Repeated Auction Games and Learning Dynamics in Electronic Logistics Marketplaces:

Complexity, Bounded Rationality, and Regulation through Information

M.A Figliozzi[a,1] **, H.S. Mahmassani**[b,2] **, and P. Jaillet**[c,3]

[a] Portland State University
College of Engineering and Computer Sciences
PO Box 751, Portland, Oregon 97207-0751

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

[c]Massachusetts Institute of Technology
Department of Civil & Environmental Engineering
Cambridge, MA 02139-4307, USA

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. Bounded rationality, learning, and behavior in a sequential auction marketplace are analyzed and simulated.

Key Words: Freight Transportation; Carrier Behavior; Game Theory; Bounded Rationality; Learning in Auctions; Carrier Management Strategies; Electronic Commerce

[1] E-mail: miguel@itls.usyd.edu.au, URL: www.itls.usyd.edu.au
[2] E-mail: masmah@umd.edu, URL: www.cee.umd.edu
[3] E-mail: jaillet@MIT.EDU, URL: www.cee.mit.edu

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).

McAffee 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 chapter is the study of transportation marketplaces (TM) 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, this chapter 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. The behavioral aspects of auction market behavior are more readily articulated without the added complexity of the combinatorial aspect. However, other market settings are possible. Markets where carriers bid on configurable bundles of loads give rise to combinatorial auctions. Nandiraju and Regan (2005) present a comprehensive survey of freight transportation electronic marketplaces.

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 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 chapter has nine sections. Next section introduces mathematical notation and describes the marketplace framework and operation. Section 3 articulates a framework to study carrier behavior. Section 4 identifies the characteristics of transportation auctions as well as associated sources of complexity and bounded rational behavior. Two sources of bounded rational behavior, knowledge acquisition and problem solving capabilities, are analyzed in Sections 5 and 6. Section 7 discusses learning in a TM setting. Reinforcement learning and fictitious play are analyzed and adapted to the particularities of a transportation marketplace. Section 8 presents different computational experiments aimed at studying the properties of different auction settings and learning methodologies. Section 9 ends with a chapter summary and conclusions.

2 Description of Transportation Marketplaces

The TM enables the sale of cargo capacity based mainly on price, while still satisfying customer level of service demands. The specific focus of the study is the reverse auction format, where shippers post loads and carriers compete over them (bidding). The auctions operate in real time and transaction volumes and prices reflect the relative status of demand and supply. A framework to study transportation marketplaces is presented by Figliozzi et al. (2003a). The market is comprised of shippers that independently call for shipment procurement auctions, and carriers, that participate in them (we assume that the probability of two auctions being called at the same instant is zero). Auctions are performed one at a time as shipments arrive to the auction market. Shippers generate a stream of shipments, with corresponding attributes, according to predetermined probability distribution functions. Shipment attributes include origin and destination, time windows, and reservation price. Reservation price is the maximum amount that the shipper is willing to pay for the transportation service. It is assumed that an auction announcement, bidding, and resolution take place in real time, thereby precluding the option of bidding on two auctions simultaneously.

Consider a TM in which n carriers are competing; a carrier is denoted by $i \in \Im$ where $\Im = \{1, 2, ..., n\}$ is the set of all carriers. Let the shipment/auction arrival/announcement epochs be $\{t_1, t_2, ..., t_N\}$ such that $t_i < t_{i+1}$. Let $\{s_1, s_2, ..., s_N\}$ be the set of arriving shipments. Let t_j represent the time when shipment s_j arrives and is auctioned. There is a one to one correspondence between each t_j and s_j (i.e. for each t_j there is just one s_j). Arrival times and shipments are not known in advance. The arrival instants $\{t_1, t_2, ..., t_N\}$ follow some general arrival process. Furthermore, arrival times and shipments are assumed to come from a probability space $(\Omega, \mathcal{F}, \mathcal{P})$, with outcomes $\{\omega_1, \omega_2, ..., \omega_N\}$. Any arriving shipment s_j represents a realization at time t_j from the aforementioned probability space, therefore $\omega_j = \{t_j, s_j\}$.

