COMBINATORIAL CLOCK AUCTION OF AIRPORT TIME SLOTS: AN AGENT-BASED ANALYSIS

By

Rui Dong

SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
B.A. IN COMPUTER SCIENCE AND ECONOMICS
AT
HARVARD UNIVERSITY
CAMBRIDGE, MA
APRIL 2005

Contents

Abstract **vi**

1 Introduction **1**

1.1 Related Work 4

1.2 Outline of Thesis 6

2 Combinatorial Auctions **7**

2.1 Past Auction Designs and Considerations 7

2.2 Combinatorial Clock Auction 8

2.2.1 Activity Rule 9

2.2.2 Monotonic Increase in Prices 10

2.2.3 Final Allocation 11

2.2.4 The various pieces of the auction 11

2.3 Application to the LaGuardia Domain 12

2.4 Examples and Discussion 13

2.4.1 Example Case 1 : Simple Example with Efficient Allocation . 14

2.4.2 Example 2 : An Undersell Problem 15

2.4.3 Example 3: Adding Budget Constraints 16

3 Valuations for the LaGuardia Domain **18**

3.1 History of the Problem 18

3.2 Models used for valuating airport slots 19

3.2.1 Budget Constraints . . . 20
3.2.2 Model 1: A Linear Optimization Model . . . 21
3.2.3 Model 2 : Bounded Rationality with Incumbents and Entrants 23

4 Vulnerabilities of the Auction and Possible Trading Strategies 30
4.1 Price Prediction . . . 31
4.1.1 Motivation for Price Prediction . . . 31
4.1.2 Process to Simulate Price Prediction . . . 32
4.1.3 Price Prediction and Budgets . . . 32
4.2 Strategy Class 1 : *'Flexibility Preservation'* . . . 33
4.2.1 Implementation of '*parking*' . . . 34
4.2.2 *OverdemandMore* - Another strategy to maintain flexibility . . 35
4.3 Strategy Class 2: "*Increasing Payoff*" . . . 35
4.3.1 *Shaving* . . . 35
4.3.2 *Demand Reduction* . . . 37

5 Experimental Setup 39

6 Chapter 6 - Price Prediction with Machine Learning 42
6.1 A Staged Approach to Price Prediction . . . 42
6.2 Price Prediction - Machine-Learning Techniques . . . 43
6.2.1 Training Data Format . . . 43
6.3 Results from Price Prediction . . . 44

7 Heuristic Strategy Equilibrium 47
7.1 Generating Payoff Matrices . . . 47
7.2 Evolutionary Search . . . 49

8 Results 51
8.1 Nash Equilibria Over Heuristic Stratgies . . . 51
8.1.1 Result 1 : *Flexibility Preservation* . . . 52

8.1.2 Result 2 : *Increasing Payoff* 55
8.2 Vulnerabilities of the Auction 61

9 Conclusion **65**

Bibliography **68**

Abstract

We provide an agent-based analysis of the currently proposed combinatorial clock auction of LaGuardia airport using results from artificial intelligence and economics. First, we build a framework in using artificial intelligent agents in the analysis of auction designs. We show, both experimentally and theoretically, that existence of budget constraints would lead to inefficient outcomes under the current auction design. Using evolutionary search to look for Nash equilibria in restricted games, we found that many of the observed agent strategies in past auction designs were still possible equilibrium outcomes for the agents competing in the combinatorial clock auction.

Chapter 1

Introduction

The Vickrey auction and the English auction are well known mechanism designs conducive to optimal outcomes to all with straightforward bidding in normal single-item, private value auctions.[32] Yet, in reality, many important markets are for multiple goods with interacting values. In such markets, agents compete for goods that interact with each other: a good may increase or decrease in value for certain agents when purchased along with another good (deemed a complement in the first case, a substitute in the second). Applicable scenario examples include large-scale auctions such as the Federal Communications Commission (FCC) spectrum auctions [18] and also day-to-day situations such as consumers bidding in simultaneous online auctions for matching components (hotel, show tickets, airflights) of a vacation package. With the advances in information technology, more such scenarios arise daily [12].

Yet, an efficient auction design for the multi-item bidding problem is deemed difficult. Indeed, an auction that allows bidders to bid values based on packages instead of items (deemed a *combinatorial*mechanism) faces a tractability problem as the number of packages increases exponentially with the size of goods[21]. Alternatively, research proposals were made for simultaneous ascending auctions, implemented in the FCC spectrum allocation auctions[18].

First run in 1994, the spectrum auctions involved huge sums of money ($617 million were generated in the first sale of 10 licenses in July 1994)[20]. Results from the auctions indicate, however, issues of concern for the efficiency of this design. Experimental results indicate possible incentives for collusion and demand reduction in the auction[3]. Moreover, theoretical work also shows concerns for the exposure problem whereby agents are forced to buy items it does not value due to losing in the auction for the complementing item [20]. In response to these problems, interest centered on integrating combinatorial bidding into the ascending auction, with FCC issuing even a request for comments on the subject[20]. These auctions, deemed *dynamic combinatorial auctions*, combines the price discovery feature of the ascending auction, and combinatorial bidding. So far, implementations of these designs have been rare [12].

In the summer of 2004, the Federal Aviation Agency (henceforth FAA), commissioned a study to better allocate airport landing slots in La Guardia airport while solving the airport congestion problem, opening the possibility of implementing dynamic combinatorial auctions on a large-scale. [1] Analysis of an application of combinatorial auctions onto airport time slots is, however, difficult. First, because of the intractability of the package bidding problem, agents participating in a simulation need to be fairly sophisticated.[2] This makes it difficult to generate experimental data with staged laboratory settings. Second, actual airline executives participating in simulations have strong interests in hiding their true valuation and strategies. In this thesis, using the currently proposed auction mechanism, we will propose instead an artificial intelligent agent-based experimental framework to better understand some of the strategic and economic aspects of using a combinatorial auction, in particular in its application to the La Guardia problem.

In this process, we believe that our paper contributes to the field in two ways: first, by proposing a complete electronic framework for the simulation and analysis

[1] In fall 2004, a proposed combinatorial auction mechanism was laid out by a team of researchers and then presented to airline executives and the FAA through a series of conferences. A mock auction hosted by the National Center of Excellence for Aviation Operations Research (NEXTOR) and based on a similar design was subsequently held in February 24-25, 2005 with industry executives, representatives from the FAA, and academics to further discuss on using an auction to solve problems at LGA.

[2] In real life, it is expected that companies will spend resources in strategizing for the auction.

of a combinatorial auction from the agent perspective; second, by laying out possible points of concern in the design of the current proposed combinatorial auction design to the LGA case through our calculation of the restricted equilibrium of possible common heuristic strategies.

Our findings in regards to the auction design include the following:

- Interaction between budget constraints and the revealed preference activity rule were found to be problematic. In the presence of budget constraints, straightforward activity rules following the revealed preference rule can be problematic.
- Undersell in the auction can be severe, at 8% in our experimentations for straightforward strategies, even after allowing for some oversell in smoothing out results(as specified in the auction design).
- Using price predicting agents improves the payoff of the auction and reduces undersell.
- Strategies to game on the activity rule by preserving flexibility for agents bidding in future rounds were generally useful in improving agent expected value.
- Effective strategies differed much within an auction depending on the valuation structure. In particular, *demand reduction* and *shaving* was found to be most effective amongst agents with complementing valuations.

A cautionary note should be attached to any interpretation of these findings. Indeed, most of the findings depend on characteristics of the valuation model. Although the best was done to ensure an interesting valuation model, results may differ once the valuation structure is changed. However, the hope it that this framework can be used to generate further results using new valuations [3].

[3] As we will discuss, Barnhart and Harsha are currently working on generating realistic models for the FAA auctions

1.1 Related Work

Related work to this research come from research in artifical intelligence and economics. The agent perspective to evaluating auction designs derives from methods in experimental economics that have used laboratory experiments to provide insights into the application of economic theory, especially in the field of mechanism design [17]. Yet, using software agents to analyze auctions is a recent idea, sprouting from artificial intelligence research. It is perhaps best exemplified in the Trading Agent Competition (TAC) created in year 2000. Annually, researchers around the world create agents to compete in a travel agency scenario where agents compete in simultaneous markets for the various components of a trip : flight tickets, hotel, entertainment facilities. TAC Classic features a combinatorial problem: the traveling agent problem of buying in simultaneous auctions various items to complete vacation packages for its clients. A variety of techniques were developed through TAC, notably, models for price learning, heuristics to bidding, framework for equilibrium searching, stochastic methods to bidding,... [14, 34, 15] We modeled our price learning mechanism according to the supervised learning model of a high-performing agent in 2001 and 2002 TAC [29].

It is important to note here that research linked to TAC differs from our research in at least two important points. First, TAC features a simultaneous auction setting and thus has yet to study the strategies for the combinatorial auction. Second, TAC can serve also as a testing ground for AI techniques in pitching into competition together (and thus ranking in efficiency) researchers from around the world. The techniques in TAC were used in deriving our strategies. Yet, because our goal is to generate general results on strategies for bidding in ascending combinatorial auctions, our general framework mirrors more closely methods for deriving restricted equilibrium.

Work in applying evolutionary equilibrium to market design analysis has been done recently by Phelps and Parsons in relation to the double-auction mechanism

[27, 28]. Reeves et al. also worked on providing a framework to calculating restricted equilibria using evolutionary equilibrium in the case of simultaneous ascending autions[10]. We extend Reeves' proposed framework to a more complex scenario of heterogeneous agents, and apply it to the FAA auction problem.

We also benefit from studies in mechanism design in generating the strategies used to evaluate the auction. Heuristic strategies such as "ZIP" and "GD" first developed for single item auctions had been used in many studies [31] . We decided, however, to instead develop heuristic strategies to model more recent strategies observed in combinatorial settings. Especially, research conducted on game theory and auctions found many interesting strategies in the game, such as collusive activity, parking, and shaving, which served as inspirations and strategy benchmarks for the heuristic strategies studied in our experiments [18, 33, 23].

Another set of literature related to our research is the set of published works on the application of the auction game to the specific case of airport slots and the expression of airline preferences. The idea of auctioning off airport slots has been proposed as early as in 1982 [24]. Ensuing models were developed for valuations, with various uses of auction mechanism. Similarly, Donohue made a model for Atlanta airport by generating revenue according to the size of the current aircraft leaving the airport [19]. Donohue assumed that airlines prefer their original schedule to any other schedule due to long-time optimization. Ou et al. discuss a variant on this scheme by generating a complete model for aircraft profits with a bidding language[11]. In the valuation model they suggested, Ou et al. based the profit associated to a slot by an airline on the type of airline involved (dominant, low-cost, or a "regular" carrier), the revenue and cost for the flight using the slot (determined from aircraft size, miles flown by the flight, and unit revenue and cost for each airline), and whether the slot is at a peak time or not. Complementarity between slots are also considered, and substitutes between slots at neighboring times are considered. Unfortunately, Ou et al. did not have the chance to test out their model experimentally in a combinatorial setting, a part our project fulfills. Our project also extends their valuation by introducing the idea of budget constraints.

1.2 Outline of Thesis

Our work is best summarized in four stages:(1) Implementation of the current auction design proposed for FAA auctions; (2) Modeling the valuations of participating airlines in the LGA auction ; (3) Modeling agent world knowledge in the form of a staged learning process of price prediction; (4) Calculation of restricted equilibrium between agents of credible heuristic strategies.

In the first stage, we implemented the auction design proposed for FAA auctions as described in Chapter 2 of the thesis. Next, two models were explored to represent agent valuations for time slots at La Guardia. First, there is a linear optimizing model that estimates desired packages by solving an integer programming problem based on existing airline network [5]. We also implemented an alternate model of packages based on the interactions between two types of agents at LGA : incumbents and new entrants. [4] This stage of the research is described in Chapter 3 of the paper.

The next part focused on analyzing the auction through various agent strategies. Chapter 4 lists the issues considered in generating our agent strategies ; while chapter 5 discusses our experimental setup for the simulatoins. First, we trained price prediction models using the supervised learning framework proposed by Stone et al. [29] The process for price prediction is discussed in Chapter 6. Chapter 7 continues our analysis by explaining our methods in generating payoff matrices for the heuristic strategies and the subsequent evolutionary search for equilibria.

The thesis finishes off on a discussion of the results of our experiments.

[4]The first model proved too long to run for our experiments, so our experiments are based solely on the second model. After a few improvements, the run-time for this first model has improved at the time this thesis is handed in. We hope to be able to run on it in further research.

Chapter 2

Combinatorial Auctions

2.1 Past Auction Designs and Considerations

Past auction designs and the considerations raised from their application serve as the motivation to the changes introduced in the FAA simulation. Of these, the most well known use of auctions to allocate items with strong complementarities and substitutes are the spectrum auctions run by the FCC [7]. Run since 1994, the FCC auctions were simultaneous ascending auctions - that is, independent auctions for each of the spectrum running at the same time. Participants were to express their preferences taking into account of the prices of all the items they might be interested in[9]. Concerns for dynamics in simultaneous ascending auctions are, however, expressed on at least two major issues.

First, the exposure problem made bidding strategies complex for agents [10]. The exposure problem involved the possibility for agents to be exposed to winning only a fraction of the package with complementing items when bidding straightforwardly its preferences. To make the problem clearer, let us consider an example:

Agent A values X and Y together at 10. Yet, agent A values either good only at 0. In an auction, agent A may have to bid both at 5, yet wins only one of the two goods. In this case, agent A would suffer a disutility of -5, worse than not participating in the auction.

As we see, the exposure problem is generated by the independent nature of each auction.

Another problem raised by the simultaneous auctions was that the transparency of prices made it easy for the agents to collude by signaling to each other through the last digits of the prices. It was thus noticed that during the auctions, trailing numbers mirroring object IDs in bids recurred, as participants tried to signal to each other and enforce collusion through the threat of punishing a rival by shifting demand to the rival's desired markets [8].

Taking account of these concerns, a different auction design is proposed to the FAA in the style of the clock-proxy auction design [4]. The clock-proxy auction design is in two parts.

First, there is a clock auction (also a price discovery phase) that is an ascending auction with package bidding; this is followed by a proxy phase which has a proxy taking in the preferences of the agents and then bidding on their behalf following a simple rule. However, it was decided in December 2004 that for the simulations before the actual auction at least [1], only the clock phase would be used as it is simple to understand and bid. It was also expected that the combinatorial bidding language should solve many of the inefficiencies observed in the independent ascending simultaneous auctions. In our experiences, we decided to follow the guidelines for these simulations.

