

How specialised are specialists? Generalisation properties of entries from the 2008 TAC Market Design Competition.

Edward R. Robinson

School of Computer Science
University of Birmingham (UK)
e.robinson@cs.bham.ac.uk

Peter McBurney

Department of Computer Science
University of Liverpool (UK)
mcburney@liverpool.ac.uk

Xin Yao

School of Computer Science
University of Birmingham (UK)
x.yao@cs.bham.ac.uk

Abstract

Unlike the classic Trading Agent competition (TAC), where participants enter trading strategies into a market, the TAC Market Design Competition (CAT) allows participants to create rules for their own double auction market and set fees for traders, which they embody in agents known as *specialists*. Although the generalisation properties of traders when the specialist (i.e., the market mechanism) is fixed have been assessed, generalisation properties of specialists have not. It is unclear whether and how a specialist might (intentionally or unintentionally) favour certain trading strategies. We present an empirical analysis of specialists' generalisation abilities in various trading environments. Our results show that specialists can be sensitive to a number of factors, including the other trading and specialist strategies in the environment.

1. Introduction

The growth of e-commerce has led to increased attention to economic markets from within computer science. Within economics, the discipline of *Mechanism Design* uses methods from mathematical game theory, behavioural economics, and computer simulation to aid in the design and analysis of marketplaces, and of the strategies used by traders within these markets. One common type of market is the *double auction*, in which multiple potential buyers and sellers place (respectively) bids and asks, seeking to engage in purchase transactions over some product or commodity. Double auctions are in use by most of the stock and commodity exchanges around the world, although often with different rules. In a global economy, stock exchanges compete against each other for trading business, and, increasingly, against new online markets not tied to traditional physical exchanges.

In such a competitive environment, the precise rules adopted by a marketplace may have important consequences—in attracting (or not) traders to the market, in rewarding (or not) particular trading strategies, and in facilitating (or not) the matching of shouts (bids and asks), and the execution of trades in the market. Thus, a detailed understanding of the different potential rules for double auction markets and their impacts is important, particularly if we seek the eventual automation of the design of market mechanisms. However, the mathematical analysis of the double auction is intractable, and so computer or human

simulations are currently needed to gain this understanding.

With this in mind, a new research tournament was launched in 2007 to promote research into automated mechanism design: the *Trading Agent Competition Market Design (or CAT) Tournament* (Gerding et al. 2007). The CAT Tournament comprises a series of artificial parallel markets, designed to mirror the competition between global stock markets. These parallel markets, called *specialists*, are created by entrants to the Tournament, and they compete with one another to attract and retain traders, who are potential buyers and sellers of some abstract commodity. The traders are software agents created and operated by the organisers of the CAT Tournament, in a reversal of the usual Trading Agent Competition structure.

In this paper, we describe simulation analyses undertaken using the CAT Tournament platform JCAT¹ with some of the specialists entered into the 2008 Tournament. Our primary research goal is to better understand the characteristics of the mechanisms used by the specialists, particularly in relation to the contexts in which they trade, in order to design more robust mechanisms. Thus, we seek to generalise specialist capabilities and strengths. Our contribution is to demonstrate that the specialists in the 2008 CAT Tournament are not robust against changes in the trader mix, in the competitor mix and in the scoring period. Changing each of these factors leads to some changes in the tournament ranks and/or the game scores achieved by the specialists. The market mechanisms employed by the specialists may thus be seen not to generalise, and so research will be needed to make these mechanisms more robust.

The paper is organised as follows. Section 2 presents a brief summary of the CAT Tournament, and Section 3 presents our research findings in detail. Section 4 presents our conclusions and proposals for future work.

2. CAT Tournament

The organization and structure of the TAC Market Design (CAT) Tournament is given in the game documents (Gerding et al. 2007). Here we briefly mention the most important aspects. A CAT game takes place over a number of simulation *trading-days*, each of which consists of a number of *rounds*. Each round lasts a number of *ticks*, measured in millisec-

¹<http://jcat.sourceforge.net/>

onds. The game uses a client server architecture, with the CAT server controlling the progression of the game. CAT clients are either traders (potential buyers or sellers) or specialists (aka markets). All communication between traders and specialists is via the CAT server.

