



Modeling Carrier Behavior in Sequential Auction Transportation Markets


M. A. Figliozzi, University of Maryland

H. S. Mahmassani, University of Maryland

P. Jaillet, Massachusetts Institute of Technology

Conference paper

Session XXX



Moving through nets:

The physical and social dimensions of travel

10th International Conference on Travel Behaviour Research

Lucerne, 10-15. August 2003

Title: Modeling Carrier Behavior in Sequential Auction Transportation Markets

Miguel A. Figliozi¹, Hani S. Mahmassani¹ and Patrick Jaillet²

¹Department of Civil & Environmental Engineering
University of Maryland, College Park, Martin Hall, College Park, MD 20742, USA

²Department of Civil and Environmental Engineering,
Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Phone: 301-405-0752

Fax: 301-405-2585

eMail: masmah@umd.edu

Abstract

Online markets for transportation services, in the form of Internet sites that dynamically match shipments (shippers' demand) and transportation capacity (carriers' offer) through auction mechanisms are changing the traditional structure of transportation markets. A general framework for the study of carriers' behavior in a sequential auction transportation marketplace is provided. The unique characteristics of these marketplaces and the sources of difficulty in analyzing the behavior of these marketplaces are discussed. Learning and behavior in a sequential Vickrey auction marketplace is analyzed and simulated. Some results and the overall behavioral framework are also discussed

Keywords

Freight Transportation, Carrier Behavior, Game Theory, Bounded Rationality, Auctions, Carrier Management Strategies, Information Technology, Electronic Commerce

International Conference on Travel Behaviour Research, IATBR

Preferred citation

Figliozi, M.A., Mahmassani, H.S. and Jaillet, P. (2003) Modeling Carrier Behavior in Sequential Auction Transportation Markets, paper presented at the 10th International Conference on Travel Behaviour Research, Lucerne, August 2003.

1. Introduction

Online markets for transportation services, in the form of Internet sites that dynamically match shipments (shippers' demand) and transportation capacity (carriers' offer) through auction mechanisms are changing the traditional structure of transportation markets. Beyond changes in market structure, Internet auctions have emerged as an effective catalyst to sell/buy through electronic marketplaces. Transaction time, cost, and effort could be dramatically reduced, creating new markets and connecting buyers and sellers in ways that were not previously possible (Lucking-Reiley and Spulber, 2001).

McAfee and McMillan (1987) define auctions as market institutions with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants. Two types of resources could be traded in transportation marketplaces: (a) loads, or demands of shippers, being "sold" to the lowest bidder-- this would be the case of extra supply looking for scarce demands; and (b) capacity, i.e. the capacity to move goods, by a given mode from location A to location B, being "sold" to the highest bidder. The buyer of such capacity could be a shipper wishing to move a load, a carrier needing the extra capacity to move contracted loads, or a third party hoping to make a profit by reselling this capacity.

The focus of this paper is the study of transportation marketplaces that enable the sale of cargo capacity based mainly on price (case a), yet can still satisfy the customer's level of service demands. Specifically, the paper considers the reverse auction format (also known as procurement auctions), where shippers post loads, triggering carrier bids. The market is comprised of shippers that independently call for shipment procurement auctions and the carriers that participate in them (we assume that the probability of two auctions being called at the same time is zero). Auctions are performed one at a time as shipments arrive to the auction market. The market generates a sequence of auctions (procurement, bidding, and auction resolution) that take place in real time, thereby precluding the option of bidding on two auctions simultaneously. Markets where carriers bid on configurable bundles of loads give rise to combinatorial auctions. The behavioral aspects of auction market behavior are more readily articulated without the added complexity of the combinatorial aspect.

In auction markets, prices are not negotiated; they are generated as the outcome of carrier bids and a predefined set of rules. These rules precisely define a strategic environment, therefore allowing the study and analysis of carriers' behavior (expressed through bids). As such, auctions provide a useful laboratory to gain insight into carriers' behavior in a freight market.

Auction-based electronic marketplaces give rise to new dimensions in the behavior of the principal freight transportation decision agents, especially with regard to learning and adaptation in a competitive bidding environment. While the area of freight demand, and the underlying behavioral dimensions, have received limited attention in the travel behavior research community (Mahmassani, 2001), behavioral considerations play a critical role in

determining the performance of auction-based electronic freight markets, and the policy implications of different marketplace rules and regulatory requirements. Furthermore, the behavioral dimensions at play in electronic freight markets are examples of more general behavior mechanisms in competitive decision situations that extend beyond the realm of freight transportation (e.g. airline schedule and pricing decisions).

This paper has six sections. The next section identifies the characteristics of transportation auctions and associated sources of complexity. Section 3 articulates a framework to study carrier behavior, and the learning processes important in this environment. Section 4 analyzes a market with Vickrey auctions and different degrees of information. Section 5 ends the paper with conclusions and possibilities for future research.

2. Characteristics of Transportation Auctions

From the carrier point of view, the cost and value of transportation services are hard to quantify. The value of a traded item (shipment) may be strongly dependent upon the acquisition of other items (e.g. nearby shipments). In addition, the value of a shipment is related to the current spatial and temporal deployment of the fleet. The geographic dispersion of both demand and supply, uncertain demand arrival rates, and realizations over time and space, contribute to a dynamic and stochastic environment. These factors further increase the uncertainty of a shipment's cost and value.

