL_{\mu_i}$$

$$\text{For } a_i \text{ a buyer: } \zeta(i, t) = v_i - RL_{\lambda_i}(t)RL_{\mu_i}$$

where $RL_{\rho_i}(t)$ is a *reward signal* which represents the most recent profits of agent a_i :

$$RL_{\rho_i}(t) = \Gamma_t(a_i) - \Gamma_{t-1}(a_i) \quad (9)$$

The function $RL_{\lambda_i} : \mathbb{N} \rightarrow \Theta_i$ represents the output of learning algorithm λ where $\Theta_i = [0, RL_{k_i}) \subset \mathbb{N}$ is the set of possible outputs from λ .

Parameter name	Semantics
$RL_{\lambda_i}(t)$	A function specifying the output from a 1-armed bandit learning algorithm
RL_{μ_i}	A scaling factor used to map learning outputs onto actual prices
RL_{k_i}	The number of possible outputs from RL_{λ_i}

Table 1: Reinforcement-learning parameters

For the Roth-Erev strategy, we set:

$$RL_{\lambda_i}(t) = RE_i(t) = \delta_i t \quad (10)$$

where $\delta_t \in \Theta_i$ is a discrete random variable distributed as:

$$P(\delta_i t = x) = RE_p(x, i, t) \quad (11)$$

<i>Parameter name</i>	<i>Semantics</i>
RE_{k_i}	The number of possible outputs
RE_{ρ_i}	The recency parameter
RE_{η_i}	The experimentation parameter η
RE_{s_i}	The scaling parameter

Table 2: Parameters for the Roth-Erev learning algorithm

<i>State variable</i>	<i>Semantics</i>
$RE_i(t)$	The output of the learning algorithm at time t
$RE_p(\theta, a_i, t)$	The probability distribution over each possible action $\theta \in \Theta_i$
$RE_q(\theta, a_i, t)$	The <i>propensity</i> for each possible action $\theta \in \Theta_i$

Table 3: State variables for the Roth-Erev learning algorithm

The propensities are initialized based on the scaling parameter RE_{s_i} ; $\forall a_i \in A$ and $\forall \theta \in \Theta_i$:

$$RE_q(\theta, a_i, t_0) = \frac{RE_{s_i}}{RL_{k_i}}. \quad (12)$$

The RE_q are then updated based on the reward signal RL_{ρ} :

$$RE_q(\theta, a_i, t) = (1 - RE_{\rho_i})RE_q(\theta, a_i, t - 1) + RE_{\rho_i}(\theta, a_i, RL_{\rho_i}(t - 1)) \quad (13)$$

and then normalized to produce a vector of probabilities; let Q_{it} denote the sum of all the propensities for agent i :

$$Q_{it} = \sum_{\theta \in \Theta_i} RE_q(\theta, a_i, t) \quad (14)$$

Then, $\forall \theta \in \Theta_i$ and $\forall a_i \in A$:

$$RE_p(\theta, a_i, t) = \frac{RE_q(\theta, a_i, t)}{Q_{it}} \quad (15)$$

4.2 Market Selection Strategy

Each trader also executes a strategy to select a market (i.e., a specialist) prior to each trading day. For some traders, on some trading days, this selection process is random. On other occasions or for other traders, selection by the trader of a market is based on the trader's prior transaction experience with the specialists in the game. Thus, when the market selection strategy is not random, it incorporates a learning component.

Traders who do not use a random market selection process treat the problem as an n -armed bandit problem, and solve this problem using either an ϵ -greedy exploration policy or a *softmax* exploration policy [16]. A trader using the ϵ -greedy exploration policy determines what it believes to be the best marketplace in terms of daily trading profit and then chooses this marketplace with a probability of $1 - \epsilon$. One of the other marketplaces is selected with a probability of ϵ . Under this policy, all those markets which are not the best market are selected with the same probability. A trader using the *softmax* exploration policy does not treat all non-best markets equally in this manner. Rather, under this second policy, when the best market is not selected, then better markets among the remaining markets are selected with a higher probability than are those remaining markets deemed by the trader to be less good.

Specific details of the market selection strategies used by traders are not provided to specialists.

5 Qualification or Trial Stage

As in previous years for the CAT Tournament, a Qualification (or Trial) Stage will be organized prior to the final 2009 CAT Game. This Trial Stage in 2009 will be held in the week of 15–19 June 2009.

The main purpose of the Trial Stage of the tournament is to test Internet connectivity and inter-operability of each entrant with the Liverpool game server. The Trial Stage is not intended to test the competitive performance of each specialist, but entrants are free to use the games of the Trials for this purpose if they wish. As in 2008, if the number of entrants is not large, we will allow each team to connect two specialists in tandem in this stage.