In an auction for shipment s_j, each carrier $i \in \Im$ simultaneously bids a monetary amount $b_j^i \in R$ (every carrier must participate in each auction, i.e. submit a bid). A set of bids $b_j^{\Im} = \{b_j^1, ..., b_j^n\}$ generates publicly observed information y_j. Under maximum information disclosure, all bids are revealed after the auction, i.e. $y_j = b_j^{\Im}$. Under minimum information disclosure, no bids are revealed after the auction, i.e. $y_j = \{\}$. Each carrier is informed only about his bidding outcome: successful or unsuccessful. The fleet status of carrier i when shipment s_j arrives is denoted as z_j^i, which comprises two different sets: S_j^i (set of shipments acquired up to time t_j by carrier $i \in \Im$) and V_j^i (set of vehicles in the fleet of carrier i, vehicle status updated to time t_j). The estimated cost of serving shipment s_j by carrier $i \in \Im$ of type z_j^i is denoted $\mathrm{c}^i(s_j, z_j^i)$. Let I_j^i be the indicator variable for carrier i for shipment s_j, such that $I_j^i = 1$ if carrier i secured the auction for shipment s_j and $I_j^i = 0$ otherwise. The set of indicator variables is denoted $I_j^{\Im} = \{I_j^1, ..., I_j^n\}$ and $\sum_{i \in \Im} I_j^i \leq 1$. Let π_j^i be the profit obtained by carrier i for shipment s_j, then $\pi_j^i = (b_j^i - c^i[s_j, \theta_j^i])I_j^i$.

Bidders have private costs when each bidder knows the cost of the object at the time of bidding. This cost is the disutility that the bidder himself obtains from the consumption, use, possession or service of the auctioned item. Let $\Im = \{1, 2, ..., n\}$ be the set of bidders and θ^i denote the private information that buyer (seller) i possesses about the value (cost) of the item being auctioned. Private values are assumed in this chapter, therefore, $\theta_j^i = \{z_j^i, \mathrm{a}^i, \mathrm{c}^i\}$ is the private information of any carrier $i \in \Im$ at time t_j. Carrier $i \in \Im$ is uncertain about $\theta_j^{-i} = \{z_j^{-i}, \mathrm{a}^{-i}, \mathrm{c}^{-i}\}$ at time t_j, the proprietary private information regarding competitors' fleet status, assignment, and cost functions respectively. The superscript *–i* is used to indicate the set of competitors of carrier *i.*

3 Determinants of Carrier Behavior

In a TM, carrier behavior is defined as a sequence of bids taken by a carrier. This section looks into the elements or factors that determine carrier behavior. These factors are: carrier technology, bounded rationality, information availability, and strategic setting. Though all the factors are somewhat related, the first two are prominently intrinsic to the carriers' own characteristics, while the last two are predominantly linked to environmental or somewhat extrinsic factors. In this section, the discussion is limited to highlight the link between them and carrier behavior.

3.1 Carrier Technology

Carrier technology or the sophistication of the dynamic vehicle routing problem (DVRP) solution has an important role in bidding. In the bidding decision making process the carrier technology determines the number of feasible schedules to be evaluated. Therefore, unsophisticated DVRP technologies seriously limit the quality and quantity of alternatives that could be evaluated (Figliozzi et al., 2004).

3.2 Auction Rules - Information Revelation

Different auction payment rules lead to different bidding functions. Information revelation rules can also play a significant role (Figliozzi et al., 2003b). The information that is revealed (before bidding begins or after each auction) can influence how, how much, and how fast carriers can learn or acquire knowledge about the strategic setting and competitors' behaviors. The information that could be available after auctions are resolved includes: bids placed, number of carriers participating, links (names) between carriers and bids, and payoffs. The information that could be available before bidding begins includes: some carriers' individual characteristics (e.g. fleet size or previous performance/profits from public financial reports), information about who knows what, information asymmetries, or common knowledge about previous items. Private information (as defined in Section 2) is not included since it involves proprietary information that usually is to the best interest of the carrier to keep private.

Two extreme information scenarios can be defined: maximum and minimum. A maximum information environment is defined as an environment where all the information, mentioned in the previous paragraph, is revealed. On the other hand, an environment where no information is revealed is called a minimum information environment. These two extreme scenarios can approximate two realistic situations: maximum information would correspond to a real time internet auction where all auction information is equally accessed by participants; minimum information would correspond to a shipper telephoning carriers for a quote. The shipper calls back just the selected carrier (if any is selected).

3.3 Strategic Setting

In this chapter, it assumed that a carrier operates in an environment determined by the other carriers' behaviors; a carrier uses a model of the behavior of the other carriers as an input to his decision problem. Under this interpretation a carrier's bidding function suits a carrier's best interest, assuming that competitors bidding functions pursues competitors' best interests. This is defined as a competitive strategic environment.