2.2 Combinatorial Clock Auction

The combinatorial clock auction is a basic ascending-price combinatorial auction, following a simple clock design. At its most basic form, the clock auction starts with prices of 0 on all objects. At each round, the auctioneer announces linear prices on individual objects. Each agent then responds with the particular bundle it would prefer given the announced prices. Taking into account total agent demand, the

[1] notably, the February 2005 simulations

auctioneer will calculate the excess demand for each item. Prices on items in excess demand will be raised in proportion to excess demand. This process is repeated until there is no more excess demand.

Traditionally, iterative combinatorial auctions solve a *winner determination problem* given the agent valuations. The combinatorial clock auction alone, without the proxy round, is, however, only combinatorial in its bidding language. This mechanism eliminates exposure risk by committing the auctioneer to selling off only bundles of items as specified by the agents, but maintains possible inefficiencies due to the linear increases of the prices. Another key difference between the clock auction design and the spectrum simultaneous ascending auctions is in the activity rule. While the spectrum auctions had an activity rule that forced reduction in quantities bid on as prices progressed, the clock auction relies instead on the revealed preference rule, hoping as such to reduce strategies that try to game on the activity rule by maintaining flexibility. Finally, to prevent collusion, inter-round prices are determined by the auctioneer, taking into account of the overdemand at each round. There is thus very little means of communications between the agents to enforce a collusive scheme.

In this section, we will begin with a description of the auction, including the activity rule, allocation, and pricing. Then, we will discuss the application of the auction to the FAA problem and possible theoretical inefficiencies that may arise from the auction.

2.2.1 Activity Rule

The activity rule dictates the rule agents must follow to make bids. Mostly, activity rules try to encourage truth revelation by forcing bidders to bid according to a downward sloping demand curve. As such, in simple ascending auctions, the activity rule restricts bids by making sure that agents cannot reenter the run for an object after dropping out of the auction. In the case of spectrum auctions, bidders could not increase the quantity of objects bid upon. Yet, this created incentives to parking - to maintain flexibility, a bidder has to game the activity rule by bidding early on many objects with low prices.

As a suggestion to fix this activity rule, Ausubel et al. propose instead to a rule based upon the principles of revealed preference. As such, consider two times denoted s and t ($s < t$). Let p^s and p^t be the price vectors at these times, and let x^s and x^t be the associated demands of some bidder, and let $V(x^i)$ be that bidder's value of the package x^i. A sincere bidder prefers x^s to x^t when prices are p^s:

$$V(x^s) - p^s \cdot x^s > V(x^t) - p^s \cdot x^t$$

and prefers xt to xs when prices are pt:

$$V(x^t) - p^t \cdot x^t > V(x^s) - p^t \cdot x^s$$

Adding these two inequalities yields the revealed preference activity rule :

$$(p^t - p^s) \cdot (x^t - x^s) < 0$$

As such, for all times t, the bidder's demand xmight satisfy RP for all times s ¡ t.

2.2.2 Monotonic Increase in Prices

Prices will monotonically increase through the auction for overdemanded items according to the level of overdemand. The current price function increases the price on an item in proportion to the overdemand observed on the item. The equation is as follows:

$$P_t = P_{t-1} + d$$

where:

$$d = a/b * C$$

- a is the overdemand

- b is the maximum overdemand that can occur given the number of participants and publicized maximum slots available (thus N_{dups}* N_{agents}.
- C is a constant, set at 800 for the simulations we run.

2.2.3 Final Allocation

The auction ends once there is no overdemand on any item. At this point, the auctioneer will go through all the rounds of the game. To minimize undersell, the final allocation will be the round with the set of bids that minimizes the maximum absolute difference between available supply and demand, subject to the constraint that no slot can sell more than 4 overdemand). In other words, it is :

$$minL = \sum_i |g_o i - g_s i| \text{ for all } i.$$

subject to the constraint:

$$g_s i - g_o i < 4 \text{ for all } i.$$

- where $g_s i$ is the number of goods sold for good i;
- and $g_o i$ is the number of goods offered, again for good i.

2.2.4 The various pieces of the auction

The setup for our mechanism is composed of three different pieces: the auction, the agent model, and the agent strategy. At each round, the auction calculates the present prices, along with the existing overdemand on each item (if there is overdmeand), and passes them to the agent strategy piece. The agent strategy then queries the model to get valuations and preferences given the current prices and predicted final prices. Having gotten the real preferences of the agent, the agent strategy then bids according to its strategy, giving back to the auction the package it wishes to bid on at the current time. The auction then verifies the bid according to the activity rule. If the bid is accepted, then the process is renewed. Otherwise, the bid is returned to the agent who is prompted to return a different bid.

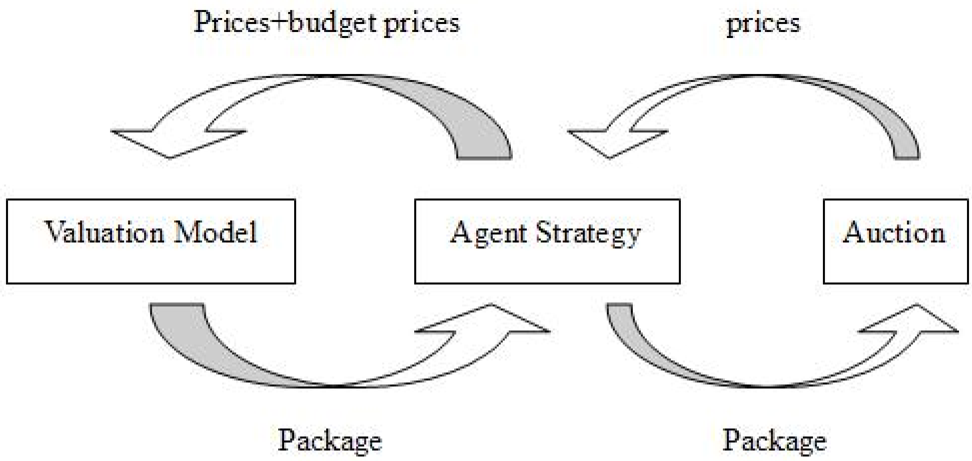

Figure 2.1: a schema of the auction process

2.3 Application to the LaGuardia Domain

The FAA application of the auction is characterized by the large number of duplicates of the same object, and a voucher system to compensate incumbent airlines for the existing facilities at LGA they hold.

Slot Allocation

Originally, the auction was designed to be held for 15-minute interval slots, with thirty-two instances of each slot. However, in the simulations presented during conferences to the FAA and the airline executives, to cut computational time, while illustrating important concepts of the auction, the number of objects was cut to hourly slots from 7 a.m. to 10 p.m. (16 different objects). In our own experiments, due to concerns for computational speed and the large amount of data needed, we used the smaller sample, with the number of duplicates also cut to 12 instances of each object, thus making the total number of items in the auction to be 192.

Vouchers

To auction off the air slots, the FAA needed to compensate incumbent airlines for possibly losing the use of existing facilities they had invested in at LGA. It has been proposed that the auction awards each incumbent a voucher for all slot times they held prior to the auction. The value of the voucher is calculated according to the final price of the slots in the auction, as follows:

$$V_i = (P_{final} \cdot x_i) * \frac{N_{current}}{N_{past}} \text{ for any agent } i.$$

where:

- x_i is the collection of slot times previously held by agent i.
- P_{final} is the vector of all the final prices of the auction
- $N_{current}$ is the current number of slots in the auction
- N_{past} is the number of slots that were regularly used before the auction. ($\frac{N_{current}}{N_{past}} < 1$ to solve the congestion problem.)

As such, airlines receive a market value of the slots they held, normalized to the congested problem. The value of the voucher is thus intended to be as monetary - the voucher can be redeemed to purchase slot times at any FAA auction (LGA or future airport auctions).

2.4 Examples and Discussion

In this section, we want to illustrate the characteristics of the auction we've described thus far with a few example cases. Through these cases, we hope also to illustrate both advantages and vulnerabilities of the auction that later guided our strategies.

2.4.1 Example Case 1 : Simple Example with Efficient Allocation

In the first example, we present an auction with two different objects, X and Y. There are two different agent participants in the auction, A and B.

Suppose A has valuation on two substitute packages:

- either X which A values at \$3
- or Y which A values at \$3

(Notice that agent A has utility for only one of the items, as the items are subsitutes of each other.)

B has valuation on one package:

- X and Y which B values at a total value of \$2 (having either one is of value 0 to B; these two items are complements for B).

It is clear here that the efficient allocation is to have agent A get item Y, while item X remains unsold, or given at price 0 to any agent.

We show that the dynamics of the clock auction proposed will lead to this efficient allocation, while a simultaneous ascending auction will not if it does not allow for quantity reduction on items whose prices have not increased, especially if agent B is *sunk-aware*. The above table summarizes the results that we should observe round-to-round from a simultaneous ascending auction and a combinatorial auction.

We suppose here, without loss of generality in the auction (assuming that the auction increases prices slowly), that the prices increase by \$1 at each round when there is overdemand.

As we see from the table, the clock auction allocates the efficient allocation. However, in the simultaneous ascending auction, if agent B is blindly responding, there is an efficient allocation with agent A obtaining item X. However, agent B then gets a disutility of -1.

Round	agent A bid	agent B bid	Price of	Price of Y	Comments
1	Y	X and Y	0	0	
2	X	X and Y	0	1	
3	X	X and Y	0	1	
4	X		2	1	END for Combinatorial Clock or Sim. Auction
5	X	X and Y	2	1	
6	Y	none	3	1	END for Sunk Awareness, Sim. Auction

Table 2.1: Example 1 : Progress of Auction

Notice here that agent B has an incentive to bid with sunk-awareness, whereby B calculates the utility from bidding while taking into account of the fact that he cannot withdraw bids on items whose prices do not increase. In this case, B would continue to bid on its package even past the optimum point as there is the same disutility of -1 for the two possible outcomes: whether it wins the entire package at a price of 3, or if it has to pay the price of -1 for something for which it has no value. If agent B bids with sunk-awareness, then the system pushes instead item Y to agent A while agent B does not pay for anything - an inefficient outcome by all means.

2.4.2 Example 2 : An Undersell Problem

A major shortfall of this combinatorial clock auction is that the auction imposes monotonically increasing linear prices - long deemed problematic. Indeed, research has shown that linear prices are generally not rich enough to capture the efficient allocation. In particular, undersell can occur. The following example illustrates this dynamic.

Suppose that we now have three agents, but still with two objects.

A has valuation on two substitute packages:

- or Y which A values at \$5

(Notice that agent A has utility for only one of the items, as the items are subsitutes of each other.)

B has valuation on one package:

Round	agent A bid	agent B bid	agent C bid	Price of	Price of Y
1	Y	X and Y	X	0	0
2	Y	X and Y	X	1	1
3	Y	X and Y	X	2	2
4	Y	X and Y		3	3
5	Y			3	4

Table 2.2: Example 2 : Progress of Auction

- X and Y which B values at a total value of \$4 (having either one is of value 0 to B; these two items are complements for B).

C has valuation on one package:

- X which C values at \$2

It is clear here that the efficient allocation of the auction has agent A winning item Y, and agent C winning item X. Yet, if we look at the progress of the bidding in the following table, the final allocation of a clock auction with no smoothing for undersell (continuing until there is no more overdemand, disregarding undersell), is merely an allocation of Y to agent A, but item X is not sold to anyone. In the FAA auction design, a dampening effect (by allowing for mild oversell), is used to help mitigate such an inefficient allocation. For more pronounced and complicated valuations models, however, allowing for limited oversell may not solve this problem.

2.4.3 Example 3: Adding Budget Constraints

A key component of our valuation model is the introduction of budget constraints. The interaction between the activity rule and the budget constraint is an important one that we would emphasize later in the auction. It is thus fitting that we also illustrate an example with budget constraints and the activity rule here.

Again, we present an auction with three different objects, X, Y and Z. There are many agents participating in the auction, notably agent A who values X and Y together at \$20; or Z alone at \$10. A has a budget constraint of \$8. It is clear that agent would start bidding on X and Y first. Assume that prices go up at uneven paces (some agents drop out and some come in on certain items). It is possible that, at a round t, the price of X and Y is \$8, whereas the price on Z is still only \$2. Suppose that at this round, many agents move from bidding on X or Y to bidding on Z. Consequently, at $t+1$, X increases its price by \$1, such that X and Y is now priced at \$9, beyond the budget of A. On the other hand, Z increases by \$3 to \$5. It is clear here that agent A would want to shift bidding to Z. However,the revealed preference rule indicates:

$$(p^t - p^s) \cdot (x^t - x^s) < 0$$

for all $s < t$.

As such, agent A would not be able to bid on Z anymore. This is the budget-revealed preference conflict that we will discuss in more details in the next chapters.

Chapter 3

Valuations for the LaGuardia Domain

3.1 History of the Problem

Overcrowding at U.S. airports had been an issue for decades. Unfortunately, some of these problems cannot be solved in the foreseeable future without a reduction in the demand. Indeed, for many airports, such as La Guardia, extensive expansion of the airport to accomodate rising demand is physically impossible (the Federal Aviation Association (FAA) predicted that by 2010, technology improvements can improve the capacity by 10%, while the predicted increase in demand for the same period is 17%)[25]. Yet, reducing demand efficiently also holds many challenges. In 1968, several US airports were subject to "High Density Rules" with limits to the number of take-offs and landings per hour. Since then, airports distribute slots at particular times, which is defined as "a reservation for an instrument flight rule takeoff or landing by an air carrier or an aircraft in air transportation".[16] This code was criticized as inefficient. Yet, when a reform to the high density rule was enacted in 2000, thereby reducing the restrictions of the code, demand soared in La Guardia with escalating delays. In September 2000, LaGuardia faced more than 9,000 delays, representing 25 percent of delays in the United States[1].

Although a re-enactment of HDR in 2001 followed by the September 11 terrorist attacks soothed the delays, it is expected for the demand to soar anew after the HDR rule expires in 2007. [2] It is clear that the status quo for LGA is unsustainable, especially with the escalating effects on other airports due to the network nature of the airline industry. Yet, reducing demand also holds also many challenges - in particular, there is a need to ensure that the rare resource that these time slots have become are used to the maximum of their value, thereby falling into the right hands perpetually. A market mechanism using auctions is currently a favored solution for a more efficient use of the airport, creating lower barriers to entry into the market.