The name, IP address and port number of the Liverpool server will be provided to registered entrants ahead of the first Trial Stage game. Intending entrants need to provide to the organizers the IP address of their local client machine from which their specialist will connect to the Liverpool server, a short name to identify their specialist, and contact details for their team (usually a *skype* alias and a phone number).⁹ The local client machine needs to have Internet access and to be able to communicate directly with the server at the University of Liverpool in Liverpool, UK. If technical or legal circumstances in the country of the client machine do not permit communication to the server in Liverpool, an intending entrant is free to use another local client remote from the entrant, e.g., a client machine physically located in another country from the team itself.

Progression by a team from the Trial Stage to the Final Stage will occur if the team's specialist participates in the qualification game with merely minimal functionality. In other words, if a specialist operates without problems in the Trial Stage, even if that specialist scores poorly on the assessment criteria, then we expect it will be permitted to proceed to the Final Stage. However, if the number of entrants to the Tournament is

⁹The exact requirements will be communicated to registered entrants via email from the 2009 CAT Tournament Gamemaster. Intending entrants are advised to ensure they are on the mailing list for these announcements by sending an email requesting this to the 2009 CAT Gamemaster at: mcburney@liverpool.ac.uk

large, performance in the Trial Stage game may be used to provide individual seedings or handicaps for the Final Stage games.

The Final Stage of the 2009 CAT Tournament will be held alongside the *Trading Agent Design and Analysis (TADA 2009) Workshop* at IJCAI 2009, in Pasadena, CA, USA. The duration and dates of the Final Stage games will be determined in the light of the experience gained in running of the Trial Stage, and will be announced by early July 2009.

6 Further Reading

The CAT Tournament has now been run twice — in 2007 and in 2008. Several refereed research publications have now been published, arising from the work undertaken by the entrants and the game organizers, and these are listed in order of publication below. We recommend to potential entrants that they consult these publications. Additions are also welcome. Current links to these publications online can be found on the CAT Information page, here:

<http://www.marketbasedcontrol.com/cat/>

- J. Niu, K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, S. Phelps and D. Shield [2008]: JCAT: A platform for the TAC Market Design Competition. In: L. Padgham, D. Parkes, J. P. Mueller and S. Parsons (Editors): *Proceedings of the Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*. Estoril, Portugal, May 2008. [11]
- J. Niu, K. Cai, E. Gerding, P. McBurney and S. Parsons [2008]: Characterizing effective auction mechanisms: Insights from the 2007 TAC Market Design Competition. In: L. Padgham, D. Parkes, J. P. Mueller and S. Parsons (Editors): *Proceedings of the Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*. Estoril, Portugal, May 2008. [9]
- K. Cai, J. Niu and S. Parsons [2008]: On the economic effects of competition between double auction markets. *Proceedings of the 10th Workshop on Agent-Mediated Electronic Commerce (AMEC 2008)*. Estoril, Portugal, May 2008. [2]
- M. L. Gruman and M. Narayana [2008]: Applications of classifying bidding strategies for the CAT Tournament. *Proceedings of the International Trading Agent Design and Analysis Workshop (TADA 2008)*. Chicago, IL, USA, July 2008. [8]
- A. Petric, V. Podobnik, A. Grguric and M. Zemljic [2008]: Designing an effective e-market: an overview of the CAT agent. *Proceedings of the International Trading Agent Design and Analysis Workshop (TADA 2008)*. Chicago, IL, USA, July 2008. [13]

- P. Vytelingum, I. A. Vetsikas, B. Shi and N. R. Jennings [2008]: IAMwildCat: The winning strategy for the TAC market design competition. *Proceedings of the 18th European Conference on AI (ECAI-2008)*. Patras, Greece, July 2008. [19]
- J. Niu, K. Cai, P. McBurney and S. Parsons [2008]: An analysis of entries in the first TAC Market Design Competition. In: L. Jain, B. Faltings, M. Gini and T. Terano (Editors): *IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2008)*. Sydney, Australia, 2008. [10]
- J. Sohn, S. Lee and T. Mullen [2009]: Impact of misalignment of trading agent strategy under a multiple market. *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA-2009)*. Boston, MA, USA, May 2009. [15]

7 Results of Past Tournaments

For reference purposes only, the entrants of the 2007 and 2008 CAT Tournaments are listed here, ranked in decreasing order of performance in each respective tournament. For the 2007 CAT Tournament, there were ten entrants:

- IAMwildCAT: University of Southampton, UK (Winner)
- PSUCAT: Pennsylvania State University, USA (Second Place)
- CrocodileAgent: University of Zagreb, Croatia (Third Place)
- jackaroo: University of Western Sydney, Australia
- Havana: Brown University, RI USA
- PersianCat: Shiraz University, Iran
- PhantAgent: Politehnica University of Bucharest, Bucharest, Romania
- Mertacor: Aristotle University of Thessaloniki, Greece
- TacTex: University of Texas at Austin, USA
- MANX: University of Michigan, USA.