A diametrically different environment is a collusive or collaborative environment. One danger of auctions is the possibility that buyers/sellers who repeatedly participate in the same auctions could engage in collusive behavior. This topic is of primordial importance in the field of Industrial Organization – general references to this area include the work of Tirole (1989) and Martin (1993). As a general rule, the more information is revealed, the easier collusion becomes. Even in minimum information settings collusion is possible. Blume and Heidhues (2006) study collusion in repeated first-price auctions under the condition of minimal information release by the auctioneer. In each auction a bidder only learns whether or not he won the object. Bidders do not observe other bidder's bid, who participates or who wins in cases in which they are not the winner. Even under these restrictive assumptions, for large enough discount factors, collusion can nevertheless be supported in the infinitely repeated game. Nevertheless, it may entail complicated inferences and full monitoring among them. Marshal and Marx (2002) analyze bidder collusion in first and second price auctions and Symmetric Independent Private Value assumptions (SIPV) assumptions. The SIPV assumptions are strong but simplify the bidding problem significantly. In general, SIVP models can be studied analytically (Krishna, 2002). As detailed in Section 4, the TM characteristics render the bidding problem intractable.

The two environments, competitive and collusive, are nonetheless connected since underlying every negotiation or agreement there is a game-like component (Raiffa et al., 2002). From each carrier's individual perspective, the incentives (and legal or market risks) of collaborating with competitors has to prevail over the profits that can be obtained when each party acts separately (competitive environment). The auction rules, e.g. first price, second price, open, closed, etc., do affect carriers strategies. For a general introduction to auction types and bidding strategies, the reader is referred to the comprehensive book by Krishna (2002).

3.4 Bounded Rationality

Bounded rationality limitations affect a) the knowledge that a carrier is able to acquire, and b) the bidding problem that the carrier can solve. Given the carrier's rational limitations, fleet technology, information available, and a competitive strategic setting the carrier ends up solving a bidding problem that it is constrained by his/her rational or computational constraints. Bounded rationality in a TM is studied in Section 4.

3.5 Framework for Carrier Behavior

Figure 1 presents a schematic overview of the process that brings about carriers' behavior in a TM. A shipper's decision to post a shipment in the auction market initiates an auction. Carriers respond to auctions postings. Carriers attempt to maximize profits by adjusting their behaviors in response to interactions with other carriers and their environment. Bounded rationality limitations are pervasive and affect how a carrier models, evaluates, and optimizes his action as indicated by the arrows in Figure 1. Carriers also must abide by the constraints and the physical feasibility specified by their assignment strategies and pool of awarded shipments.

In this framework, carriers' learning and knowledge about other competitors' behavior types evolve jointly over time and their strategies at a given moment are contingent on interactions that have occurred or will occur in a path-dependent time line. 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. 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 may also be the only means of communication for a carrier to "identify" or "manipulate" other competitors.