3.2 Models used for valuating airport slots

There are two major challenges in generating valuation models for airlines. The first is in understanding the actual microeconomics that guide the airline sector. There are clearly strong incentives for airlines to protect this information. It is thus natural that, in the conferences used to discuss the use of auctions so far, airlines have not been straightforward in showing the actual valuations they have on the slots. The second challenge is in the sheer number of possible bundles on which an agent can calculate preferences on. Indeed, the number of bundles possible is exponential in the number of objects available in the auction - and as such intractable. For the FAA simulations, two models were developed to estimate the value of air slots hoping to bypass these two problems.

The first proposed model, and the one originally intended to be used for the simulations was a maximization problem model formulated as a integer programming problem. The model was developed by Barnhart and Harsha in a partnership research project and designed to take into account of the existing network of airlines, complete with the design of flows of fleets of airplanes from one airport to another. Unfortunately, the model was not completed in time for the running of the experiments. [1]

[1]Due to the NP-hardness of the problem, the querying of model preferences took very long computational time to arrive at a solution. At the time of the end of writing of this paper, the

Therefore, a second model based upon interactions between incumbent airlines and new entrants was developed to both reduce the runtime for performing the experiments, and also to simulate the decisions a human bidder with limited instructions (thus bounded rational) might have to make while bidding in the auction. Indeed, if, in the first model, valuations had to be computed electronically by sophisticated computers at each round, it is possible, given the precedence of the FCC auctions, that bidders in an actual FAA auction could be executives who have in mind a finite number of business plans each with a specific complement and substitute structure.

3.2.1 Budget Constraints

An important characteristic common to both the simple bounded rational model and the linear optimization model is the presence of budget constraints - a characteristic we'd like to take a moment to justify.

In standard auction literature, budget constraints have not been much considered. Indeed, although budget constraints are constantly held out in consumer utility maximization models, most auction designs for large-scale commercial auctions assume that company bidders should be constrained only by the productive power of the item bidded upon. This is an assumption for an efficient financial market, and thus, for unconstrained liquidity. In reality, however, it is unclear that companies with troubled recent history (especially ones in the current airline industry), can actually raise the money necessary to win a slot on which they can make money in the future. This is an especially critical case given the comments made by airline executives participating in the National Center of Excellence for Aviation Operations Research (NEXTOR) conferences on the cash constraints under which they currently operate.[2]

---

valuation model has become more tractable, with solutions for auctions for 4 time slots and a restricted number of 3-5 cities and 3 fleets in the agent network considerations. However, there was no time to run tests with the strategies on it.

[2]Two conferences have been held so far. The first one was held in November 2004, discussed administrative solutions to LGA such as a Passengers Bill of Rights that will force airlines to compensate travelers financially for delays. The second was the mock auction held on February 24-25, 2005.

While the conclusions to be drawn from the various arguments on efficient markets are unclear, there has been research that the complications that arise from budget constraints can be potent [6]. As such, given the current particular position of the airline industry, with major airlines having filed bankruptcy protection within the past two years, we believe that the presence of budget constraints are an important concern to address and study for an LGA auction.

3.2.2 Model 1: A Linear Optimization Model

In the first model, an integer programming problem is posed with the goal of optimizing an objective function in a network of cities. Although we did not get to test this model with our strategies, the model was plugged and integrated into the auction system. We present here the model both in hopes of being able to pursue further agent-based research on it in the future, and to show also the potential for agent-based analysis in studying even complex valuation functions.

The objective function in the model predicts the profits an airline can make given a bundle b_m at prices p_m. As such, the object is :

$$\max \left(\sum_{i,j,t} Fare^{t}_{i,j} P^{t}_{i,j} - \sum_{i,j,k,s,t} C^{t}_{i,j,k} x^{s,t}_{i,j,k} - \sum_{m} p_m \max\{0, b_m - v_m\} + \sum_{m} r.p^{F}_{m} \max\{0, v_m - b_m\} \right)$$

where :

- the first term refers to the expected revenue from representative fares flying from city i to city j with P passengers.
- the second term estimates the cost of operations for maintaining the route between city i to city j scaled by the number of flights x flown daily.
- the third term is the cash paid to buying the bundle b_m in the auction at prices p_m and with vouchers v_m;

- and finally, a fourth term to indicate the amount obtained from the unused vouchers provided to the agents in the auction design.

Constraints applied to the objective are the following:

- the budget constraint that takes into account of the vouchers incumbents get from flying:

$$\sum_m p_m b_m \leq B + \sum_m p_m \min\{v_m, b_m\} + \sum_m r.p_m^F \max\{0, v_m - b_m\} \qquad (3.1)$$

  There is a differentiation here in the valuation between the amount of vouchers used in the auction (the second term), and non-used vouchers. This is due to the FAA rule that vouchers can only be redeemed within an FAA auction which reduces the liquidity of non-used vouchers.

- the bidding limit and operational limit constraints which restricts the maximum number of slots an airline can bid on (due to supply of objects in auction and a limit on the number of flights the airline can operate in an airport at any single time):
  Bidding Limit

$$\sum_m b_m \leq L \qquad (3.2)$$

  Operational Limit

$$b_m \leq O_m \qquad (3.3)$$

- and a series of balancing constraints that conserve the flow of aircraft to ensure that fleets used for all flights arrive and leave an airport within a reasonable window of time so that all flights requiring the specific fleet find it available. These constraints make for the network effects of the airport.

  This model provides a fairly complex model of airport valuations with strong considerations of network effects in the airport. At the time of the writing of this thesis, the model runs in reasonable time for small problems with 4 airslot times, and small network considerations (with 3-5 cities, and 3 fleet types). The challenge is still present for expanding the model to make the runtime reasonable for the size of the real auction problem.

3.2.3 Model 2 : Bounded Rationality with Incumbents and Entrants

For this simpler model, we took inspiration from the work from Ou et al. [11] and built a bidding language in the form of XOR-AND-XOR to express a valuation function. We will describe the bidding language in more details at a later point in this chapter, after we present the microeconomics employed in building numbers for the model itself.

Valuation Tree

To use a tractable bidding language, we employ a valuation tree, which is expressed in the tree-based bidding language proposed by Parkes et al.[13] The language is one of XOR-AND-XOR, expressing both substitutes and complements in the auction.

The tree representing the valuation has goods as leaves. The leaves are the only object in the tree with a value on them. This is important in the calculation of the total value attached to the tree.

One level above the leaf level is an XOR of substitutes. In particular, neighboring slot times can also be used to substitute each other, albeit at a discount. Thus, the XOR would have children the substitute goods, each with a different value should it be won.

The parent of the XOR is an AND. This represents all the slots that the agent needs to satisfy a particular business plan. This is the complementarity factor: an agent needs to satisfy all the nodes to obtain the implied utility from these nodes. Notice that an agent gets zero utility should a single node be missed.

Finally, at the top of the tree is an XOR of all the business plans the agent can follow. In particular, the plans differ by the *cities* that the agent plans to serve in each case. Currently, the number of business plans an agent considers depends on the number of *cities* the agent plans to fly to, fluctuating between 1 to 15. We will return to this point later, after discussing the role of *cities* in the model design.

An example of a valuation tree can be illustrated as follows:

BUSINESS PLAN 1

Flight 1: (800, 200) XOR (0900, 130) XOR (0700, 130) (City 1)

AND

Flight 2 : (1200, 200) XOR (1300, 130) XOR (1100, 130) (City 1)

AND

Flight 3: (1700, 200) XOR (1600, 130) XOR (1800, 130) (City 1)

AND

Flight 4: (800, 100) XOR (0900, 70) XOR (0700, 70) (City 2)

AND

Flight 5: (1600, 100) XOR (1700, 70) XOR (1500, 70) (City 2)

XOR

BUSINESS PLAN 2 :

XOR...

Incumbents and Entrants

In the literature on airport slot valuation, generally three types of airlines are identified: a *dominant* airline, or hub airline; a *regular* airline, and a *low-cost* airline. As 92% of the travelers are local origin/destination traffic, LaGuardia does not serve as hub to any airline [25]. Indeed, while there are three major carriers at LaGuardia, covering about 40.5% of the market share, none of them dominate the market. Indeed, American carries about 20% of the passengers, with USAirways and Delta each contributing about 10%. The remaining passengers are distributed among a few smaller carriers, each with less than 7% of the market share, notably United, Spirit, Northwestern [22].

Thus, given this special nature of LGA, we decided to split the market instead into two types of airlines : the incumbent airline that already has slots at LGA, and the new entrant airline. We differ these airlines, first by their CASM and RASM (the cost per available seat-mile and revenue per available seat-mile, respectively, two constants generated randomly); and also by their general strategy. Below is a table with numbers taken from [11].

Type	Revenue per ASM	Cost per ASM
Regular	U[0.09, 0.11]	U[0.08, 0.10]
Low-Cost	U[0.08, 0.09]	U[0.06, 0.08]

Table 3.1: Differences between the incumbents and entrants

The incumbent airlines (apart from Spirit) are regular airlines, whereas entrants are all deemed to be low-cost carriers. In general, the two types of agents also differ in their schedules. Indeed, we assume that the incumbents have basically non-overlapping schedules that they wish to maintain post-auction (due to the costs of changing schedules), whereas the entrants have identified one to three cities with which they would like to start serving La Guardia. More precisely, at the beginning of a simulation, a past schedule is built encompassing all the incumbents, splitting on average the allocation of slots according to the present market share of each airline.

For this part, we assign the owner of a past slot by sampling uniformly the incumbent airlines for a particular slot, such that the probability an airline gets assigned the slot is equal to the market share it currently holds at LGA. The market share is divided as follows:

Airline	Market Share
American	45%
Delta	12.5%
Usair	13%
Incumbents 1-6 (each)	7%

Table 3.2: Incumbent Market Shares

For the entrants, we assume that entrant airlines are mostly economy airlines looking for a few slots to pair with their existing networks. As such, their desired number of slots is sampled randomly from an uniform distribution $U[1, 21]$. In addition to this base demand, additional demand is also generated randomly for both the incumbent and economy airlines to bias demand towards currently observed bias in scheduling (overdemand in morning and early evening).

After a desired schedule is assigned, the values attached to the air slots is generated by assigning cities that each agent flies to. The value attached to a slot is highly dependent on the city to which the airline flies to from this slot. We assume that all flights leaving from the same city exhibit strong complementarity as airlines are able to maintain their presence in that city. (This is especially true of frequent shuttle services, such as the Delta shuttle between Boston and LGA). Each city is characterized by the following variables:

- numflights per day: Distributed as $U[1, 10]$
- scheduleOfFlights: a set of times agent finds most desirable to fly to this city, sample randomly from the slots available to the agent.
- complementarityFactor: The complementarity factor agent gets from being able to fly the number of flights per day to the city. Distributed as $U[0, 0.5]$.

- distance : Sampled from N $(1170, 1000)$ This is the distance between the city to LGA.
- numSeats: Sampled from $U[30, 200]$. This is the number of seats the average aircraft connecting this city and LGA has.

Once the entire existing schedule has been generated for the incumbents, values at the leaves are generated as follows:

$$Profit(i, x_{i,j}) = V(i, x_{i,j}) * Comp_{i,j}$$

$\forall$ participants i, and time slots x for city j to which airline flies to.
where:

- $Comp_{i,j}$ is the complementarity factor for winning all the slots airline currently flies to city j.
- $V(i, x_{i,j})$ is the value depending on the network of i, such that :
  - if x is the time in the schedule that airline i wishes to fly to city j on :
  $$V(i, x_{i,j}) = (RASM_i - CASM_i) * distance_j * numSeats_{j,i}$$
  - Otherwise, for all x such that $C(x \pm 1, i) = C(y, i)$ exists, we have:
  $$V(i, x_{i,j}) = (RASM_i - CASM_i) * distance_j * numSeats_{j,i} * subsFactor$$
  where $0 < subsFactor < 1$, is the substitute factor for neighboring timeslots to the one desired.

In other words, an agent values slots depending on the network of cities it currently has planned for. Substitutes are for neighboring times of most desired time slots.

There are ten incumbents and four entrants in any given simulation. In building this model, we took care of encompassing three key characteritics: first, credible complements and substitutes are expressed for all agents; second, there is a budget constraint on the model related to the voucher scheme prescribed by the auction design; and third, that microeconomics explain the values generated for the various time slots.

Budget Constraint

In our bounded rational model, budget constraints are generated first by an assigned cash constraint on the agents which is increased for the incumbents by the vouchers on previously held time slots, as described in the auction design in the last chapter.

The initial cash constraint is generated by by randomly sampling from a Gaussian distribution built upon the efficient average payoff of agents in the auction. In particular, we generated 200 different games that we solved for the efficient allocation by converting the games into integer programming problems, and solving with CPLEX. From there, we calculated the average payoff ($\bar{P}$) and variance (σ_p^2) for the system, under efficient allocations. The budget for the incumbent is distributed as follows:

$$B \sim N(\bar{P}/4, \sigma_p^2/16)$$

while the budget for the entrant is the following:

$$B \sim N(\bar{P}/2, \sigma_p^2/4)$$

We chose here the budget for the incumbent as half the budget of the entrant for two reasons. First, given the nature of the incumbents at LGA (mostly regular airlines who have filed bankruptcy in the last two years), it was natural to find them more budget constrained. Second, incumbents had access to vouchers in the auction which would increase the budget available to them in the game in general. The budget chosen as such should be binding, but should also allow for the same efficient allocations as the normal circumstances.

Critique of Model 2

The model provides a simple and flexible framework built upon microeconomic principles. In particular, the idea of multiple business plans may capture well a certain part of the bounded rational reality of the managers actually participating in the auction, with a constrained number of predefined plans to take their companies. This representation also cut on the computational time to run the experiments we will describe in chapter 5.

To provide basic micro-foundations to justifying the values assigned to the slots, we used a network of cities and fleets, combined with the airline ASM. Credible complements and substitutes are expressed in the valuation tree in terms of business plans and neighboring slots, respectively. We also built a budget constraint to reflect possible liquidity constraints in the current airline industry.