2.1. Complexity of Transportation Auctions

Auctions, as a device to match supply and demand, provide a powerful mechanism to allocate resources, especially when the latter have uncertain or non-standard value. Auction analysis can be quite challenging, especially in a stochastic setting such as transportation. In this kind of setting, carriers face a complicated decision problem, which stems partly from the strategic inter-dependency among competitors' decisions, costs, and profits. Auctions have been widely studied by economists, leading to recent advances in the theoretical understanding of different auction types and designs. These models have been mainly focused on one-time auctions with symmetric risk-neutral agents that bid competitively for a single or multiple units. Optimal bidding strategies have been found in many auction environments, however the case of *sequential auctions*, with bidders with multiunit demand/supply curves, remains intractable (Krishna, 2002). Another source of complexity arises from the need to solve fleet management problems (vehicle routing problems with time windows, penalties, etc) to obtain the cost information for a shipment. These are NP hard problems, which cannot generally be solved optimally for realistic fleet sizes in a dynamic and stochastic environment.

2.2. Carrier Decisions

A carrier participating in a sequential auction market is required to make decisions. Bidding decisions, for example, require a carrier's choice, and therefore reveal information about a carrier's decision making process. We define a *strategic decision* as the investment of resources, for the purpose of learning about or influencing competitors, to improve profits by manipulating future auction outcomes. Strategic decisions can be sub-classified as *identifying* or *signaling* decisions. *Identifying* decisions are characterized by attempts to identify or discover a competitor's behavior. *Signals* are decisions that aim to establish a reputation or status for the carrier. The importance of *signaling* has been recognized in game theory; a class of games called "*signaling games*" studies the strategic interaction between asymmetrically informed players. This class of games usually has two players; the first mover is endowed with private information, and the second mover observes what action (or signal) has been taken (Cho and Kreps, 1987).

We define *operating decisions* as decisions that are not strategic but aim to improve a carrier's profit level. This type of decision, for example, includes the estimation of a shipment value, the rerouting of the fleet after a successful bid, the reaction to unexpected increase in travel times, etc. Bids can be the result of either strategic or operating decisions, depending on the complexity of the marketplace, the information available, and the rationality of the carrier. This idea is developed further in section 4. *Carrier Behavior* is defined as a sequence of bids taken by a carrier. *Behavioral Type* denotes a process that, given information regarding the internal state (of the carrier) and the strategic environment (current and past), will output an action (a bid).

2.3. Finding an "Optimal" Bidding Policy

In an auction, profits are highly dependent on the quality of the bidding policy. A bidding policy is a function that uses information about the state of the carrier, the characteristics of the shipment for auction, the marginal cost of serving the shipment, auction type, and beliefs about the competitors and environment, to produce a bid value. A bidding policy is similar to a revenue management policy; however it is ultimately more general in that a bidding policy uses beliefs about competitors, not simply relying on demand and internal state data.

There are many "reasonable" bidding policies, however the sources of complexity discussed above generally preclude finding a policy that "optimizes" the entire auction/assignment problem. In some settings, it may even be difficult to find a policy that works "reasonably well". In some special cases, as the one analyzed in section 4, an optimal bidding policy can be found. However, in general, carriers have the opportunity of *learning* about the environment and about other players with each auction. Learning is crucial as it can influence behavior; this issue is analyzed in the next section.

3. Carrier Behavior Framework and Bounded Rationality

This section presents a framework to articulate the behavioral processes followed by carriers in the context of interactive auction-based marketplaces; the role of learning and possible mechanisms are discussed in this context. In this discussion, carriers are assumed to be, inherently, profit maximizers, whose goal is to achieve higher profit levels. As such, they prefer outcomes that yield higher expected profits. If two outcomes yield the same expected profit, carriers will prefer outcomes that yield a higher expected market share. However, as noted, the complex nature of the decision environment (stochastic, dynamic, partial information, combinatorial problem, limited cognition) will often preclude optimal decisions (ex-post). Carriers can be modeled as intelligent agents that determine their interaction with other agents and with their environment on the basis of history (experience), learning, expectations of future consequences of current actions, and evolving strategies.