In the standard CAT installation, four different trader strategies are provided. *Zero Intelligence – Constrained (ZIC)* traders (Gode and Sunder 1993) essentially place random bids and asks, within constraints. These traders ignore both the current state of the market and the market history. *Zero Intelligence Plus (ZIP)* traders (Cliff 1997) are modified versions of *ZIC* traders that seek to remain in profit in competition with other traders, using some market history. *RE* traders (Erev and Roth 1998) use a learning algorithm based on a model of human learning, with the most recent surplus or loss guiding the trader’s shouting strategy one step ahead. Finally, *GD* traders (Gjerstad and Dickhaut 1998) use past marketplace history of submissions and transactions to generate beliefs over the likelihood of any particular bid or ask being accepted, and use this belief to guide shouting strategies. *ZIC* are the least, and *GD* are the most, sophisticated of these four types. In addition, all types of traders in the standard CAT installation use an n-armed bandit strategy (Robbins 1952) for selecting which specialist to register with on each new trading day.

Specialists have freedom to set market rules in six broad policy areas: *Charging policy*: What charges and fees are imposed by the specialist on traders? *Quote policy*: What limitations, such as a lower bound for bids, does the specialist impose on shouts by traders? *Shout accepting policy*: When does a specialist accept a shout made by a trader? *Matching policy*: How does the specialist match bids and asks made by traders? *Pricing policy*: What are the transaction prices for matched bid-ask pairs? *Clearing policy*: When does the specialist clear the market and execute transactions between matched bids and asks? Following the 2007 CAT Tournament, (Niu et al. 2008a) undertook a series of simulations to infer the policies of the specialists entered in the game. In addition, the research teams behind two of the 2007 CAT specialists have written about their strategies (Petric et al. 2008; Vytelingum et al. 2008).

In the CAT Tournament, specialists know that each trader is one of the four types, but they do not know the overall proportions of each type. Accordingly, the design of a specialist seeking to win the game cannot be optimised for only a subset of trading strategies. In addition, the scoring metric used by the game is multi-dimensional. Games are scored using an unweighted average of three criteria: the proportion of traders attracted to the specialist each day (market share); the proportion of accepted shouts which are matched (transaction success rate); and the share of profits made by the specialist. As with trader types, this multi-dimensionality creates challenges for the optimal design of specialists, since these criteria may conflict. A game-winning strategy may focus on scoring highly on different criteria at different times in the life-cycle of a game, or against different types of traders.

Given this game structure, it is easy to see that some specialists may perform better with traders of a particular type,

and/or against competing specialists using particular policies. Because the actual CAT Tournaments are only conducted over a limited number of games (typically, three), the performance of a specialist in the Tournament is not necessarily a good guide to that specialist’s general ability, i.e., with other trader mixes, or in competition against different specialists. Niu *et al.* (2008b), for example, have shown that some of the well-performing 2007 CAT specialists have weaknesses in other situations. Specialists may, therefore, be considered *brittle* (or obversely, *robust*), if their performance greatly depends (or does not) on the competitive and trader context.

In this paper, we explore via simulation the extent to which the performance of specialists in the 2008 Tournament may generalise across competitive and trader contexts, with the aim of creating more robust specialists.

3. An Empirical Evaluation of the Generalisation Ability of the Entries

General Experimental Setup

Each JCAT simulation consists of a single tournament that runs for a number of trading days, with, in our experiments, 10 rounds per day and 500ms per round. All experiments were carried out with both the JCAT server and all clients situated on the same local machine. The trading population size was set at 400 traders, filled with traders taken from the four types described in Section 2. Buyers and sellers were split as evenly as possible in the different trader sub-populations. Each of the traders’ private values for goods were decided by drawing from a uniform distribution, a value between 50 and 150. In order to achieve statistically significant results for each tournament variation, each variation is repeated 15 times, using the same configuration.

The specialist agents we used in our experiments were downloaded in a pre-compiled form from the TAC Agent Repository². Specifically, we used entries from the 2008 competition, except when this was not possible³. The following specialists were included in our experiments: *CrocodileAgent*, *DOG*, *iAmWildCat 2008*, *Mertacor1*, *Mertacor2*, *PSUCAT*, *PersianCAT* and *jackaroo*. In the rest of this section we will often refer to them as *CR*, *DO*, *IA*, *M1*, *M2*, *PS*, *PC* and *JA* respectively.