However, the simplicity of the model glosses over many aspects of reality. For example, the difference between the incumbents and entrants may be more pronounced than model: it seems more realistic for entrants to have more broader valuations (maybe not for a flight at a particular time, but rather for a number of flights per day at many different types of time). If this is the case, another valuation generation scheme would need to be used. Also, the AND component of the model at the second level seems overly rigid- a human agent may have more flexible business plans than a specific number of slots that must be won. An "OR" may be more expressive - but also brings greater computational complexity. In the end, however, we aimed for a model that can capture key characteristics of the auction - and this model serves well for this purpose.

Chapter 4

Vulnerabilities of the Auction and Possible Trading Strategies

To better understand the effect of common strategies seen in past auctions, as well as to explore possible gaming effects of the current activity rule, we devise two sets of experiments. In the first part, we focus on the strategy of price prediction. Indeed, one of the most common techniques TAC agents used to solve the exposure problem when playing in simultaneous ascending auctions was to learn forms of price prediction and bid accordingly. Price prediction was also useful in general when activity rules can be constraining even in terms of truth expression. Thus, although the combinatorial clock auction does not have an exposure problem in the style of simultaneous auctions, we found that the interaction between the binding nature of the activity rule and the budget constraints made it useful for the agent to price predict its budget. Our first part of the experiments was thus to build this baseline strategy of price prediction and compare it with the straightforward maximization strategy.

Apart from price prediction, we also wanted to better understand the effects of other strategies found useful in the auction. For this part, we devised a set of parametrized trading strategies within which we calculate a restricted evolutionary equilibrium. The strategies we paid attention to can be classified into four categories:

'straightforward' bidding which is based on simple maximization ; *'flexibility retaining'* in order to gain more flexibility in future bidding; *'shaving'* to try to free-ride in cooperating bids; and *'demand reduction'*. The following section discusses in details the role of price prediction in a trading strategy, and the logic and implementation behind each of the other strategies in the restricted equilibrium experiments.

4.1 Price Prediction

4.1.1 Motivation for Price Prediction

Price prediction can be seen generally as a profit predictor that can help guide agent decisions in the bidding. This is particularly useful if the activity rules can constrain the bids the agent is allowed to bid upon in the future. As such, in simultaneous ascending auctions, price prediction can help prevent exposure to sunk losses by helping the agent make decisions on a broader level than myopically responding. It is thus not surprising that price prediction is among the most commonly used strategies to calculate the optimal bid in the TAC competitions[14].

The interaction between the budget constraints and the revealed preference activity rule motivate for us the use of price prediction. As we saw in the auction design, the revealed preference rule forces participants to always bid on packages whose price increase is less than the price increases of packages participants had previously bid upon. This was based on the revealed preference principle - that participants must prefer a package at a time than others if its utility on that package was higher. Yet, this activity rule does not take into account of a budget constraint. Indeed, with a budget constraint, a participant may have a higher utility from consuming a package A than package B at time B, but cannot do so because of the budget constraint. In this case, if the revealed preference rule is enforced strictly, the agent may be forbidden to express its true preferences under budget constraint [1].

[1] Refer back to example presented in chapter 2

This is particularly true if budget constraints are particularly constraining, or unknown. This is the case for the LaGuardia auctions as incumbents' budget constraint depend directly on the amount of money generated from the vouchers. In real life, market knowledge may help give bidders an idea of the amount to expect from the auction as well as other information as to what the agent can expect to win. In the trading agent problem, this knowledge can be simulated through price predictions - indeed, if the agent is able to predict the final prices, it could use that information to better optimize its bid.

4.1.2 Process to Simulate Price Prediction

Most forms of price prediction look at history of games played[14]. However, because price prediction formed the base strategy of our games, there were no examples of past games to learn from. As such, we devised a 6-stage-learning process for the price prediction. In the first stage, games were played without the activity rule and with all agents myopically optimizing at each round. The results from these games were then learned through a machine learning algorithm. The resulting model was then used as the price predicting algorithm for games played in the subsequent stage. This process was repeated six times. The process and algorithm would be discussed in deeper details in the next chapter.

The resulting model served the basis for price prediction in all our heuristic strategies. Although we were worried first that the accuracy of the price prediction algorithm is lessened as agents using heuristic strategies are introduced in the environment, we found the error of prediction in these circumstances quite stable.

4.1.3 Price Prediction and Budgets

As described in Stone's paper[29], there are many ways to use predicted prices. In particular, one may bid according to the predicted prices, using them as prices when bidding in the auction. Also, predicted prices can be calculated as a distribution through which one can calculate the expected optimal bid.

In our case, however, we wanted a simple baseline agent with price prediction with which to test the various heuristic strategies - one with which a parallel can be drawn between this agent to the traditional straightforward agent often used in auction research. Moreover, while price prediction could also be used to guide bidding on the likely objects agent can win, we felt that this would cut on the benefit of clock auction itself- price discovery. As such, using price prediction to predict budgets but myopically calculating utility from the current prices proved a good compromise.

It is clear that a baseline strategy involving only budget constraints does not resolve the problematic interactions between the budget constraints and the activity rule. This permitted us to test further heuristics that can further improve on this problem.

4.2 Strategy Class 1 : *'Flexibility Preservation'*

The first heuristic class of strategy addresses further the problem of interacting budget constraints and activity rule discussed earlier. Interestingly, this problem has similar qualities to the constraints imposed by the spectrum activity rules which dictated that agents could not increase the number of objects bid upon. One of the most common ways agents used to game these constraints was to try to bid to preserve flexibility in the beginning of the auction rather than to reveal their true preferences[3]. This is a technique often referred to as *parking*.

In the aftermath of the spectrum auctions, researchers found that "*parking*" was one of the main problems created by the activity rule. Parking is commonly defined as a bid "that the bidder does not intend to win- made to preserve enough eligibility so that the bidder can bid on its desired markets later in the auction."[8] Yet, parking can also be used for collusion enforcement and demand reduction purposes - such that the bidder prevents price increases on the items it cares about, pushing instead rivals onto concentrating their capital in their territory.

In general, parking was found problematic as it could sometimes force price distortions that exclude the high-value bidder from winning the item it should get. It was actually to eliminate incentives for parking that the proposal for the FAA auctions included an alternate activity rule based upon principles of revealed preferences as presented in chapter 2. Yet, in support of the problem found with interacting budget constraints and the revealed preference activity rule, we found in our experiments that parking remained a dominant strategy for incumbents in the auction.

4.2.1 Implementation of '*parking*'

In our implementation of a parametrized parking strategy, we define the bid on a package as follows:

1. Choose an α
2. For each round:
   - $bestBid = getBestBid(P_{current}, P_{predicted)}$ which is the true package wanted given present prices following straightforward base strategy.
   - if $Mean\left(\frac{P_{predicted} - P_{current}}{P_{predicted)}}\right) > \alpha$ then:
     - choose bid such that : $Q_{bid} = Q_{bestBid}$ (quantity of objects bid upon is the same as the quantity of object agent would have bid upon given these bids) AND $P_{bid} < P_{bestBid}$
     - In other words, agent bids on objects whose prices are lowest at the present round up to the quantity of objects agent would have bid on if it was playing the straightforward strategy.
   - otherwise, bid $bestBid$.

4.2.2 *OverdemandMore* - Another strategy to maintain flexibility

Another strategy we also devised to try to test the value of more flexibility for the agent was to weigh utilities on bundles of items to take into account of the flexibility gained for the agent. In this strategy, given two bundles close in utility value, the agent would bid upon the package that allows it to preserve more flexbility in the future given the current activity rules. The proxy we used to measure flexibility was the overdemand of the price of the item. As such, higher overdemand meant faster growing prices that would let agent switch onto the slower price-increasing bundles in the future. As such, we measure the weighted utility $U'(x)$ of an agent for bundle x as such:

$$U'(x) = U(x) * \frac{1}{k*O(x)+1}$$

where:

- $O(x)$ is the average overdemand of x.
- $U(x)$ is the actual utility of x.

4.3 Strategy Class 2: "*Increasing Payoff*"

The second strategy class tries to increase payoff using two common strategies observed in auctions that can be gamed upon. The first is *shaving* the utility when bidding, and the second is *demand reduction.*

4.3.1 *Shaving*

Shaving is often deemed helpful in auctions where the Vickrey outcome is not guaranteed with straightforward bidding. In combinatorial auctions, shaving is seen most useful to try to free-ride on a bidder with complementing valuations. To illustrate an situation where shaving may be useful, let us consider the following example:

Let this be a world of two agents, two objects again. Agent A values item X at 4. Agent B values item Y at 5. Agent C values item X and Y together at 4, with no value for either one by itself.

It is clear here that the optimal allocation is to give X to A and Y to B. However, because C only values the two items together, if B bids up to his value (4), A will not have to bid up X at all (except at the first round to express interest) to have Y allocated to him. In this case, it might have been advantageous for either A or B to only bid up to a shaved value of their true valuation.

Implementation of *Shaving*

At the most basic level, shaving involves determining an α such that agents bid only up to a value $V(x)$ in the auction for every bundle x such that:

$$V(x) = U(x) * \alpha$$

The first implementation that we wanted to do was to try shaving while varying *alpha* (possibly from 0.5 to 1) in the various games to create a payoff matrix of the various profiles and expected valuation an agent playing *alpha* can obtain. However, we found that the resulting payoff matrix would require too many entries (the size of a payoff matrix for a two-strategy game with the 14 agents we used in the model is 50; for a 3-strategy game, it would be 1800). Especially, since each game needs an average of 4 minutes to complete, we did not have enough time to populate a payoff matrix with the number of entries required.

Instead, we chose α to be the theoretical Vickrey payment the agent could expect in a Vickrey auction: $\alpha = (N-1)/N$ where N is the number of agents participating in the auction. It is certain that in the future, research into the variation on the effect variations in the shaving have on the expected payoff of the auction should be interesting.

4.3.2 *Demand Reduction*

The final strategy we tested was *demand reduction*. Demand reduction refers to a collusive process by which bidders drop out of the race for items that have yet to have reached the bidder's threshold value. Such an incentive can be illustrated in the following example:

Assume a two-agent, 3-item world. Both agents A and B either value a package of two of the objects at \$5, or value each object at \$2. If both are prepared to bid the two cheapest up to a total of \$5, there is an incentive for them to instead agree to separate the three objects amongst themselves in the first round such that both have a utility of either \$5 or \$2. This can be done by having one of them drop out of the auction.

This is an unstable equilibrium - either agent has an incentive to deviate and capture a greater part of the pie. Yet, in the FCC auctions, demand reduction has been seen as a problem, as evidence of collusive behavior enforcement has been observed amongst the participants [3]. In particular, agents may be threatening punishment on agents bidding in their territory. In the present auction, collusive behavior has been made more difficult to enforce as the pricing is done by the auctioneer (there is thus no means of communication between agents in the auction), and agents do not see which agents are bidding on which objects. However, there is still clearly an incentive for agents to try to demand reduce each other using information available during price discovery. We try to calculate the restricted equilibrium between demand reduction and straightforward bidding by making demand reduction a parametrized heuristic.

Implementation of Demand Reduction

We used overdemand as a proxy to implement demand reduction. Indeed, it seemed that, given two packages of close value, it could be advantageous for the agent to move to the least overdemanded item. This would reduce the demand on the former package, and could lead to a better distribution of the items amongst the agents. Indeed, if we took the former example, if one of the agents moved to bidding on one of the cheapest items since there is no overdemand on it, instead of bidding on both, the auction could be led to a better allocation for both of the agents than if they competed neck to neck up to the end.

Our model based on overdemand is very close to the model used for a heuristic to preserve flexibility. It is the following:

$$U'(x) = U(x) * \frac{1}{k*O(x)+1} k < 0, andsmall$$

where:

- $O(x)$ is the average overdemand of x.
- $U(x)$ is the actual utility of x.

It should be noticed that this strategy is close to *OverdemandMore* which we used to preserve flexibility in bidding. Interesting, as part of our experiments, we found that both strategies are relevant strategies to the agent.

Chapter 5

Experimental Setup

In this chapter, we'd like to give a general overview of the general experimental setup of our simulations, along with some technical details. Our system is, at its core, a Monte-Carlo simulation of auction games with two types of agents (incumbent and entrants). We have already discussed the generation of valuation models and the auction mechanism. After the system is running, we set a straightforward baseline strategy. We saw in chapter 4 that we defined this baseline strategy by a price prediction model(in our case, a budget-predicting agent trained with supervised learning). Hereon, heuristic strategies are generated to form a complete agent strategy. Finally, we calculate the restricted equilibrium between these heuristic strategies and the default strategy by generating payoff matrices and calculating the resulting evolutionary equilibrium.

In a round, the process for determining the bid for an agent is as follows: first, prices are past from the auction to the strategy agent. The strategy agent passes the current status into the price prediction box which returns back a set of predicted final prices. The strategy agent then prompts the valuation model with the current prices, the predicted final prices, and also any heuristic bias on utility the valuation model is to incorporate when looking for the bundle optimizing utility. For example, if the strategy is for the agent to shave its true value, the strategy agent passes the shaving constant α along with an indication that the valuation model shaves all its

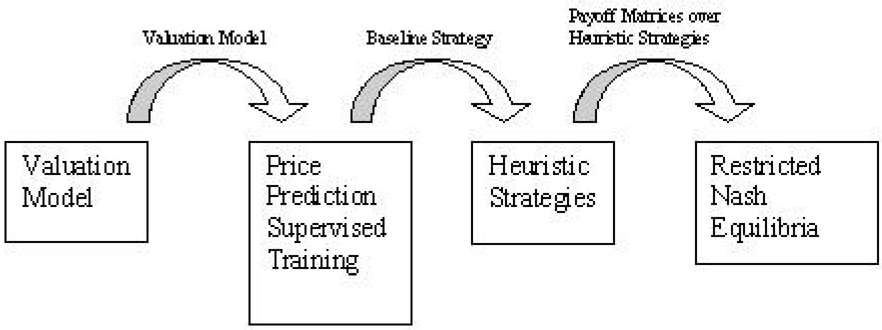

Figure 5.1: A schema of the experimental process.