3.1. Framework for Carrier Behavior

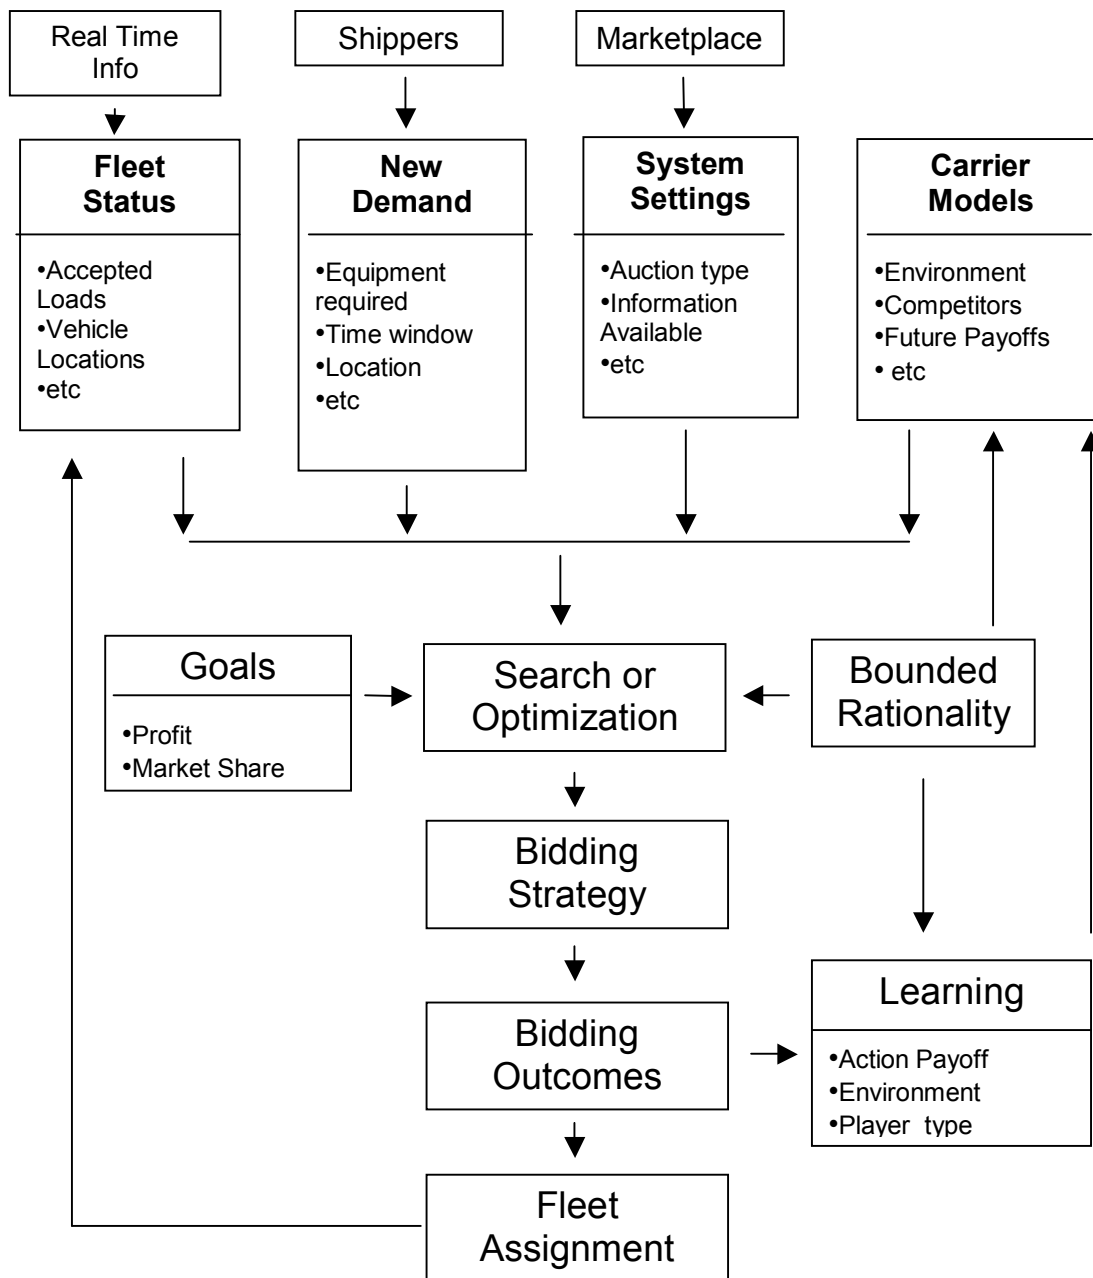
Figure 1, adapted from Figliozzi et al. (2003), is a schematic representation of the processes and interaction sources within the context of a sequential auction marketplace. A shipper's decision to post a shipment initiates an auction. The principal events in the market are the arrival of shipments, the subsequent bidding process, and bid resolution. Shippers are viewed here as agents that generate a stream of shipments, with corresponding attributes, according to predetermined probability distribution functions. They are rational agents because they know the exact value of the reservation price of their shipments as a function of its attributes (origin-destination, commodity type, stock out costs, time window, etc.). Furthermore, shippers maximize profits by setting the right reservation price, which is the highest price a shipper is willing to pay a carrier for serving a given shipment. The shipper achieves a profit (saving) when paying less than the reservation price. A rational shipper rejects transportation services exceeding the reservation price (the shipper does not incur loss).

Carriers will attempt to maximize profits or gain market power by adapting their behavior in response to interactions with other carriers and with their environment. Further, they will act according to the constraints and the physical feasibility specified by their (fleet) assignment strategies and pool of awarded shipments (that remain to be served). Past decisions are binding and limit the future actions of carriers, therefore behavioral rules are state-conditioned and the carriers co-adapt their behavior as the marketplace evolves over time.