As defined in Section 2, each day every specialist receives a score based on three criteria. A specialist’s tournament score is the sum of its daily scores for all scoring days. Specialists are ranked in descending order of their total scores, with the specialist in rank 1 declared the winner of that tournament. To show that some specialists’ performances can be sensitive to a number of factors, and in some cases generalise poorly, we measure differences in specialist performance across different tournament variations. Specifically, we look for two differences: firstly we measure the *qual-*

²<http://www.sics.se/tac/showagents.php>

³We were unable to use the following specialists: *BazarganZebel* (binary corrupted); *Hairball* (not submitted to repository); *MANX* (requires unavailable libraries); *MyFuzzy* (unable to disable X output).

itative impact tournament variations have on each specialist's performance, which we achieve by comparing rankings of specialists' mean scores for different (comparable) tournament configurations. Secondly, we measure the *performance impact*, i.e., the change that tournament variations have on each specialist's score. It is clear that since the differences introduced between tournament variations—e.g., the proportions of trader types in the trading population—itsself contributes to the performance of the specialists, one cannot simply compare the mean scores of specialists over tournament variations and confidently state how the specialists are able to generalise between the two cases.

Our response to this challenge is to define a new statistic to measure the performance of one specialist relative to others, across a diversity of tournament variations. Our statistic, which we call the normalised performance delta of a specialist, denoted $\hat{\delta}_i$, provides a metric for analysing how a given tournament configuration affects the performance of the specialist. To calculate this statistic, we first calculate for each specialist i , the normalised mean score $\hat{\mu}_i$.

$$\hat{\mu}_i = \frac{\mu_i}{\sum_{j=1}^m \mu_j} \quad (1)$$

For a single specialist i , given two normalised scores $\hat{\mu}_i^x$ and $\hat{\mu}_i^y$ from two tournament variations x and y , we can calculate the absolute difference d_i^{xy} between the two scores.

$$d_i^{xy} = |\hat{\mu}_i^x - \hat{\mu}_i^y| \quad (2)$$

d_i^{xy} gives us an understanding of how, with respect to other specialists in the tournament, a specialist i 's performance has changed from one tournament variation to the next. Finally, for each specialist i we calculate the normalised difference value $\hat{\delta}_i$.

$$\hat{\delta}_i = \frac{d_i^{xy}}{\sum_{j=1}^m d_j^{xy}} \quad (3)$$

In order to ascertain some statistical significance to specialist mean score values generated from multiple tournament runs, two-tailed paired t-tests of equality of means were performed on certain pairs of specialists, in order to attempt to identify whether the reported rankings were distinct. In such cases both the *t-value* and *p-value* (using $n - 1$ df.) are reported. These tests assume normality in the distribution of the underlying random variables. However, even if the variables are from other distributions, the tests are still approximately correct (Moses 1986) (p.197–198); (Devroye 1986) (p.362).

Over-fitting to trading population

In the following set of results, we show that some specialists' performances are sensitive to different mixes of trader types in the trader population, and as such some specialists may be over-fitted to specific types or mixes of traders. For this set of experiments, all of the specialists available

were used. We ran a number of tournament variations, each of which consisted of differing proportions of individual trader types. In these tournaments, all 500 trading days were counted as scoring days. Overall, we found that several of the specialists' final rankings were affected by variations, particularly `jackaroo`, `Mertacor1` and `Mertacor2`.

Table 1 shows the results of two tournament variations, which we refer to as 'justGD' and 'noGD'. The justGD variation consisted of a trading population made up of entirely GD traders, with equal buyers and sellers. In the noGD variation, the trading population was composed of 133 ZIP, 133 RE and 134 ZIC traders. There were an equal number of ZIC buyers and sellers, while the RE (ZIP) sub-populations had 67 (66) sellers respectively.

Spe	Just GD Traders		No GD Traders		Rank	$\hat{\delta}$
	μ	σ	μ	σ		
PC	232.2	5.86	271.1	5.51	1, 1	0.355
M1	230.9	6.55	193.9	3.17	2, 4	0.364
JA	218.9	6.06	213.5	2.65	3, 3	0.064
M2	207.4	5.09	215.7	6.05	4, 2	0.067
IA	165.8	1.45	165.1	5.15	5, 6	0.016
DO	164.4	2.24	173.7	2.48	6, 5	0.078
CR	24.9	17.76	19.1	8.28	7, 7	0.057
PS	16.2	0.66	16.3	0.43	8, 8	0.001

Table 1: Mean, standard deviation, rank and $\hat{\delta}$ values for a set of tournaments with just GD traders and a set of tournaments with no GD traders.