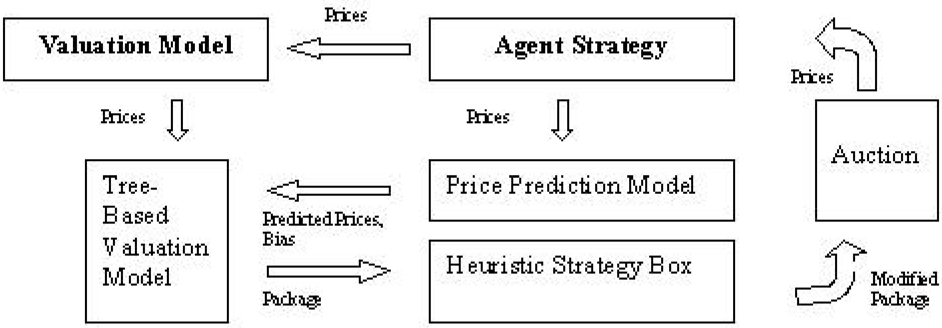

Figure 5.2: The process of selecting a bid.

values. Finally, the heuristic strategy modifies the bid according to the strategy and passes it back to the auction.

Different AI techniques and algorithms are used at each step of our experiments. The model optimization itself is formulated as an integer programming problem that we solve using CPLEX - necessary due to budget constraints and activity rules. For price prediction, we use machine-learning weka packages to train models. Parametrized heuristic strategies that can easily be changed through one parameter are then generated to simulate traditionally popular auction strategies. Running simulations over these restricted parametrized strategies generates for us payoff matrices for the various strategy profiles of a game. To find the Nash equilibria within this small set of bidding strategies, we employ replicator dynamics.

Finally, to evaluate results for efficient outcome, we calculate the efficient outcome of games by solving the winner determination problem over all agents' honest valuations. This is calculated following the winner determination process as described in Parkes et al. [13] whereby the problem is formulated as a mixed-integer problem (MIP) solved using CPLEX, with the structure of the valuations trees captured in the formulation of the problem.

In the following chapters, we will explore in more details the process of developing price predicting models and restricted game equilibria.

Chapter 6

Chapter 6 - Price Prediction with Machine Learning

Many techniques for price prediction have been developed for the TAC games[14]. Our approach follows the machine learning process developed by Stone et al. and implemented for top-scoring agents in TAC from 2001-2002[29]. However, while Stone et al. were able to use past data from past TAC games and pre-competition seeding games TAC participants played, we had to generate data and train our model instead using a staged price prediction process. In this chapter, we will first describe this staged learning process, then go through the learning algorithm itself, to present evidence of the effectiveness of our approarch.

6.1 A Staged Approach to Price Prediction

Most price prediction algorithms are based on past games. In our case, however, we needed to have a budget predicting algorithm in order to play games and generate results relevant to an environment where all agents actually budget predict. In particular, we needed prior results with which we can train a price prediction model first, before playing the games to generate the results to train the prediction model. A staged approach where the environment is gradually generated proved a useful approach.

In our staged experience, the prior used was generated by running straightforward myopically optimizing agents in auctions from which the activity rule wasn't used. A price predicting agent was then trained from this prior. In stage 2, we used the agent trained from stage 1 to run further simulations and generate results. This method is repeated for stage 3 through 6. As the stages progressed, we expected that the environment the agent bid in became more akin to the actual environment of price predicting agents, and that the error of price prediction converges, which was observed experimentally.

6.2 Price Prediction - Machine-Learning Techniques

6.2.1 Training Data Format

To set up the our price prediction as a learning problem, we gathered training data from games run at each stage. In this training data was defined a set of features or attributes known at each round and that can thus be used to predict the final price. The following basic attributes were used to predict the price on an item t:

- t ; the ID of the time slot we're price predicting.
- *currentPrice*; the current price of the item t.
- *overDemand* ; the current level of overdemand of the item t.
- *averageOverDemand* ; the average level of overdemand of the auction (averaged over all the items on sale in the auction).
- *numRounds* ; the number of rounds that have past since the beginning of the auction.
- *averagePriceChange* ; the average change in price of item t over the past three rounds.

A training data instance is created for each item each round in a game, storing all these attributes and the end price of the item of the game. Over the six stages, we generated around 40,000 training instances at each stage. We found this amount of data appropriate for the model to learn a model adapted to the instances.

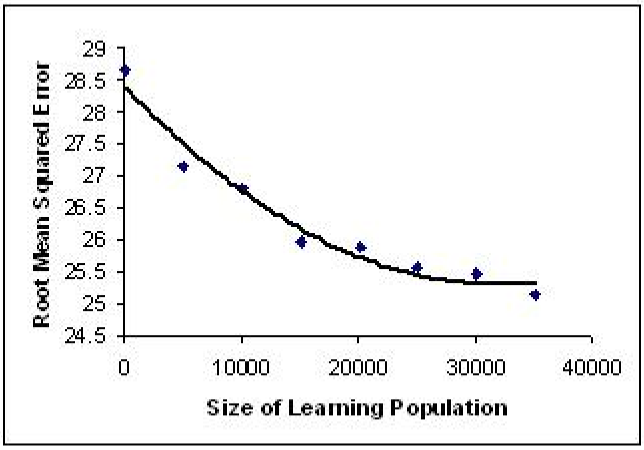

Figure 6.1: Convergence of Learning within A Stage

This set of training data was then used on a learning algorithm. Following Stone et al, which found M5 trees to be most useful in price prediction during TAC games, we tried both the M5Rules and M5P tree algorithms implemented in weka. After experimentation, we found M5P regression trees to be most effective in price predicting in our games.

6.3 Results from Price Prediction

The results from our staged price predicting experiences are presented in figure 6.2. We see that the price prediction stabilized after 4 stages in our staged error process. With a root mean error of around 0.25, we found the final model to be effective at predicting the final prices of the auction.

As shown in figures 6.3 and 6.4, we also found that, as the price prediction became

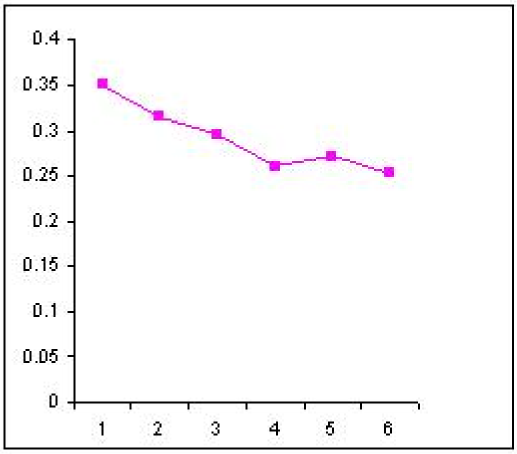

Figure 6.2: Convergence of price prediction results

more effective, underdemand in the auction fell, while the average payoff increased. This indicates that the accuracy of budget constraint price prediction in itself may be important in improving the efficiency in the outcome of the auction. It is also further indication that myopic straightforward bidding does not lead to the efficient Vickrey allocation.

Finally, the price prediction algorithms were fairly robust to the heuristic strategies we introduced. Indeed, we found that the price prediction error remained low in all strategy cases, with the root mean squared error at less than .28 across the board.

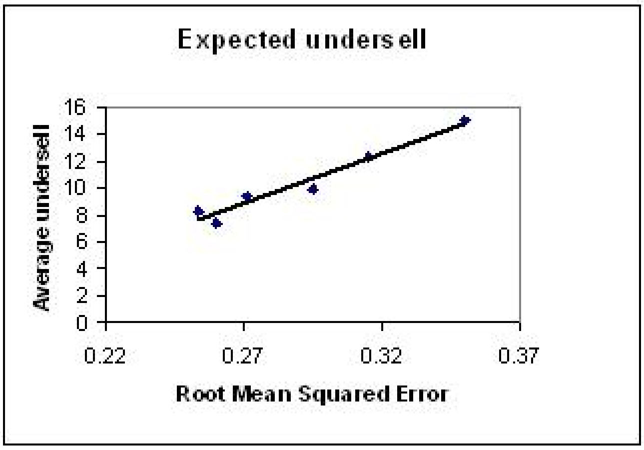

Figure 6.3: Relationship between undersell and quality of price prediction

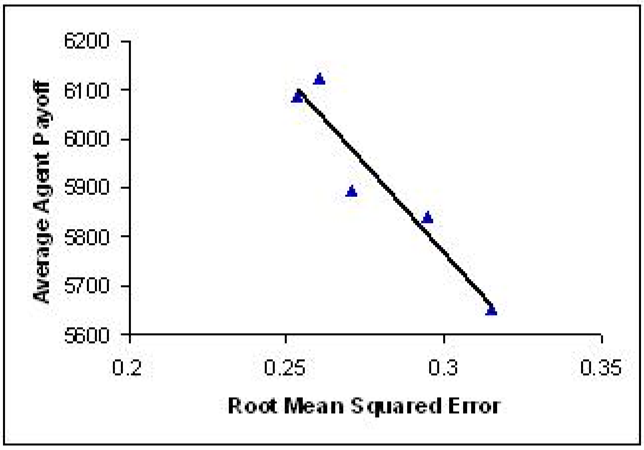

Figure 6.4: Relationship between average agent payoff and quality of price prediction

Chapter 7

Heuristic Strategy Equilibrium

To evaluate the heuristic strategies proposed in chapter 4, we follow the methodology proposed by Reeves et al. in calculating the restricted equilibrium for a small set of bidding strategies[10]. As such, we first convert this game from extensive to strategic form, using Monte Carlo simulation to generate an expected payoff matrix for every combination of the strategies playing against each other. Given the expected payoff matrices, we calculate existing Nash equilibria through replicator dynamics.

7.1 Generating Payoff Matrices

Reeves et al. generated payoff matrices for what they call "restricted games", characterized by specific auction rules, distribution of the domain, and a finite set of strategies permitted to agents[10]. We have discussed in detail each of these of these characteristics of the auction in application to our research in the earlier chapters. To estimate entries of the payoff matrix, we repeatedly sample valuations for agents and simulate games for each game profile.

An entry in the profile matrix is characterized by a *strategy profile* and an agent strategy. A strategy profile indicates the environment in which the simulation is played. For example, a strategy profile of $(0, 0, 0, 0, 0, 0, 0, 0, 0, 0; 1, 1, 1, 1)$ for a simulation indicates that 10 incumbent agents play the strategy 0, while 4 entrant agents play the strategy 1 in the game. Notice here that symmetry amongst the agents helps to reduce the number of strategy profiles to populate the payoff matrix. As such, a profile of

$$(0, 0, 0, 0, 0, 0, 0, 1, 0, 0; 1, 1, 1, 0)$$

is equivalent to

$$(0, 0, 0, 0, 0, 1, 0, 0, 0, 0; 1, 1, 0, 1)$$

as both indicate 9 incumbent agents playing strategy 0, 1 playing strategy 1; and 3 new entrant agents playing strategy 1, 1 new entrant agent playing 0.

An entry in the matrix would be characterized by a strategy profile and an agent strategy, indicating the expected payoff for an agent playing a certain strategy in a particular environment of strategy profile.

To calculate each entry in the payoff matrix, payoffs for the specified profile and strategy are averaged over all simulations played. To reduce variances in results due to the abnormalities in the generation of the valuations, the same set of valuations are used to calculate the payoff matrix of every single profile.

Payoff matrices were the part the most time-consuming of our simulations, and the limiting factor on the number of strategies we could vary per game. For their simulations, Reeves et al. used perfectly symmetrical agents[10]. This symmetry made it easier to generate payoff matrices - a profile of (0, 1) was deemed the same as a profile of (1,0) . Moreover, their simple problem reduced greatly the run-time of a single auction. Still, they found that it took many cpu-weeks to generate payoff matrices of 286 profiles - ten-players with four strategies.

Yet, our valuation really held two very different types of agents - the incumbents and the entrants. Moreover, the model we developed had 14 agents. This greatly constrained the number of strategies we could feasibly run. Indeed, for 2 strategies, the payoff matrix has 55 different strategy profiles. For 3 strategies, this number jumps to 2200, a number that would take many cpu-months to calculate payoff matrices for. Due to the time constraints, we decided to run simulations only for games of 2

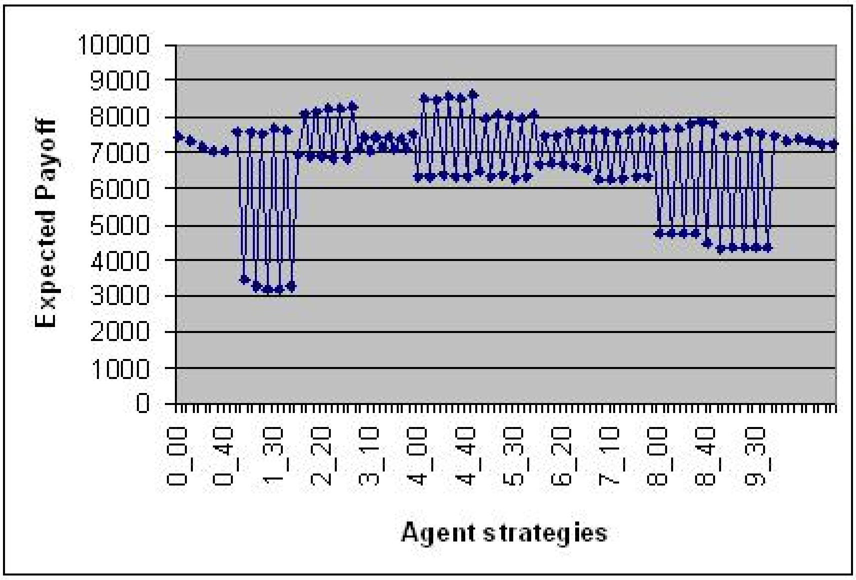

Figure 7.1: Example of a Payoff Matrix for Incumbents: each point presents an entry in the matrix. Neighboring data points indicate different strategies. On the x-axis, X_Y0 the expected return of an incumbent agent playing strategy 0 when X number of incumbent agents playing strategy 1 and Y number new entrant agents playing strategy 1. Each X_Y0 is followed on the right by X_Y1 (the payoff for strategy 1 in that same environment)

strategies, for all the heuristic strategies described in chapter 4, 1100 games run per simulation.