In this framework, carriers' beliefs, experiences, and knowledge about other competitors' behavioral types evolve jointly over time; their strategies at a given moment are contingent on interactions that have occurred in a path-dependent time line. Carriers' internal events are the assignment, pickup, and delivery of loads, mostly operational decisions. Carriers repeatedly engage in bidding interactions, modeled as non-cooperative games. However, these repeated bidding interactions are also the only means of communication for a carrier to convey "signals" or "reputations" to the other competitors.

The sources of complexity and limitations alluded to previously affect how a carrier models, evaluates, and optimizes his action, as indicated by the arrows in figure 1. These limitations generally result in the bounded rationality (Simon, 1972) of the actors in their decision process, as discussed next.

Figure 1 Carrier behavior in a sequential auction marketplace



3.2. Bounded Rationality

In a sequential auction market environment, each carrier is aware (though with incomplete and imperfect knowledge) that his actions have significant impact upon his rival's payoffs, and vice-versa. This interdependence is naturally modeled using tools based on game theory. Player rationality, unbounded computational resources, and shared knowledge are assumptions necessary to sustain a Nash equilibrium. These assumptions are unlikely to hold in our marketplace.

As mentioned in Section 2, the complexity of the game is such that a player's insight about the effect of his own actions in future games is vague. In general, players largely lack the quantitative knowledge of how their actions may affect competitors' future play. It is difficult to form and sustain shared knowledge when a player has no clear understanding of how his actions affect his own and competitors' future rewards. Therefore, market participants can be thought of as boundedly rational agents (Simon, 1972), not "fully rational" agents as classically assumed in game theory.

The repeated interaction among carriers allows the possibility of learning about strategies, the environment, and the competition. Carriers can analyze with *different degrees of sophistication* (bounded rationality) the history of play and estimate the possible future consequences of current actions. From the strategic point of view, cognitive and computational limitations will be evidenced in these areas:

- i. *Identification*: the carrier has limited ability to discover competitors' behavioral types, which may require complex econometric techniques;
- ii. *Signaling*: limited ability to "read" or "send" signals that convey a reputation
- iii. *Memory*: limited ability to record and keep past outcome information or memory to simulate all future possible paths in the decision tree
- iv. *Optimization*: even if carriers could identify competitors' behavior, their ability to formulate and solve stochastic optimization problems is likely limited.

Strategic play under boundedly rational behavior has not been sufficiently addressed by researchers in game theory, and definitely not in the context of freight electronic marketplaces. Research in the area of learning in games is actively seeking to explain how agents acquire, process, evaluate or search for information.

3.3. Market Information Levels and Learning

The information revealed after each auction can influence the nature and rate of the process by which carriers learn about the "game" and their competitors' behavior. This information includes:

- i. Actions (bids) placed.
- ii. Number of players (carriers) participating

- iii. Links (name) between carriers and bids
- iv. Individual characteristics of carriers (e.g. fleet size)
- v. Payoffs
- vi. Knowledge about who knows what, information asymmetries, or shared knowledge about previous items.

In the freight marketplace considered, revealed payoff information consists of revenue, not profit, because the cost of providing service usually depends on private information about each carrier's fleet status. In a *maximum information* environment, all the above information is revealed. On the other hand, an environment where no information is revealed is a *minimum information* environment. These two extremes approximate two realistic situations. Maximum information would correspond to a real time internet auction where all auction information is accessed by participants. Minimum information would correspond to a shipper telephoning carriers for a quote, and calling back only the selected carrier.

Learning implies that there is information or expertise to be acquired. In a transportation marketplace, the main object of learning would be to identify an optimal bidding policy, or at least one that works reasonably well most of the time. Methods that can be used for learning depend on the level of information available. A brief discussion of learning techniques that can be applied to the two above extremes of the information spectrum is presented here. The discussion is not intended to be an exhaustive survey; rather it presents the main ideas and intuition behind the most prevalent methods.

3.3.1. Minimum information setting

Since knowledge of the payoff structure and competitors' behavior is extremely limited (or even non-existent), from a single carrier's perspective the situation is modeled as a game against nature; each strategy has some random payoff about which the carrier has no a priori knowledge. Learning in this situation is the process of moving (in the action space) in a direction of higher profit. Experimentation (trial and error) is necessary to identify good and bad directions. Learning methods commonly used in this setting are:

- i. Genetic Algorithms: the adoption of a strategy based upon its recent success. There is experimentation, with small probability successful strategies "mutate". Good strategies are retained to be used as future "seeds", bad strategies "die" (Goldberg, 1989).
- ii. Reinforcement Learning: based on psychologists' models of learning. Strategy choice is stochastic; the probability of the agent's action is based on how well that action has performed in the past. The better the past performance, the higher the probability of being played. (Roth and Erev, 1995)

3.3.2. Maximum information setting

In this setting all information is fully revealed and can be used with different degrees of sophistication to model competitors' behaviors. The learning literature has mostly focused on this case. Some learning methods that model other players are:

- i. *Fictitious Play*: agents use a backward-looking procedure to estimate the frequency of play of other players (probability of bidding now given by the frequency of past bids) Agents select an optimal response to this past frequency of play (Brown, 195; Monderer and Shapley, 1996)
- ii. *Machine Learning*: each player is modeled as a machine with a finite number of states. Each machine carries in it a rough model of what the other machines may do (Binmore, 1987)
- iii. *Rationalizable and Bayesian Learning*: players form beliefs about their environment and other players. In simple games, this type of learning, starting with "reasonable beliefs", converges to a Nash Equilibrium. Here, initial beliefs play an important role (Milgrom and Roberts, 1990; Kalai and Lehrer, 1993)
- iv. *Rule Learning*: players have different levels of internal simulation; having this in mind they model other players' intentions and actions. This is an example of cognitive process where players truncate an internal simulation of the model of the other players (Stahl, 1995).