For a typical tournament variation, we found that in each of the n repetitions, scores, and thus rankings, were quite similar, leading to low σ values. It is extremely unlikely that scores would ever be identical over all runs due to the stochastic nature of the JCAT environment. Table 1 highlights the fact that the overall rankings for the two tournament variations were different, most notably with changes in the middle and lower portions. Qualitatively, we found of particular interest was the change in rank between M1, M2 and JA. In the justGD case, M1 was rank 2 and M2 rank 4, while in the noGD case the ranks were swapped to 4 and 2 respectively. In the justGD case, a paired t-test showed that the average scores of M1 and M2 were significantly different, with a *t-value* of 9.36 and a *p-value* < 0.0001 . The average scores of M1 and M2 were 230.9 and 207.4 respectively. In the noGD case, the t-test resulted in a *t-value* of 14.04 and a *p-value* < 0.0001 . Average scores in the noGD case were 193.9 for M1 and 215.7 for M2.

Further, in the justGD case we found that M1 and JA had ranks of 2 and 3 respectively, while in the noGD case they had ranks of 4 and 3. In the justGD case, for M1 and JA, a t-test resulted in a *t-value* of 19.09 and a *p-value* < 0.0001 , with mean scores of 230.9 for M1 and 218.9 for JA. In the noGD case, a t-test reported a *t-value* of 4.20 and a *p-value* < 0.0001 , with mean scores of 193.9 for M1 and 213.5 for JA.

Finally, we note that even a simple change of the trading

population, i.e., justGD, can make a previous winner, PC, lose its winning edge. Statistically, PC is not the clear winner in the justGD case. A t-test of equality of means between PC and M1 in the justGD case showed a t -value of 0.46 and a p -value of 0.65, with mean scores of 232.2 for PC and 230.9 for M1. A counterpart to this situation is the noGD case, where PC clearly outperformed M1. Here the t -value was 38.31 and p -value < 0.0001 , with mean scores of 271.1 for PC and 193.9 for M1. This highlights a situation where either PC or M1 are particularly sensitive to the proportions of GD traders in the population. The δ values for PC (0.355) and M1 (0.364) were considerably larger than those of the other specialists, showing a disproportionate change in performance over the two cases for both specialists.

In Table 2 we show the results of two different tournament variations: ‘justZIC’ and ‘noZIC’. In the justZIC variation we used all specialists and a trader population containing only ZIC traders (equal number of buyers and sellers). The noZIC variation used all specialists, but the trader population contained 133 GD, 133 RE and 134 ZIP traders, with an equal numbers of ZIP buyers and sellers and 67 (66) GD (RE) sellers respectively.

Spe	Just ZIC Traders		No ZIC Traders		Rank	δ
	μ	σ	μ	σ		
PC	276.2	11.47	244.1	4.53	1, 1	0.419
JA	218.6	4.71	214.6	2.49	2, 3	0.071
M2	213.9	8.20	224.7	5.77	3, 2	0.109
M1	192.1	2.72	207.0	3.76	4, 4	0.160
DO	170.7	2.62	180.0	2.32	5, 5	0.095
IA	161.6	5.15	170.7	4.09	6, 6	0.094
PS	16.9	0.46	16.2	0.47	7, 8	0.010
CR	16.8	5.07	20.4	7.90	8, 7	0.041

Table 2: Mean, standard deviation, rank and δ values for a set of tournaments with just ZIC traders and a set of tournaments with no ZIC traders.

We found another situation when it is not statistically clear that the rankings between two specialists are the same across the two tournament variations, indicating that there were generalisation problems. In the noZIC case, we found that M2 (rank 2) outperformed JA (rank 3). A paired t-test of equality of means found the scores statistically different, with a t -value of 5.83 and a p -value < 0.0001 . The mean scores were 224.7 for M2 and 214.6 for JA. However, in the justZIC case we found that the mean scores between the two specialists, and thus the rankings, were not statistically distinct. A paired t-test results in a t -value of 1.75 and a p -value of 0.10. The mean scores were 213.9 for M2 and 218.6 for JA.