7.2 Evolutionary Search

After payoff matrices are generated, we calculate the Nash equilibria through replicator dynamics. The replicator dynamics were first credited to Taylor and Jonker [30] and Schuster and Sigmund [26]. Our work in this part follows closely the framework applied to calculate restricted equilibriums in the market-based scheduling mechanism developed by Reeves et al [10]. In this framework, an iterative (evolutionary) algorithm is developed to find the Nash equilibria of the system presented.

Because we have two types of population (the incumbents and the entrants), we varied the proportions in each population independently from each other. The total strategy profile is thus a joint strategy with N_I incumbents and N_NE new entrants.

An initial population proportion in a population (either incumbent or entrant) is chosen at time $t = 1$. At each generation, the proportion of people playing strategy s_i in population i, $p_g(s, i)$, is updated such that:

$$p_g(s_i, i) \propto p_{g-1}(s_i, i) * (EP_{s_i} - W)$$

and normalized such that:

$$\sum p_g(s_i, i) = 1 \text{ over all strategies } s_i \text{ for population } i.$$

where EP_{s_i} is the expected payoff for playing strategy s in the population i, and W is the minimum value in the payoff matrix, serving as a lower bound on payoffs.

To calculate EP_s, the average of all the payoffs for s is taken, weighted by the probabilities that these profiles can occur given the proportions p_{g-1} observed for all the different strategies. The payoff from each profile is weighted by the the factor M calculated as :

$$M = M_I * M_{NE}$$

where M_I and MNE are the probabilities for the incumbent strategy profile $(n_1, ...n_s^I)$ where n_s^I is the number of incumbent participants playing strategy s_I. Similarly, the new entrant strategy profile $(n_1, ...n_s^{NE})$, for n_s^{NE} entrant participants playing strategy s_{NE}. As such, the weights are calculated following:

$$M_I = \frac{N_I!}{n_1!..n_s^I!} * p_{g-1}(1, I)^{n_1} \dots p_{g-1}(s_I, I)^{n_s^I} \text{ and}$$
$$M_{NE} = \frac{N_{NE}!}{n_1!..n_s^{NE}!} * p_{g-1}(1, I)^{n_1} \dots p_{g-1}(s_{NE}, I)^{n_s^{NE}}$$

Once a population update reaches a fixed point, then the population has reached a candidate mixed population equilibrium. Although reaching a fixed point is a necessary condition for finding a Nash equilibrium, it is not a sufficient condition. As such, we verify directly that the fixed point is a best response to itself to make sure that it is indeed a Nash equilibrium.

Chapter 8

Results

Through our experiments we obtain two sets of results. In the first set, we analyze the Nash equilibrium generated for the heuristic strategies. In most cases, the Nash equilibrium yields more complex results than expected. In the second set, we analyze the problems and vulnerabilities of the auction as such observed through these experimentations.

8.1 Nash Equilibria Over Heuristic Stratgies

In general, the Nash equilibria in heuristic strategies indicated the use of heuristic strategies. However, the strategies employed differed greatly from type of agent to type of agent. This indicated to us the importance in the valuation model in determining the interactions between agents within the model. In this section, we will present the results we found in each heuristic case, followed by a tentative theoretical explanation to these equilibria. These results, however, should be verified using alternate valuation models.

8.1.1 Result 1 : *Flexibility Preservation*

In general, we found the strategies for flexibility preservation useful for the agents in the auction. The strategies differed, however, on the proxy chosen to preserve flexbility within each type of agent.

Overdemand as Proxy for Price Increase

One of the ways to preserve flexibility was to use overdemand as a proxy for flexibility. Indeed, overdemanded packages tended to have higher price increase - and thus promised more flexibility in the future. This was the heuristic we named *Overdemand-More* in chapter 4. Experimentally, we found that entrants who overweighted utility on packages with fast increasing prices had the advantage in playing the heuristic, while the incumbents had two pure equilibria with equal expected outcome- to all play the heuristic, or to all not play the heuristic. We can see these effets from the figures 8.1 and 8.2:

In figure 8.1, we see that when started with a mixed population of 50/50, both incumbents and new entrants would move to a pure equilibrium of playing the heuristic strategy. Yet, in figure 8.2, when starting the initial population with the majority playing the straightforward strategy, the incumbent stayed at using the default strategy, while the entrants moved completely towards the other strategy - using the heuristic to preserve flexibility. We attribute the results for the incumbents to the advantage observed experimentally that incumbents may obtain in playing the opposite strategy for demand reduction (moving to the items with the least fast increasing prices), thereby reducing potential gains from preserving flexibility this way.

Parking

On the other hand, in the implementation of the *parking* strategy, incumbents were found to prefer to use the heuristic strategy of parking, while entrants preferred to use straightforward bidding. This can be seen again in figures 8.3-8.5.

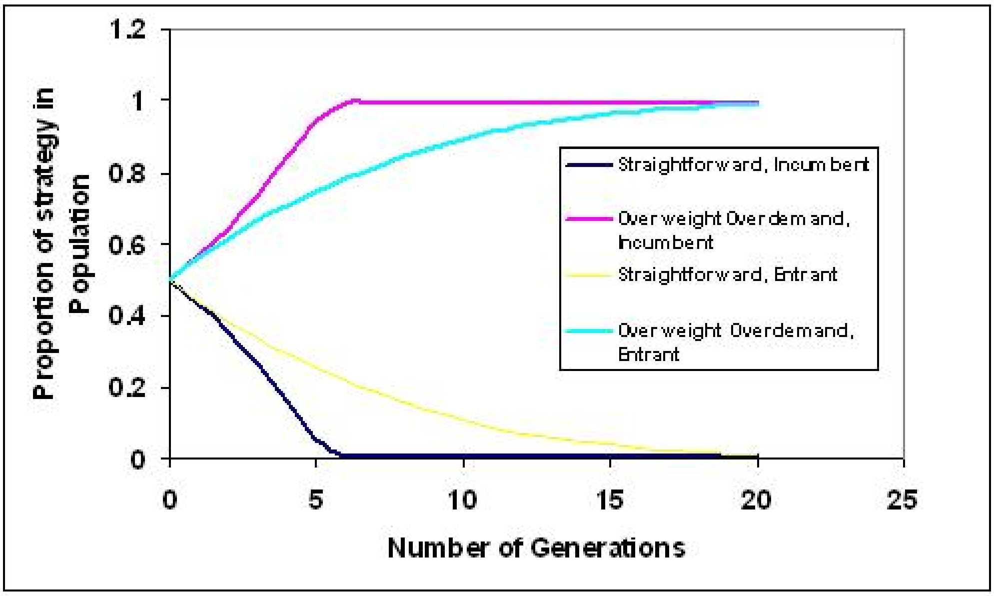

Figure 8.1: Evolution of proportions of strategies in population after generational updating with an initial proportion of 50/50)

At first, this seemed surprising. Yet, intuitively, it is natural that parking can only work if *some* agents in the auction do reveal their true preferences. Otherwise, prices will only be raised on all items without any agent gaining true flexibility or true information from it. It is possible, therefore, that given the structure of the valuations, it is more profitable for the entire system for the entrants to reveal their valuations first. This is probably due to two factors.

First, entrants probably have a higher expected profit on the flights they take interest in than incumbents [1] This makes them more likely to win the item anyway. It is thus better for them to express this interest in the beginning to keep the incumbents off from increasing the price on the item - while incumbents are better off keeping off since they are less likely to win it in the first place. The edge entrants have in bidding is further confirmed by the higher average payoff entrants have compared to

[1]Due to the cost structure in the valuation model, entrants are deemed economy airlines with lower costs, which, according to our model, translates into lower CASM, and probably higher profits.

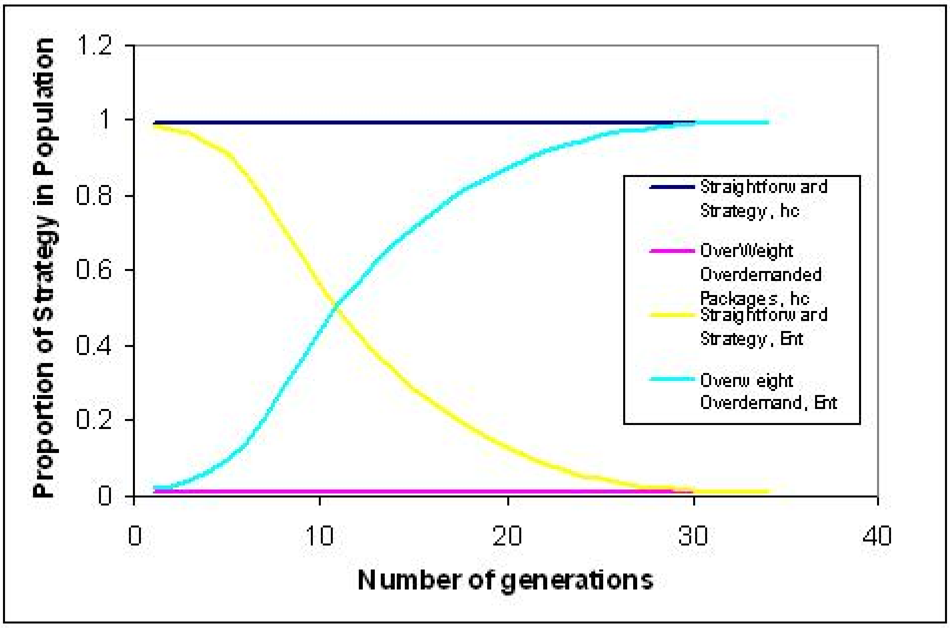

Figure 8.2: Evolution of proportions of strategies in population after generational updating with initial proportion of all default strategies.)

the incumbents (at an average of 9953 for entrants, versus 6301 for incumbents at equilibrium in this set of experiences), despite entrants starting off with less valuations.

Second, due to the way our valuation is constructed, the incumbents have more total number of business plans than the entrants. Subsequent analysis shows that the entrants had on average 45% the number of business plans as the incumbents. This was an unintended consequence due to the incumbents valuing larger packages and more slots than the entrants had at the beginning. Yet, this can translate into more flexibility - and therefore more gain in waiting for more information before deciding what to bid on. This characteristic of the model may not reflect reality (it is very possible that new entrants, without any established schedules are the actual ones with more flexibility) - an overlook when generating the model that should be corrected in future research. Nevertheless, our experiences showed that cooperation between

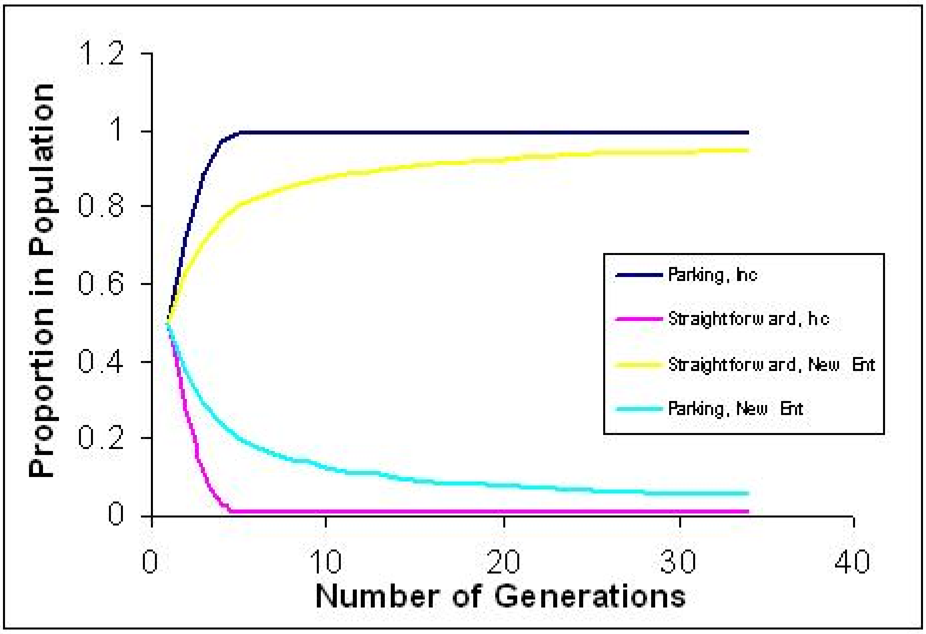

Figure 8.3: Evolution of proportions of strategies in population after generational updating with initial proportion of 50% straightforward agents, 50% parking agents.)

these two different types of agents in terms of parking, can lead to better results for both.

8.1.2 Result 2 : *Increasing Payoff*

For our strategies to increase payoff, the results were harder to interpret. For shaving, there was a single equilibrium with incumbents shaving, while entrants preferred the straightforward results. For demand reduction, the results were equally divisive - incumbents had two equilibria, while entrants were quite indifferent to either strategy.

Shaving

We were quite surprised by the results for shaving. Indeed, shaving normally is profitable for small agents for whom free-riding upon an agent with complementing

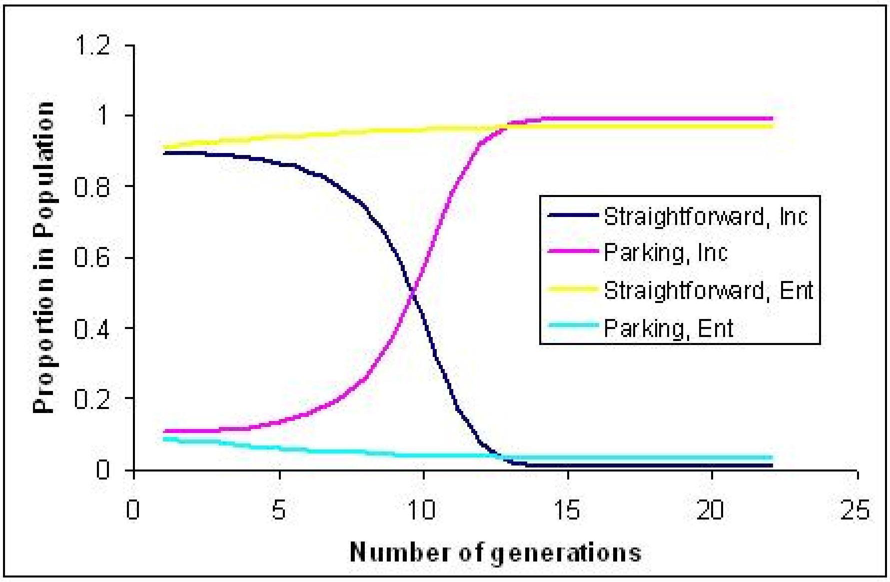

Figure 8.4: Evolution of proportions of strategies in population after generational updating with initial proportion of 100% straightforward agents.)

valuations is possible to outbid a large agent. In our case, however, incumbents (agents with more valuations in general), seemed to do better than entrants. This is illustrated in figures 8.6-8.8.