A common characteristic of the above-mentioned learning methods is that their properties have been studied primarily in simple games, which are quite different from the complex environment of a transportation marketplace. No learning method adequately describes how to "identify" other players. Methods that use internal simulation of other players (e.g. rule learning, rationalizable) assume players have a determinate behavior type like Nash equilibrium players, or that they can rationalize. These methods neither take into account how players come up with beliefs or models about other players, nor analyze the difficulty or cost of coming up with such models.

Another learning method that is not systematic is called "Rules of Thumb" learning. This method is not systematic in the sense that it lacks continuous updating of beliefs or population types; rather a player must use different rules of thumb that are reactive to the history of play. These actions may also convey a strategic decision of *signaling* a reputation (to other players). The difficult issue in this method is how to design or select the right "rule of thumb" for a given problem. A classical rule of thumb (and robust in many games) is *Tit for Tat*. In our marketplace, with two carriers (A and B, where carrier B plays Tit for Tat), this rule of thumb can be defined as:

- i. Carrier B computes the average bid value of carrier A over the last T auctions, called \bar{a}
- ii. Carrier B obtains $\alpha = \text{maximum} \{ \bar{a} / mc, 1 \}$, where mc is B's average marginal cost in the last T periods.
- iii. B's next bid will be equal to his marginal cost multiplied by α

This method of bidding sends a signal that conveys: “whatever you bid against me, I will bid against you later”. Researchers began to pay serious attention to this method of rule learning after the experimental work of Axelrod (1984), who asked experts to prepare computer programs for playing the repeated prisoner’s dilemma. The winner of the experiment was Tit for Tat. Despite its plain character, Tit for Tat demonstrated that simplicity and signaling could be of major importance in a boundedly rational environment. The next section considers the Tit for Tat rule and other learning methods in a Vickrey sequential auction environment.

4. Analyzing a Second Price Auction Marketplace

This section examines the behavior two carriers competing in a sequential auction marketplace. The auction format is a 2nd price auction, also known as a Vickrey auction (Krishna, 2002). First is a description and analysis of Vickrey auction properties. The second part focuses on bounded rationality demands of this auction type. The third part describes learning and simulation results in markets with different degrees of information. The market setting used is the same as described in Figliozzi et al. (2003), and summarized in appendix A. Although the market setting is simple and stylized, it provides the necessary features that capture the most important stochastic elements of the problem: stochasticity of reservation prices, origins and destinations, time windows, and competitors’ bids and costs. The simulation results and analysis are applicable only to situations with *low demand arrival rates* (low enough to keep any carrier from getting close to capacity). In conclusion, we analyze this marketplace from a strategic perspective, mentioning some policy implications of the findings.

4.1. Second Price Auctions (Vickrey Auctions)

This is a one shot (or period) auction that takes place as follows:

- i. Each agent (carrier) submits a single sealed bid
- ii. The winner is the agent with the lowest bid
- iii. The item (shipment) is awarded to the winner
- iv. The winner pays the value of the second lowest bid.
- v. The other agents do not win, pay, or receive anything

A powerful characteristic of (one shot) Vickrey auctions is that the optimal strategy is simply to bid the true cost of serving the shipment. A Vickrey auction guarantees that the shipment will always be awarded to the carrier with the lowest service cost for that shipment (assuming carriers that are profit maximizers at each bid instance, that there is no look ahead, i.e.

ignoring the future) regardless of the participants' beliefs about the shippers and other carriers.

The proof is strikingly simple and was first stated by Vickrey (1961). Overbidding (bidding over the marginal cost) will decrease the probability of winning while keeping the payment constant (2nd price does not depend on the winner). Underbidding (bidding below the marginal cost) will increase the probability of winning only for those situations where another agent has a lower cost; therefore, winning in those situations brings about a loss (the winner's payment is less than his/her cost).

4.2. Vickrey Auctions and Bounded Rationality

The assumption that carriers are profit maximizers at each bid instance allows for a simple and elegant treatment of carrier strategies since they render the tracking or updating of carrier beliefs about the shippers or competitors essentially irrelevant. The "bounded rationality" requirements for this type of behavior are quite modest:

- i. Memory-wise, the carrier is not required to store any data about past actions and outcomes;
- ii. Modeling-wise, the carrier does not have to identify competitors' "type";
- iii. Adequate "introspection" and rational thinking is needed to understand that bidding over or below the marginal cost is not optimal (optimal in a myopic way).
- iv. Introspection might be considered optional. What it is required is just "knowledge" about how 2nd price auctions work. Such knowledge is readily acquired by any carrier that hires an auction consultant or by a manager that can read and understand a basic auction theory book.