Table 2 also highlights a clear example of the performance impact that changes in the trader population can have on a specialist. The mean score for PC in the justZIC case was 276.2, yet this dropped 8.83% to 244.1 in the noZIC case. Of course, the makeup of the trader population can have a significant impact on the scores of specialists, but in these

two cases we see that other specialists’ scores did not vary proportionally as much as PC’s. By considering the performances changes of the other specialists, PC’s normalised delta value δ was 0.419, which was considerably higher than the others’.

We highlight in Table 3 other qualitative impacts that changes in the trader population had on specialists. In these tournament variations we considered the cases ‘justZIP’ and ‘noZIP’. In the justZIP variation we used all specialists and a trading population consisting of only ZIC traders (equal number of buyers and sellers). In the noZIP variation all specialists were used but the trading population contained 133 GD, 133 RE and 134 ZIC traders, with an equal number of ZIC buyers and sellers, and 67 (66) GD (RE) sellers respectively.

Spe	Just ZIP Traders		No ZIP Traders		Rank	δ
	μ	σ	μ	σ		
PC	255.7	12.18	255.9	5.17	1, 1	0.051
JA	231.3	5.02	211.5	4.37	2, 3	0.281
M1	208.0	5.86	199.0	4.19	3, 4	0.150
M2	189.0	5	221.9	6.94	4, 2	0.347
IA	169.7	3.16	173.5	4.34	5, 6	0.008
DO	164.1	4.86	179.7	2.4	6, 5	0.147
CR	17.0	6.35	16.1	5.04	7, 8	0.014
PS	16.2	0.65	16.3	0.4	8, 7	0.003

Table 3: Mean, standard deviation, rank and δ values for a set of tournaments with just ZIP traders and a set of tournaments with no ZIP traders.

In the justZIP case we found that M1 (rank 3) outperformed M2 (rank 4). A t-test resulted in a t -value of 10.35 and a p -value < 0.0001 . The mean scores for M1 and M2 were 208.0 and 189.0 respectively. However, in the noZIP case we again see a different outcome, with the ranks changed to 2 for M2 and 4 for M1. In this case, a t-test resulted in a t -value of 9.56 and a p -value < 0.0001 , with means of 221.9 for M2 and 199.0 for M1.

Over-fitting to other specialists

In this set of results, we show that when the proportions of traders in the trader population remain fixed, some specialists’ performances are sensitive to the presence of other specialists in the marketplace. For this set of experiments, we used a fixed trader population containing an equal mix of GD and ZIC traders. Since ZIC traders are the least, and GD the most, sophisticated trader types, using this mix may offer the most diverse trading environment. In these tournaments, all 500 trading days were counted as scoring days.

In Table 4 we see the effects that removing the bottom five specialists had on the remaining three. We particularly note the effect that lower-ranked specialists had on the performance of JA and M2. For example, when all specialists were present, JA outperformed M2. A t-test of equality of means showed a t -value of 4.12 with a p -value of 0.0010.

Spe	All Specialists		Just PC, JA & M2		Rank	δ
	μ	σ	μ	σ		
PC	267.34	5.99	299.74	6.96	1, 1	0.361
JA	217.46	4.44	243.5	3.51	2, 3	0.293
M2	206.93	6.44	256.72	5.82	3, 2	0.347
M1	198.31	2.79	–	–	4, –	–
DO	174.45	2.81	–	–	5, –	–
IA	170.59	3.51	–	–	6, –	–
CR	18.86	5.21	–	–	7, –	–
PS	16.08	0.29	–	–	8, –	–

Table 4: Mean, standard deviation, rank and δ values for a set of tournaments with all specialists and a set of tournaments with just PC, JA and M2. The trader population consisted of an equal mix of GD and ZIC traders.

Mean scores for JA and M2 were 217.5 and 206.9 respectively. Alternatively, when the lower five specialists were removed, and the remaining three compete with the same traders, we found that the rankings of JA and M2 switched. In that case, a t-test resulted in a *t-value* of 8.30 with a *p-value* < 0.0001. Mean scores for JA and M2 were 243.5 and 256.72 respectively.

In Table 5 we can observe the effect that *removing* the top three specialists had on the remaining bottom five. In a qualitative context, IA did significantly better than DO when the top three specialists were not present, with a t-test revealing a *t-value* of 5.80 and a *p-value* < 0.0001. The average score for IA was 215.0, while DO scored 208.7. However, when the top three specialists were present, DO ranked higher than IA (*t-value* = 3.31, *p-value* = 0.0051). The average score for IA was 170.6, while DO scored 174.5.