One explanation to this result would be that the size of a bundle does not matter as much as two agents' potential for complements. In our case, because incumbents stick to their own schedules, they have more potential for complementing each other on the goods desired than the entrants. As such, it is more advantageous for an incumbent agent to deviate from the straightforward strategy to shave and force another incumbent agent with complementing valuations to win the goods than for the entrant agent to do so. Instead, entrant agents may benefit from consistently bidding their true valuations stealing the items that incumbents would win, had they bidded truthfully.

One cause of concern of this model is the reduction in general system payoff at the equilibrium. The race to shaving on the incumbents' side may lead to inefficient

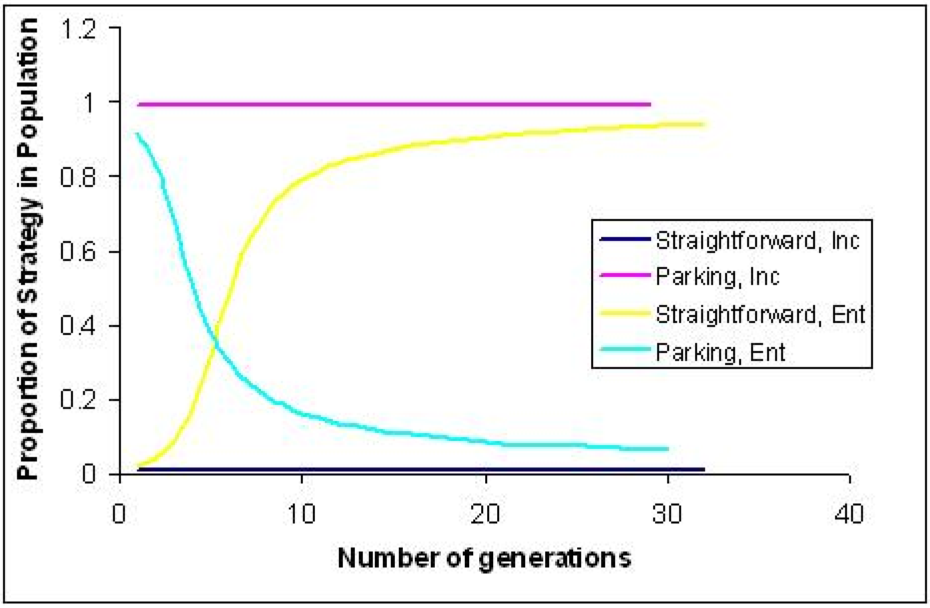

Figure 8.5: Evolution of proportions of strategies in population after generational updating with initial proportion of 100% parking agents.)

allocations in the auction. Indeed, while agent payoff for straightforward strategy in the experiments performed was 7000 for incumbents, and 6800 for entrants, the equilibrium payoff was reduced for both at 6500 for incumbents, and 5900 for entrants.

Demand Reduction

For demand reduction, the entrants were mostly indifferent to either strategy - playing a mixed equilibrium hovering at around 50/50. The incumbents, on the other hand, had two pure equilibria again, but with a greater probability of playing the heuristic than not. This can be seen in figures 8.9-8.11 figures indicating the equilibria with different values for original population proportions.

Part of these results may be attributed to the conflicting goals of flexibility preserving and demand reduction. Indeed, we've measured above the equilibrium for

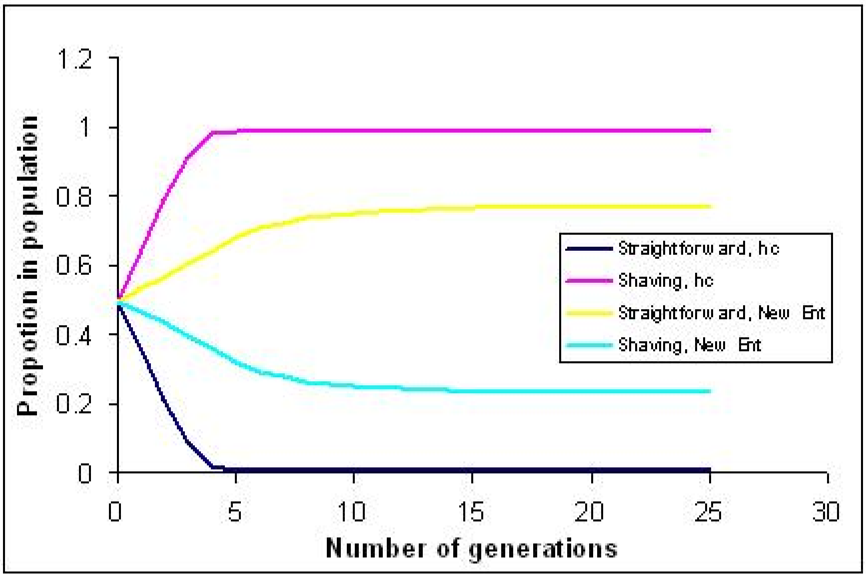

Figure 8.6: Evolution of proportions of strategies in population after generational updating with initial proportions of 50% straightforward agents, 50% shaving agents.)

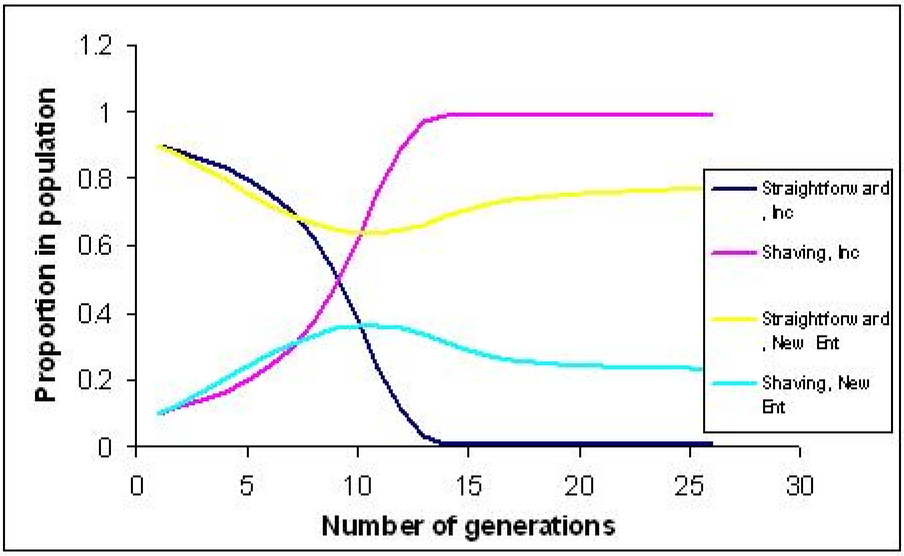

Figure 8.7: Evolution of proportions of strategies in population after generational updating with initial proportion of 90% straightforward agents.)

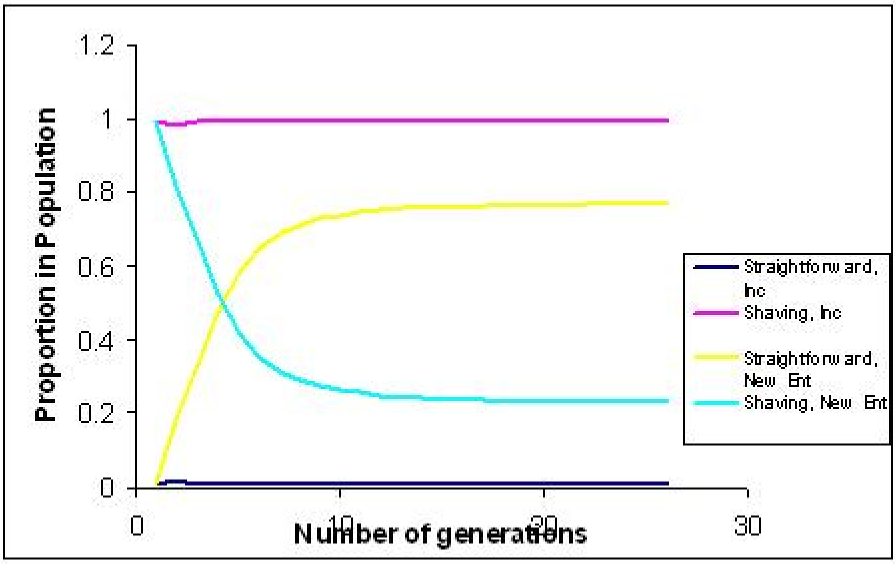

Figure 8.8: Evolution of proportions of strategies in population after generational updating with initial proportion of 100% shaving agents.)

using overdemand as a proxy for flexibility preserving, leading to the same two equilibria for incumbents as in this case. It is interesting, however, in either case, that the equilibria involve playing the same strategy as everyone else.

The difference in results between the incumbents and the entrants can again be attributed to their valuation structures. Indeed, there is an existing niche for the incumbents designed in the auction - incumbents had preferences mostly for complements of one another since they wanted to stick to their original schedule. For the entrants, the advantage in trying to demand reduce is less obvious - their valuation bundles may intersect with many other agents' preferences and ways to avoid confrontation are fairly complex. In this case, entrants may benefit instead from bidding their true value, picking up items that incumbents lost due to shaving.

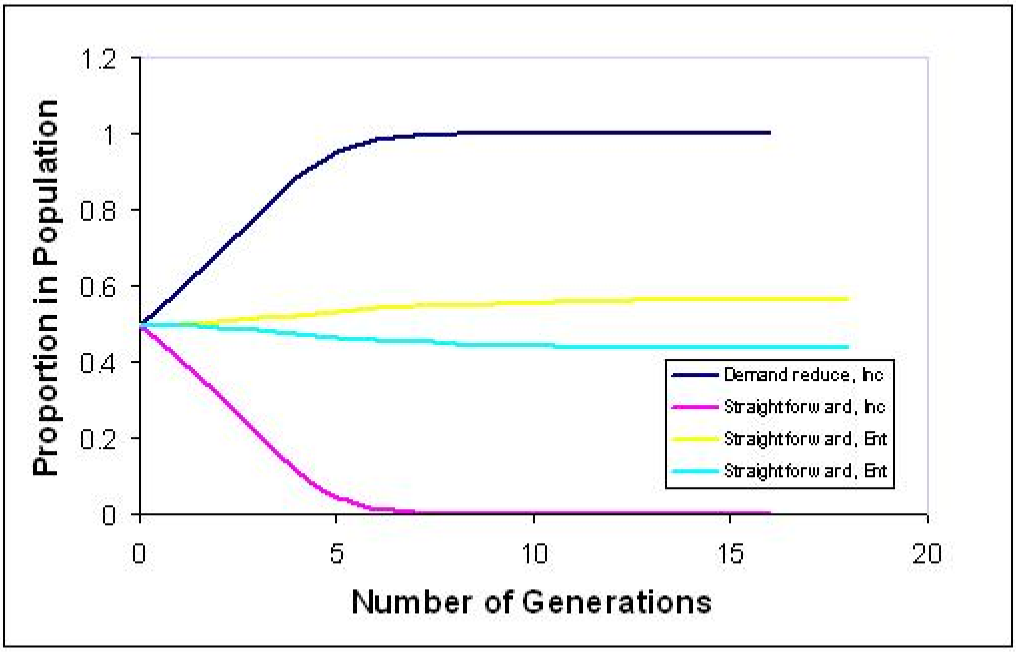

Figure 8.9: Evolution of proportions of strategies in population after generational updating with initial proportion of 50% straightforward, 50% demand reducing agents.)

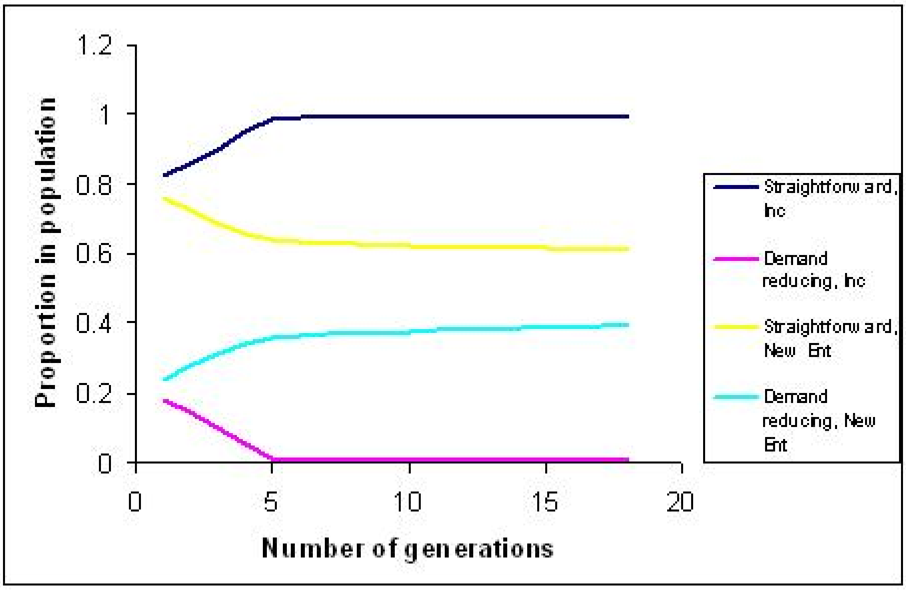

Figure 8.10: Evolution of proportions of strategies in population after generational updating with initial proportions of 20% demand reduction, 80% straightforward.)

8.2 Vulnerabilities of the Auction

There are strong indications through our work that bidding myopically is not the equilibrium strategy in this auction, and may not lead to the equilibrium result. Although any result from these experiments is tied also closely with the quality of the valuation model, there are also causes for general concern.

First, the problem from the interaction between the budget constraint and the activity rule discussed in chapter 4 is mostly confirmed in the experiments we ran. As shown in chapter 6, payoffs in the auction improved as price prediction became more accurate. The same happened to the problem of undersell.