4.3. Learning and Simulation Results

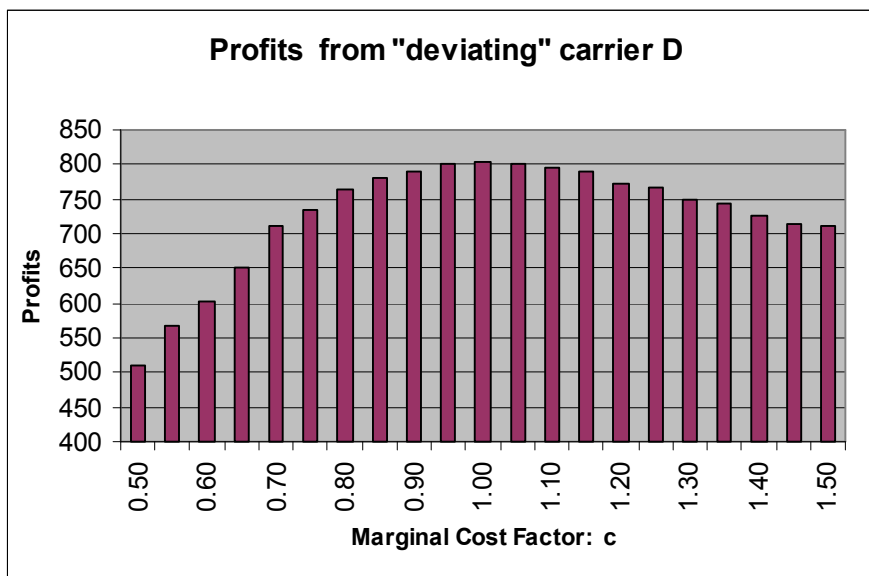
The first concept examined through simulation is optimality of marginal cost bidding. Our setting is different (sequential) and complex enough to introduce some reasonable doubt about the applicability of results obtained for one shot auctions. To test the optimality of marginal cost bidding a numerical experiment was carried out. In the experimental setting described in appendix A, two carriers with different bidding strategies compete. One carrier (called MC) bids only his marginal cost. The other carrier (called D for "deviating") bids his marginal cost multiplied by a constant marginal cost factor c , in other words deviates from the marginal cost policy by a constant factor. Figure 2 shows how the profits of carrier D vary with 21 different values of c , ranging from $c = 0.50$ to $c = 1.50$ in intervals of 0.05 units. These results are the product of averaging 30 runs with 10^4 load arrivals each.

The figure shows that the maximum profit is reached when the marginal cost factor $c = 1$, i.e. under marginal cost bidding. Even though the profit curve is fairly flat around $c = 1$, the result

generalizes to other settings. A similar shaped curved is obtained in settings with more than 2 carriers in the market (e.g. 2 carriers bidding marginal cost and one carrier deviating). The same shape is obtained when adding or subtracting a constant to the marginal cost, in this case profits are maximized when the constant is equal to zero.

The intuition behind the optimality of marginal cost bidding is closely related to the proof that marginal cost bidding is optimal in a single shot auction. With low arrival rates, higher profits are mostly dependent on the number of bids won or market share. Bidding over marginal cost could result in winning fewer bids while keeping payments the same on average. Bidding below marginal cost could result in winning more bids; however the additional won bids occur when the competitor has lower marginal costs on average and therefore bring on a loss to the under-bidder carrier. The danger of underbidding is greatly exacerbated when competitors have better fleet assignment technology and therefore lower marginal costs.

Figure 2 Deviating carrier profits for different deviation factor values



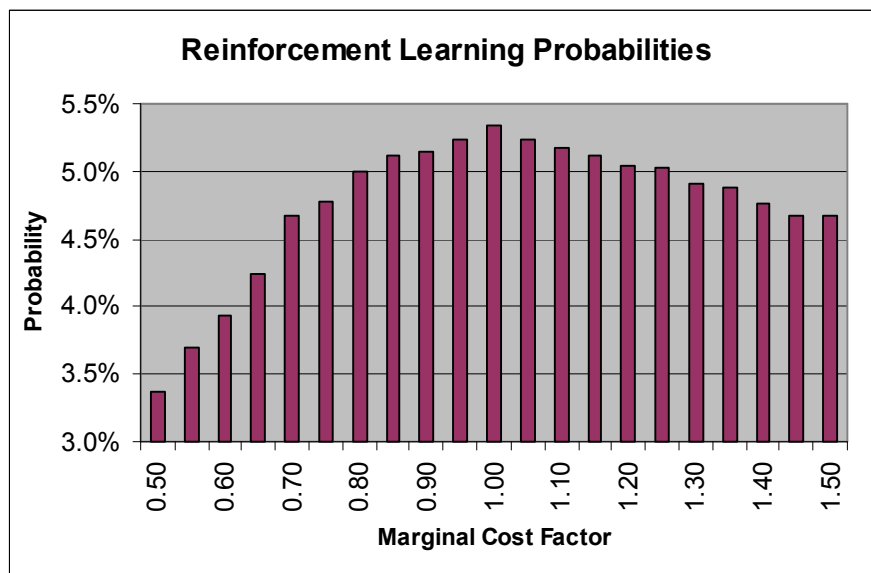
Finally, the results shown in Figure 2 indicate that marginal cost bidding is not only optimal, but also a Nash equilibrium. If two or more carriers are bidding marginal costs, no individual carrier has an incentive to deviate up or down because the deviation will diminish his profits.