With respect to the performance impact that the two tournament variations had on specialists, it is clear that the inclusion (or likewise, exclusion) of the three specialists PC, JA and M2 clearly affected the performance of M1 more than any of the remaining four. When the top three specialists were introduced, all of the other specialists' scores were lower, however a high normalised delta value of 0.525 for M1 indicated it was considerably more sensitive to their presence.

Over-fitting to scoring period

Our final results show that some specialists' performances were affected by the choice of trading days used as scoring days. In Sections 3and 3we maintained a fixed scoring period of 500 days and varied either the specialist or trader populations. In these results we take simulations from the previous sections, and adjust the scoring period from which the specialists' scores are generated. Specifically, we measured performance over the scoring day period 1–100 and 100–380.

Table 6 shows the effect changing the scoring period had on the performance of the specialists. In these two variations we used all specialists and a trading population containing just ZIC traders, with an equal number of buyers and

Spec	All Specialists		No PC, JA or M2		Rank	δ
	μ	σ	μ	σ		
PC	267.3	5.99	–	–	1, –	–
JA	217.5	4.44	–	–	2, –	–
M2	206.9	6.44	–	–	3, –	–
M1	198.3	2.79	376.3	4.17	4, 1	0.525
DO	174.5	2.81	208.7	1.5	5, 3	0.198
IA	170.6	3.51	215.0	3.86	6, 2	0.217
CR	18.9	5.21	32.9	10.22	7, 7	0.044
PS	16.1	0.29	18.2	0.32	8, 8	0.016

Table 5: Mean, standard deviation, rank and δ values for a set of tournaments with all specialists and a set of tournaments with all specialists except PC, JA and M2. The trader population consisted of an equal mix of GD and ZIC traders.

sellors. For all specialists, as you would expect, scores for days 1–100 were much lower than the 100–380 tournament because the 1–100 scoring period is less than a third shorter. We found that many of the specialists' ranks changed between the two cases, and highlight specifically JA and M2. If specialists were ranked according to their score from the earlier period, we found that JA had a rank of 3, while M2 had a rank of 2. A t-test tells us these means were statistically different; the t-test resulting in a *t-value* of 10.42 and a *p-value* < 0.0001, with mean scores of 46.0 for M2 and 39.8 for JA. Alternatively, if specialists were ranked using the scoring period 100–380 the ranks of M2 and JA switched. Again, in this case we found a t-test of equality of means revealed that the scores are statistically distinct, resulting in a *t-value* of 3.34 and a *p-value* of 0.0048.

Spec	Day 1–100		Day 100–380		Rank	δ
	μ	σ	μ	σ		
PC	50.8	1.94	158.6	7.34	1, 1	0.190
M2	46.0	1.86	118.8	5.19	2, 3	0.025
JA	39.8	0.81	124.4	3.34	3, 2	0.150
M1	39.5	1.01	107.7	1.80	4, 4	0.056
IA	35.0	1.19	89.1	3.72	5, 6	0.011
DO	34.0	1.09	96.4	1.46	6, 5	0.069
PS	16.9	0.46	0.00	0.00	7, 8	0.251
CR	16.8	5.07	0.00	0.00	8, 7	0.249

Table 6: Mean, standard deviation, rank and δ values for a set of tournaments with scoring days 1–100 and a set of tournaments with scoring days 100–380. The trader population consisted of just ZIC traders.

Discussion

We have presented in this paper a methodology that allows us to evaluate and compare agents' generalisation abilities

systematically and quantitatively. We believe that this approach is the first to allow such properties to be measured—especially in a coevolutionary context. Further, there are several reasons why such an approach is important. Firstly, without such an approach it is hard to know the true performance of any particular agent strategy. This is because without a methodology such as ours one cannot see how well a strategy performs against unknown or previously unseen competing strategies, i.e., how well a given strategy generalises and thus its ‘robustness’; in our results we have shown that some seemingly strong strategies can have weaknesses when competing with certain other strategies. This is interesting because it is clear from the literature that often ‘best strategies’ or ‘best results’ are published, and without any further elaboration on their generalisation ability these results can be very misleading. Secondly, although it may seem obvious that the performance of any given strategy—whether in the CAT game or some other competitive multi-agent system—will depend on the strategies being used by other agents, it is *never* clear how strong this dependency is, or the resulting quantitative values such a dependency provides. Finally, we note that our approach’s significance can be demonstrated by emphasising that our results concur with the assumption that the performance of agents will depend on the strategies being used by their competitors, thus it is a sensible one.