For the 2008 CAT Tournament, there were eleven entrants:

- PersianCat: Shiraz University, Iran (Winner)
- MANX: University of Michigan, USA (Second Place)
- jackaroo: University of Western Sydney, Australia (Third Place)
- PSUCAT: Pennsylvania State University, USA

- Mertacor: Aristotle University of Thessaloniki, Greece
- IAMwildCAT: University of Southampton, UK
- DOG: Kings College, London, UK (Best New Entrant)
- CrocodileAgent: University of Zagreb, Croatia
- MyFuzzy: Individual entry, Cairo, Egypt
- BazarganZebel: Individual entry, Iran
- Hairball: Individual entry, USA.

References

- [1] K. Cai, E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps. Overview of CAT: A Market Design Competition. Technical Report ULCS-09-005, Department of Computer Science, University of Liverpool, Liverpool, UK, 2009. Version 2.0.
- [2] K. Cai, J. Niu, and S. Parsons. On the economic effects of competition between double auction markets. In O. Shehory and H. La Poutre, editors, *Proceedings of the 10th Workshop on Agent-Mediated Electronic Commerce (AMEC 2008)*, Estoril, Portugal, 2008. AAMAS 2008.
- [3] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England, 1997.
- [4] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In B. Nebel, editor, *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI 2001)*, pages 1169–1176, Seattle, WA, USA, 2001. Morgan Kaufmann.
- [5] I. Erev and A. E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–881, 1998.
- [6] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behaviour*, 22:1–29, 1998.
- [7] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [8] M. L. Gruman and M. Narayana. Applications of classifying bidding strategies for the CAT Tournament. In W. Ketter, A. Rogers, N. Sadeh, and W. Walsh, editors, *Proceedings of the International Trading Agent Design and Analysis Workshop (TADA 2008)*, Chicago, IL, USA, 2008.

- [9] J. Niu, K. Cai, E. Gerding, P. McBurney, and S. Parsons. Characterizing effective auction mechanisms: Insights from the 2007 TAC Market Design Competition. In L. Padgham, D. Parkes, J. P. Mueller, and S. Parsons, editors, *Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*, New York, NY, USA, 2008. ACM Press.
- [10] J. Niu, K. Cai, P. McBurney, and S. Parsons. An analysis of entries in the first TAC Market Design Competition. In L. Jain, B. Faltings, M. Gini, and T. Terano, editors, *IEEE-WIC-ACM International Conference on Intelligent Agent Technology (IAT 2008)*, Sydney, Australia, December 2008. IEEE-WIC-ACM.
- [11] J. Niu, K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, S. Phelps, and D. Shield. JCAT: A platform for the TAC Market Design Competition. In L. Padgham, D. Parkes, J. P. Mueller, and S. Parsons, editors, *Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*, New York, NY, USA, 2008. ACM Press.
- [12] J. Niu, A. Mmoloke, P. McBurney, and S. Parsons. CATP: A Communication Protocol for CAT Games. Technical Report ULCS-09-012, Department of Computer Science, University of Liverpool, Liverpool, UK, 2009. Version 2.0.
- [13] A. Petric, V. Podobnik, A. Grguric, and M. Zemljic. Designing an effective e-market: an overview of the CAT agent. In W. Ketter, A. Rogers, N. Sadeh, and W. Walsh, editors, *Proceedings of the International Trading Agent Design and Analysis Workshop (TADA 2008)*, Chicago, IL, USA, 2008.
- [14] S. Phelps. *Evolutionary Mechanism Design*. Phd thesis, Department of Computer Science, University of Liverpool, Liverpool, UK, 2008.
- [15] J. Sohn, S. Lee, and T. Mullen. Impact of misalignment of trading agent strategy under a multiple market. In S. Das and M. Ostrovsky, editors, *Proceedings of the First Conference on Auctions, Market Mechanisms and Their Applications (AMMA 2009)*, Boston, MA, USA, 2009. *Forthcoming*.
- [16] R. S. Sutton and A. G. Barto. *Reinforcement Learning*. MIT Press, Cambridge, MA, USA, 1998.
- [17] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the Third ACM Conference on Electronic Commerce (EC 2001)*, pages 206–209. ACM, 2001.
- [18] P. Vytelingum. *The Structure and Behaviour of the Continuous Double Auction*. PhD thesis, School of Electronics and Computer Science, University of Southampton, 2006.
- [19] P. Vytelingum, I. A. Vetsikas, B. Shi, and N. R. Jennings. IAMwildCAT: The winning strategy for the TAC Market Design Competition. In M. Ghallab and C. D. Spyropoulos, editors, *Proceedings of 18th European Conference on Artificial Intelligence (ECAI 2008)*, pages 428–434, Patras, Greece, 2008.

- [20] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent games. In P. Gmytrasiewicz and S. Parsons, editors, *Proceedings of the AAAI Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT 2002)*, Edmonton, Canada, 2002.
- [21] B. Widrow and M. E. Hoff. Adaptive switching circuits. *IRE WESCON Convention Record*, 4:96–104, 1960.