What happens when carriers are completely ignorant about the mechanics of the 2nd price auction? Without any previous knowledge, can they learn to bid optimally in a low arrival rate setting? To answer these questions another numerical experiment is carried out. In the same experimental setting two carriers compete. Minimum information availability is assumed.

Carriers implement a reinforcement learning methodology. In order to facilitate comparisons with the previous experiment, carriers can choose 21 different actions ranging from bidding 0.50 of the marginal cost to 1.50 of the marginal cost in intervals of 0.05 units. In reinforcement learning, an agent chooses an action with a probability that is directly proportional to the profit that such action has achieved in the past. Therefore probabilities are also a proxy for profit levels. At the beginning of a simulation a carrier starts with uniform prior probabilities over the possible actions; figure 3 shows the probabilities after 10^4 arrivals.

Reinforcement learning probes each action's profit level, thus the shape similarity between figures 2 and 3. Bidding marginal cost is again favored, though it will not likely be played with probability one since the other actions have positive profit. However, in a stationary environment like this, a carrier could infer from these results that it is optimal to bid the marginal cost. Bounded rationality assumptions about the carrier will ultimately determine whether such an inference is possible.

Figure 3 Reinforcement learning probabilities of bid selection at steady-state



4.4. Strategic Analysis

It has already been mentioned that a carrier could learn the behavioral type of another player in order to manipulate future outcomes. Is this possible in our case? Let's assume that two carriers (denoted A and B, respectively) compete in a setting like the one described in the Appendix, but with *maximum* information availability. Both carriers are playing marginal cost

bidding. If carrier B learned to play marginal cost and stuck to it, it would be unproductive for carrier A to try to manipulate results. This is because playing marginal cost is a Nash equilibrium; any deviation from it would be detrimental for carrier A.

A completely different situation takes place if, after experimentation, carrier A learns that B's behavioral type is "Tit for Tat"; carrier A can then raise its bids to:

$$\text{bid} = \text{marginal cost} * x ,$$

where x = the value that a monopolist would use if he/she were allowed only to set prices equal to marginal cost multiplied by a constant number.

After a finite number of bids T (equal to the number of bids that carrier B uses to determine the level of its Tit for Tat answer), carrier B would be imitating A, that is, multiplying its marginal costs by a constant close to x . Tit for Tat then would be the strategy chosen by an astute player that hopes to be "discovered" by the other competitor, which would drive prices and profits up.

The irony of this result is that a setting with full information (i.e. where all bids can be observed) could easily lead to an uncompetitive result. Technology like the Internet can connect shippers and carriers nationwide in real time, increasing the size and scope of the market, moving them closer to ideal perfect markets. However, the same enabling technologies may also facilitate anticompetitive behaviors like the one just mentioned.

The complexity of playing or identifying Tit for Tat play in a market with 2 carriers is extremely low; at the same time this simple example shows the power of strategic play. Carriers who are aware of how competitors act can manipulate or at least "signal", "communicate", or "anticipate" each others' behaviors and exploit the situation to the detriment of shippers or less skillful carriers.

The example illustrates (a) the benefits of "identifying" behavioral types and (b) the fact that "identification" may be quite simple for some settings and strategies. However, determining if these points are easily generalized to other settings remains unclear. The following observations seem to indicate that "identification" of behavioral types is not easy in general:

- i. The example considers just two carriers. If the number of carriers is greater than two, identification of behavioral types and market manipulation may be impossible. For example, considering three bidders, and define Tit for Tat as follows: imitate the competitor with lowest bid average in the last T periods. In this setting, one carrier cannot identify if the other two carriers are playing Tit for Tat or marginal cost. Even if the other two carriers are playing Tit for Tat, the prices would not be driven up when one carrier experiments. More than unilateral experimentation is required.
- ii. With more than two players, slightly holding back information can make the process of "identifying" competitors' behavioral types much more difficult. For

example, just presenting the bid results (numbers) without carrier names makes it harder to link carriers and bids.

- iii. Assuming that a behavioral type could be identified, achieving such identification in an optimal way (i.e. how should one bid to discover it) can be a complex problem in its own right. Once that behavior type is found, using it to optimize present and future bids may, yet again, be an extremely difficult problem.
- iv. Vickrey auctions reveal too much information about the carriers' costs and operations, and characterization by bidding (one shot) does not depend on competitors' valuations or costs. On the other hand, first price auctions' optimal bidding involves speculating about competitors' costs.

These arguments show that, in general, strategic thinking is not possible for bounded rational carriers. When this is the case, the marketplace has to be modeled as a game against nature or as a simple statistic to represent competitors' possible actions based on past outcomes. A carrier that cannot be strategic has to use one of the non-strategic learning tools available, like fictitious play or reinforcement learning, depending on the extent to which information might be available.