Previously, Vetsikas and Selman (2003) presented a methodology for deciding on the best strategy that their agent, *WhiteBear*, should use in the TAC classic game⁴. The authors decompose the overall bidding-agent problem into a smaller sub-problems, e.g., decide on the quantity of a good to buy, and then determine the best strategy given those quantities. One of the reasons that TAC is so hard for an agent is that there are multiple concurrent auctions taking place using a variety of rules. Their approach was to tackle each auction independently by generating a set of boundary strategies for each auction or good (an example of a boundary strategy might be ‘bid-low’ or ‘bid-high’). From these partial strategies they then generated intermediate strategy variants, which lay between the boundary strategies. Multiple experiments were carried out to test different strategies in the presence of differing numbers of other strategies. Although this approach is sensible when trying to decide on the best strategy from an initial set of possible strategies, it is unclear how well such a best strategy would fair against unseen strategies, i.e., other competitors in the trading competition. Indeed, unlike the approach in this paper, theirs provides no such mechanism to test the robustness of a strategy in the presence of unseen competitors.

Wellman et al. (2005) explore the idea of using a reduced strategy space to allow them to evaluate potential strategies taken from their TAC agent, *Walverine*. They hypothesise that although a strategy’s performance depends on the strategies being used by other competitors, it may be relatively insensitive to the number of competitors using identical competing strategies. To that end Wellman et al. introduce a methodology that explores a reduced number of pos-

sible profiles—where a profile defines the type of strategies in use in a population and their frequency. They find that by reducing the strategy space, they can efficiently search the remaining space to find the strategies that can outperform the others. Although such an approach is very appropriate when considering a fixed initial space of potential strategies (in their case 35) as with the methodology in Vetsikas and Selman it is hard to say how well the found strategies would generalise to previously unseen ones, such as those faced in competitions.

Given the two previously discussed methodologies there are several reasons why our approach is novel. Firstly, the CAT game is more complicated than the original TAC game because there are two interacting populations—specialists and traders, both of whom contain (or could contain) members able to learn and/or evolve. Thus our approach is the first to look at a situation where there are *coevolving* populations of competing agents, and offer insights into the robustness of the strategies in those populations. Secondly, to even define what ‘robustness’ of a trading or specialist strategy means is not straightforward in such a coevolutionary context, let alone to actually measure such a property in a systematic and rigorous way. Our approach, unlike the previously discussed ones, is the first to look at how well such designed strategies generalise to previously unseen strategies, using the CAT competition as a case-study. Indeed, we believe that our approach can be used for analysing the generalisation properties of strategies in not only other agent-based competitions, but in any agent-based simulation models within an economic domain.

4. Conclusions and Future Work

To our knowledge, the generalisation property of specialists (i.e., market mechanisms) has not been studied in the literature. Traditionally, economic mechanism design theory has dealt with the generalisation issue by seeking mechanisms which are incentive-compatible (i.e., which encourage truth-telling) and by assuming that all traders are always rational and self-interested. An incentive-compatible mechanism with rational, self-interested traders should generalise across trader strategies and types. However, because computer software is always resource-constrained and bug-prone, the designers of computational mechanisms cannot assume that traders always act rationally or in their own self-interest. Moreover, economic mechanism design theory has not considered competition between mechanisms, and so has not explored generalisability of mechanisms over competing mechanisms. Within computer science, Niu *et al.* (2007) considers the choices automated traders will make when faced with several competing online marketplaces, showing that although traders typically gravitate towards low-charging markets, such markets may lose their dominance when not all traders are experienced or do not all have access to full information.

Thus, the research reported here has tried to explore the generalisation properties of market mechanisms, using the 2008 CAT Tournament specialists as the basis. It is unclear whether and how a market mechanism might (intentionally

⁴<http://www.sics.se/tac/page.php?id=3>

or unintentionally) favour certain trading strategies, or facilitate or inhibit other competing market mechanisms. A specialist which performs well under one particular tournament setup may not perform well if the setup changes slightly. Therefore, it is essential to study and understand any hidden bias that a specialist might have built into it.