Moreover, results from running heuristic strategies for preserving flexibility indicate that preserving flexibility is, generally, an advantageous strategy to play in the game. This was especially true for agents who have more options of packages to consider, less competitive, and therefore more likely to need to switch business plans

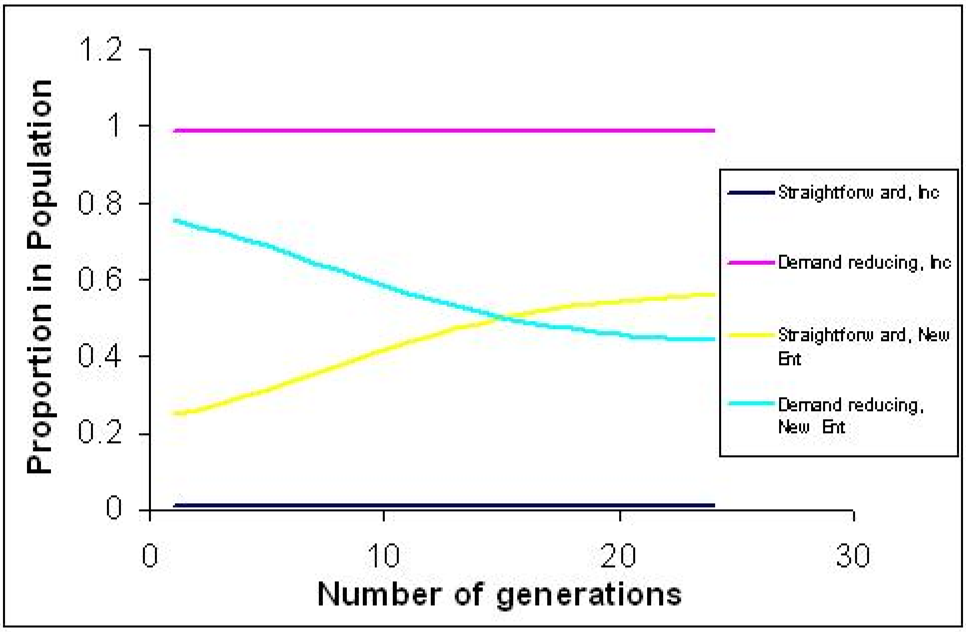

Figure 8.11: Evolution of proportions of strategies in population after generational updating with initial proportions of 100% strategizing agents.)

within the game (the incumbents, in our study). The profile of these agents fits well with the characteristics we'd expect to find within agents for whom the activity rule is more likely to be constraining.

Another problem occurs purely as the undersell problem due to linear pricing. It is possible that the problem of undersell is worsened by our bounded rational model which only considers a finite number of business plans. Further studies should test this framework we developed on variations of the valuation model. However, no strategy we explored so far has been able to correct for the undersell, this is a problem that the auction design needs to take more heed to. Indeed, at 5.2%, the undersell is quite significant, especially in the context of rare resources such as the LGA air slots.

Finally, we found that straightforward bidding, even with budget prediction, leads to a suboptimal outcome, with the expected total utility of the system reduced (figures 8.12 and 8.13). However, the Nash equilibrium outcome for strategies that aimed to preserve flexibility (*OverdemandMore* and *Parking*) was found to improve on the expected utility. Expected total utility in the equilibrium for *Demand reduction* was found to even slightly surpass the utility of the expected optimum. However, we did not verify here whether the items followed the optimum allocation, and the increase could potentially come from decreases in prices. We also found that shaving could reduce the total utility from the straightforward strategy with budget prediction. This can be a cause of concern as shaving was a Nash equilibrium for the incumbents, even as it reduced their expected utility by 9%.

We recognize, however, that our heuristic strategy experiments are far from complete - using the parametrized function to make compete together various degrees of shaving would yield a more robust and complete answer. Moreover, the number of games played to calculate the payoff matrices were still quite limited in context. Yet, the results do give an indication of trends and causes of concern for the combinatorial clock auction.

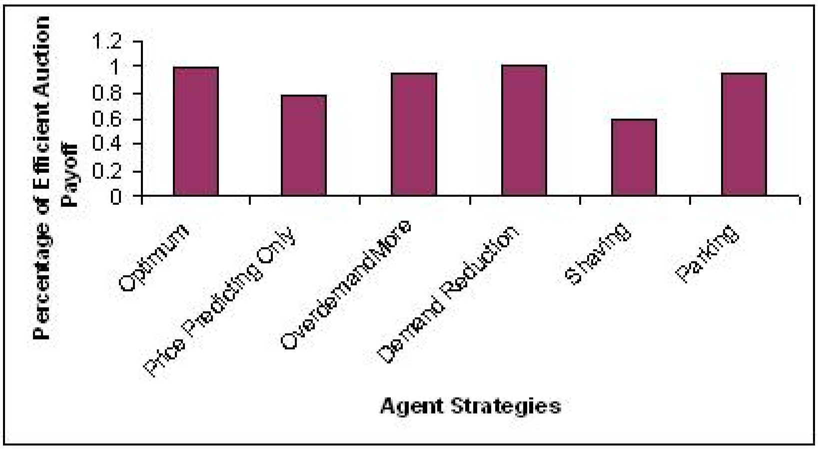

Figure 8.12: Comparison of expected system payoffs under various strategies. (Optimum indicates efficient allocation, normalied at 100% or 1.0. Results from Nash equilibrium of Heuristic strategies presented here.)

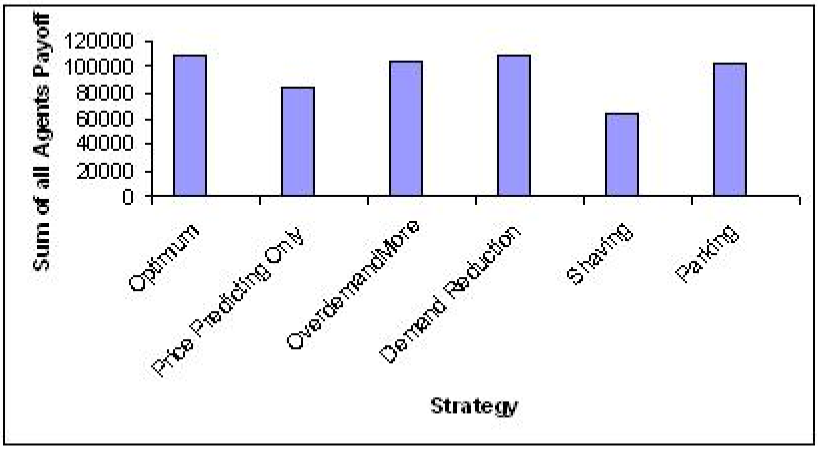

Figure 8.13: Comparison of total system payoffs under various strategies. (Optimum indicates efficient allocation, results from Nash equilibrium of Heuristic strategies presented here.)

Chapter 9

Conclusion

In this paper, we provide an agent-based analysis of the currently proposed combinatorial clock auction of LaGuardia airport. Our framework is inspired from research results from various sources. In our general agent-based framework design, we took ideas developed in response to the trading agent problem [10, 29, 14, 34]. We also extend previous work on FAA valuation models in application to LaGuardia airport[19]. To model our agents, we provide adaptations of the supervised learning techniques developed in TAC to the LGA problem [29, 14], introducing their use in predicting budget constraints in the LGA case. Then, heuristic strategies are developed from observed agent behavior in auctions [3, 23, 33]. Finally, we extend Wellman's [10] description of finding equilibria in homogeneous populations with restricted strategy spaces to a heterogeneous population in the combinatorial auction framework.

Our analysis of the model provided results on different levels of the model. One major problem found in the model relates budget constraints and the proposed activity rule. In particular, we showed, both experimentally and theoretically, that agents would be forbidden from bidding their honest valuation by the activity rule if they are subject to budget constraints. Accordingly, agents who have better information of the market (who price predict better) have greater expected value in joining the auction. Moreover, heuristic strategies to preserve flexibility also help in raising the utility generated by the market.

Another potential problem is in terms of undersell. Undersell is probably exacerbated by the budget/activity rule problem such that markets with better informed agent could expect less undersell. However, the percentage of undersell, even after the completion of our staged price learning, was still important, at 5.2%.

In terms of agent strategy, we found that many of the observed agent strategies in past auction designs were still possible equilibrium outcomes for the agents competing in the combinatorial clock auction. The results here, however, are a bit confounding and differ according to characteristics of the agent valuation model. We proposed from the observed results that strategies such as *shaving* and *demand reduction* may be affected by how complementary agent valuations are within the game. Moreover, benefits to *parking* are also probably affected by the flexibility within agent valuation.

From this study, we feel that agent-based analysis of the auction can be useful in generating results that have been difficult to generate in live experiments. In our case, this was caused by two factors. First, due to industry secrets, airline executives may be reticent in showing their true strategies in pre-game simulations. This was already observed in the first simulation in February 2005. Second, due to the complexity in managing package bidding, controlled experiments with random subjects are hard to implement with complex valuation models. This last problem is inherent in combinatorial auctions which are difficult to solve. In these cases, automated agent-based analysis can be useful in generating experimental results.

Future Work

This paper leaves multiple directions to pursue related agent-based analysis of mechanism designs. We list here some possibilities:

- Due to the large number of entries to populate in a payoff matrix, we did not have the time to explore restricted equilibriums on more than two strategies. Yet, we offered parametrized heuristics for three large classes of strategies. At a future date, experiences should be made on varying these parameters to calculate the optimal α in each case. This is especially important for the shaving strategy.

- The work here is performed based on a single valuation model. However, we found that the optimal agent strategies is much dependent on characteristics of the valuation model. Future work should test our results on alternate valuation models.
- Designing or extending new heuristic strategies to test the model.
- Extending our agent-based analysis to other forms of auctions, notably full combinatorial auctions with winner determination.
- The study of interaction between budget constraints and other possible activity rules.

We hope that our work will encourage further work in the agent-based analysis of auctions.

Bibliography

[1] Congressional testimony, 2000.

[2] The passengers bill of rights game: Congestion management at lga, 2004. Game Handbook for November 3,4,5,2004 at George Mason University.

[3] L. Ausubel and P. Cramton. Demand reduction and inefficiency in multi-unit auctions. 2002. mimeo, University of Maryland.

[4] Lawrence M. Ausubel, Peter Cramton, and Paul Milgrom. The clock-proxy auction: A practical combinatorial auction design. Technical Report 04mit5, University of Maryland, Department of Economics - Peter Cramton, 2004.

[5] Cynthia Barnhart and Pavithra Harsha. Current working project at the Massachusetts Institute of Technology, 2005.

[6] Jean-Pierre Benoit and Vijay Krishna. Multiple-object auctions with budget constrained bidders. *Review of Economic Studies*, 68(1):155–79, 2001. available at http://ideas.repec.org/a/bla/restud/v68y2001i1p155-79.html.

[7] Peter Cramton. The efficiency of the fcc spectrum auctions. Technical Report 98jlefcc, University of Maryland, Department of Economics - Peter Cramton, 1998. available at http://ideas.repec.org/p/pcc/pccumd/98jlefcc.html.

[8] Peter Cramton and Jesse Schwartz. Collusive bidding in the fcc spectrum auctions. Technical Report 1210, Econometric Society, August 2000.

[9] Peter C Cramton. Money out of thin air: The nationwide narrowband pcs auction. *Journal of Economics and Management Strategy*, 4(2):267–343, 1995.

[10] Jeffrey K. Mackie-Mason Anne Osepayshvili Daniel M. Reeves, Michael P. Wellman. Exploring bidding strategies for market-based scheduling.

[11] Allan Sumiyama Elaine Ou, Jeff Shneidman. An imperfect valuation model for airport landing slots. class project, CS286r, Harvard College, 2004.

[12] Wedad Elmaghraby. Auctions and pricing in e-marketplaces. forthcoming in the Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era, 2004.

[13] David C. Parkes et al. Ice: An iterative combinatorial exchange.

[14] Amy Greenwald. The 2002 trading agent competition: an overview of agent strategies. *AI Mag.*, 24(1):83–91, 2003.

[15] Amy Greenwald and Justin Boyan. Bidding under uncertainty: Theory and experiments. *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, 2004.

[16] Legal Information Institute. Us code, 2004. http://www4.law.cornell.edu/uscode/.

[17] John H. Kagel and Alvin e. Roth. *The Handbook of Experimental Economics*. Princeton University Press, 1995.

[18] Lawrence M. An efficient ascending-bid auction for multiple objects. mimeo, University of Maryland, 1995.

[19] K.Hoffman. M. Ball, G.Donohue. Auctions for the safe, efficient and equitable allocation of airspace system resources. In P.Crampton et al., editor, Combinatorial Auctions, Ch.22.

[20] Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2), 2000.

[21] Noam Nisan. Bidding and allocation in combinatorial auctions. *Proceedings of the 2nd ACM conference on Electronic commerce*, 2000.

[22] The Port Authority of New York and New Jersey. September 2004 traffic report.

[23] Ryan Porter and Yoav Shoham. On cheating in sealed-bid auctions. Stanford University, 2004.

[24] S.J. Rassenti. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13(2), 1982.

[25] Federal Register. Notice of alternative policy options for managing capacity at laguardia airport and proposed extension of the lottery allocation. Technical report.

[26] P. Schuster and K. Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 1983.

[27] Elizabeth Sklar Steve Phelps, Simon Parsons. Applying genetic programming to economic mechanism design: Evolving a pricing rule for a continuous double auction. A 2-page poster of this paper appears in the Proceedings of AAMAS '03, 2003.

[28] Peter McBurney Steve Phelps, Simon Parsons. Automated trading agents verses virtual humans: An evolutionary game-theoretic comparison. To appear in AMEC VI, 2004.

[29] Peter Stone, Robert E. Schapire, Michael L. Littman, János A. Csirik, and David McAllester. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research*, 19:209–242, 2003.

[30] P. G. Taylor and L. Jonker. Evolutionary stable strategies and game dynamics. *Math. Biosci.*, (40), 1976.

[31] Gerald Tesauro and Rajarshi Das. High-performance bidding agents for the continuous double auction. *Proceedings of the International Joint Conference on Artificial Intelligence*, 2001.

[32] Hal R. Varian. Economic mechanism design for computerized agents. 2000.

[33] Robert J. Weber. Making more from less: Strategic demand reduction in the fcc spectrum auctions. *Journal of Economics and Management Strategy*, 1997.

[34] M. Wellman, P. Wurman, K. O'Malley, R. Bangera, S. Lin, D. Reeves, and W. Walsh. A trading agent competition, 2000.