4.5. Analysis and Policy Implications

In general, learning and adapting in a complex environment is unavoidable to remain competitive. In the marketplace setting considered, different amounts of information did not influence the final outcome of learning. This seems to be the result of this particular setting and the properties of second price auctions. Different starting levels of knowledge may lead to different outcomes. In this respect, determining what level of rationality may be too naïve or too rational is essentially context/problem dependent.

Stronger and more robust results occur in a setting where the same outcomes are reached independently of the initial assumptions about rationality and knowledge of the carriers. These robust markets are desirable from a "market engineering" perspective, which seeks good market outcomes from a societal point of view. A deeper understanding of carriers' behavior would be useful in designing markets and answering policy motivated by questions like the following: Can the market perform well across the spectrum of bounded rationality and learning? How does the number of carriers impact the results? How do carrier asymmetries and market power exercise affect market performance? What information should be revealed or concealed in order to increase market efficiency?

5. Conclusions and Future Research

Behavioral studies of the freight transportation sector have been mostly demand driven. Shippers' choices of mode, port, and location have been widely studied. Most freight studies

tacitly assume that supply simply follows demand. Though this is true to some extent, the behavioral aspects of transportation service suppliers have been somewhat neglected. Carriers' decisions may have a long term effect in the transportation system involving large sunk costs (facilities, vehicles, etc). Oligopoly environments (a small number of suppliers) can allow exercise of market power or collusion.

Both shippers and carriers can be considered utility maximizers; however, carriers' decisions or choices tend to be loaded by an important strategic component. This research provides a framework for studying carrier behavior where the complexity and characteristics of the decision making environment, plus the rationality endowment of the decision maker (carrier), shapes behavior and learning.

The Vickrey auction example illustrates the importance of rules and information availability. The study of idealized marketplaces may provide important insights. Similarly, empirical analysis of transaction data (when available) will provide a critical test for theoretical and simulation results, further energizing this area of travel behavior research.

Acknowledgments

This study is supported by a grant from the National Science Foundation, Program: NSF 00-42, Award Number: CMS-0085691. The content of the paper remains solely the responsibility of the authors.

6. References

- Axelrod, R. (1984) *The evolution of cooperation*, Basic Books, New York.
- Binmore K. (1987), Modeling Rational Players, I, *Economics and Philosophy*, Vol. 3, pp 179-214.
- Brown G. (1951), Iterative solution of games by fictitious play. In *Activity analysis of production and allocation* (Koopman T. ed), Wiley, New York.
- Cho I. K., Kreps D., (1987), Signaling games and stable equilibria, *Quarterly Journal of Economics*, Vol. 102, pp 179-221.
- Figliozzi M, Mahmassani H., and Jaillet P. (2003), A Framework for the study of Carrier Strategies in an Auction Based Transportation Marketplace, forthcoming in *Transportation Research Record*.

- Goldberg D. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison Wesley, Massachusetts.
- Kalai E. and Lehrer E. (1993), Rational learning leads to Nash equilibrium, *Econometrica* Vol. 61, pp 1019-1045.
- Krishna, V. (2002), *Auction Theory*, Academic Press, New York.
- Lucking-Reiley, D. and Spulber, D. (2001), Business-to-Business Electronic Commerce *Journal of Economic Perspectives*, vol. 15, no. 1, pp. 55-68.
- McAfee, R. Preston, and John McMillan (1987), Auctions and Biddings, *Journal of Economic Literature*, Vol. 25, pp 699-738.
- Mahmassani, H.S. (2001), Freight and Commercial Vehicle Applications Workshop Report, in Hensher, D.A. *Travel Behaviour Research: The Leading Edge*, Pergamon.
- Milgrom P., Roberts J. (1990), Rationalizability, learning and equilibrium in games with strategic complementarities, *Econometrica* 58, pp 1255-1277.
- Monderer D, Shapley L. (1996), Fictitious Play property, *Journal Economic Theory*, Vol. 68, pp 258-265.
- Roth, A. and Erev, I. (1995), Learning in extensive form games: experimental data and simple dynamic models in the intermediate term, *Games and Economic Behavior*, Vol. 8, pp 164-212.
- Simon, H. (1972), Theories of Bounded Rationality, in *Decision and Organization* (C. McGuire and R Radner, Eds), Amsterdam: North Holland.
- Stahl, D. (1994), Evolution of smart n players, *Games and Economic Behavior* Vol. 5, pp 604-617.
- Vickrey, W., (1961) Counterspeculation, auctions and competitive sealed tenders, *Journal of Finance*, Vol. 16, pp 8-37.

Appendix A: Simulation Settings

Carriers are identically implemented

- Geographic Area : 1 by 1 unit square area
- Shipment Origins and Destinations \approx Uniformly distributed
- Shipment Type: Truckload movements only
- Shipments Reservation Price \approx Uniformly distributed (1.45, 1.5)
- Shipment Time Window Length: 2 units of time
- Earliest Delivery Time = arrival time
- Latest Delivery Time = arrival time + Time window length

- Truck Speed : 1 (unit of distance/unit of time)
- Total Fleet size: 12 vehicles (constant) serving the market
- Number of Carriers: 2, 3, 4, and 6 with fleet sizes of 6, 4, 3, and 2 trucks respectively
- Poisson Arrival Rate: 2 shipments per unit of time (very uncongested conditions)