In order to achieve this we ran many JCAT simulations using a variety of configurations, including changing both the trader and specialists populations as well as changing the period used to generate specialist scores. This paper shows for the first time how changes in the competition configuration can have an impact on specialists' performances in both a qualitative and quantitative context.

Our results showed that specialists can be sensitive (and specialised) to a number of factors in the competition, including the trading strategies used in assessing specialists, other specialists that one competes against and the scoring period. Such results indicate that an appropriate evaluation of such a competition (and other similar ones) would need a theoretically sound framework, which can measure specialists' generalisation abilities quantitatively. They also point out the importance in analysing the relationship between a winning specialist and the particular competition setup used, so that insights into what makes a specialist better/worse can be gained.

Although we have studied three major issues that have a significant impact on specialists' performance from the viewpoint of generalisation, there are other issues to be considered, e.g., the performance metric used. It is very interesting to study as the next step the potential trade-offs between market share and total profit. A multi-objective approach could be used to study these trade-offs. In terms of a theoretical framework for measuring specialists' generalisation quantitatively, we will investigate the possibility of adopting the one for measuring strategies' generalisation ability (Chong, Tino, and Yao 2008).

Acknowledgments

This work is supported by the Market Based Control of Complex Computational Systems project, EPSRC grants: GR/T10671/01 & GR/T10657/01. The authors are grateful for the financial support received. We would also like to acknowledge insightful discussions and constructive feedback from Peter Lewis, Michael Wooldridge, Tomasz Michalak, Andrew Dowell and Jinzhong Niu.

References

Chong, S. Y.; Tino, P.; and Yao, X. 2008. Measuring generalization performance in co-evolutionary learning. *IEEE Transactions on Evolutionary Computation* 12(4):479–505.

Cliff, D. 1997. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, Hewlett-Packard Research Laboratories, Bristol, UK.

Devroye, L. 1986. *Non-Uniform Random Variate Generation*. New York, NY 10010: Springer-Verlag.

Erev, I., and Roth, A. E. 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88(4):848–881.

Gerding, E.; McBurney, P.; Niu, J.; Parsons, S.; and Phelps, S. 2007. Overview of CAT: A market design competition. Technical Report ULCS-07-006 Version 1.1, Department of Computer Science, University of Liverpool, Liverpool, UK.

Gjerstad, S., and Dickhaut, J. 1998. Price formation in double auctions. *Games and Economic Behavior* 22(1):1–29.

Gode, D. K., 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):119–137.

Moses, L. E. 1986. *Think and Explain with Statistics*. Reading, MA: Addison-Wesley Publ. Co.

Niu, J.; Cai, K.; Parsons, S.; and Sklar, E. 2007. Some preliminary results on competition between markets for automated traders. In Collins, J., ed., *Proceedings of the AAAI 2007 Workshop on Trading Agent Design and Analysis (TADA 2007)*.

Niu, J.; Cai, K.; Gerding, E.; McBurney, P.; and Parsons, S. 2008a. Characterizing effective auction mechanisms: Insights from the 2007 TAC market design competition. In et al., L. P., ed., *Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2008)*. New York, USA: ACM Press.

Niu, J.; Cai, K.; McBurney, P.; and Parsons, S. 2008b. An analysis of entries in the first TAC market design competition. In et al., L. J., ed., *IEEE-WIC-ACM International Conference on Intelligent Agent Technology (IAT 2008)*. Sydney, Australia: IEEE-WIC-ACM.

Petric, A.; Podobnik, V.; Grguric, A.; and Zemljic, M. 2008. Designing an effective e-market: an overview of the CAT agent. In Ketter, W., ed., *Proceedings of the AAAI 2008 Workshop on Trading Agent Design and Analysis (TADA 2008)*.

Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 58(5):527–35.

Vetsikas, I. A., and Selman, B. 2003. A principled study of the design tradeoffs for autonomous trading agents. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 473–480. New York, NY, USA: ACM.

Vytelingum, P.; Vetsikas, I. A.; Shi, B.; and Jennings, N. R. 2008. IAMwildCat: The winning strategy for the TAC market design competition. In et al., M. G., ed., *18th European Conference on AI (ECAI-2008)*. Patras, Greece: IOS Press.

Wellman, M.; Reeves, D.; Lochner, K.; Cheng, S.; and Suri, R. 2005. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*. Pittsburgh, Pennsylvania, USA: AAAI.