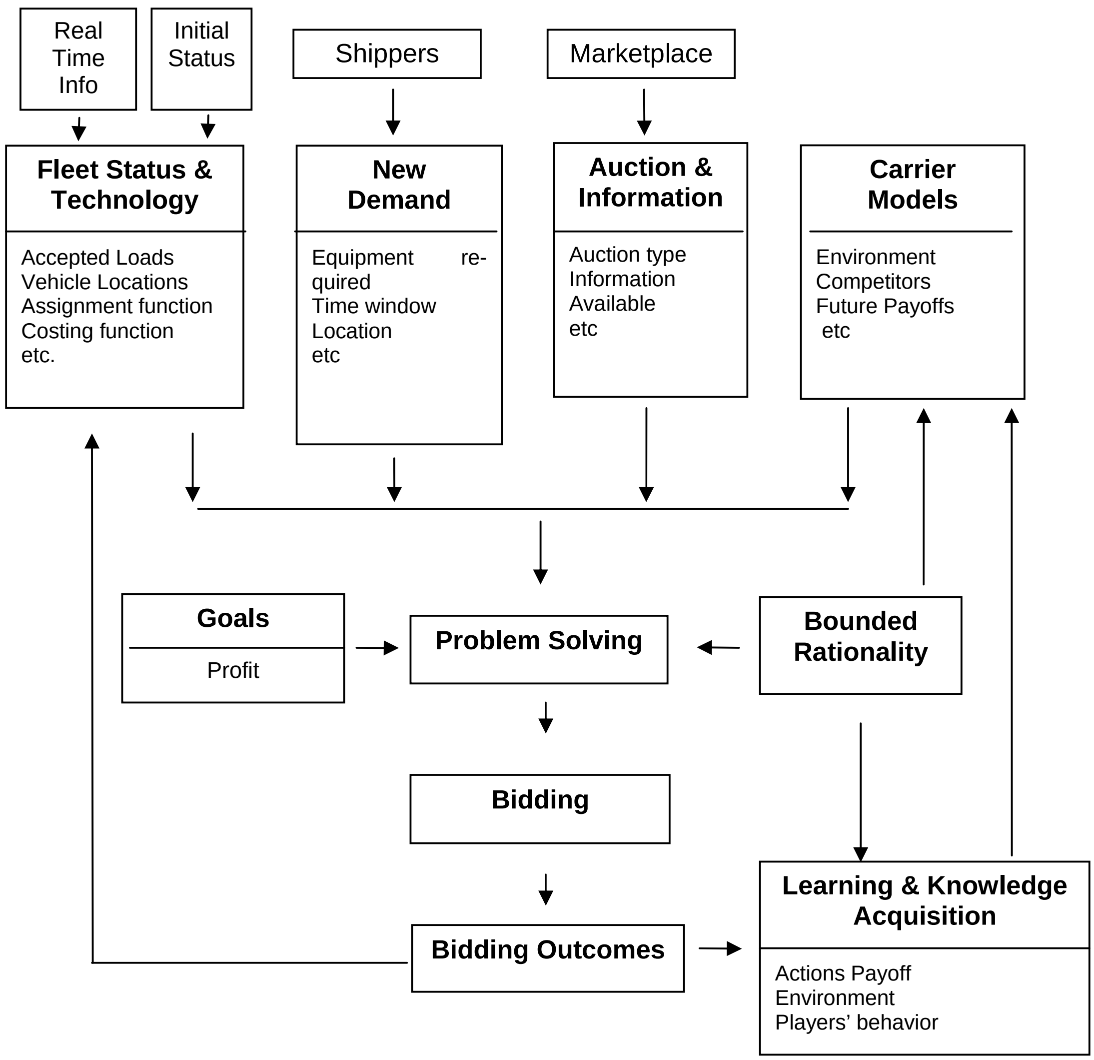

Figure 1: Carrier behavior in a sequential auction transportation marketplace (TM)

4 Bounded Rational Behavior in the Transportation Marketplace

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.

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.

Competition in a TM is an ongoing and sequential process, and thus naturally represented as an extensive-form game. The standard definition of rationality (for economists at least) requires that agents automatically solve problems that are in fact beyond the capabilities of any agent (Colinsk, 1996). The problem is intractable and well beyond the conceptual and computational abilities of ordinary humans or decision support systems. In addition, response time limitations or the framing effects and cognitive limitations of the human mind prevent bidders from behaving rationally. The framing and cognitive limitations of human thinking have been widely studied and reported (Camerer, 1995, Kagel, 1995), mainly in the psychology and economics literature. Therefore, the basic motivation for studying models of bounded rationality in TM environments comes about from the implausibility of perfect rationality models.

When the complexity of the auction problem precludes bidders from implementing a full game theoretic approach, computational agents (or human beings with the help of decision support systems) need to simplify or alter the original choice or decision problem. Bounded rational behavior, as studied in this research, is born out of these simplifications or alterations to the original insurmountable problem. This chapter attempts to provide a behavioral framework to understand how carriers can tackle the overwhelming complexity of the problems they face in a TM (complex detailed stories, numerous current options, future infinite contingent options, and the potential consequences).

Bounded rational bidders solve a less complex problem than fully rational bidders. The type of problem they solve is directly influenced by available response time, existing computational/material resources, and their own cognitive/decision-making process. Although the result of bounded rational deliberation would not necessarily be an equilibrium solution, the bounded rational response would have more bearing on how ordinary carriers or human decision makers would act in sequential auction TM. The introduction of bounded rational decision makers radically alters the notion of equilibrium and decision making. Game theory assumes that players know the prevailing equilibrium and act consequently. For bounded rational agents, the equilibrium, if any, is not known beforehand, it is built.

Bounded rational behavior is born out of simplifying a (complex) problem or the cognitive/material limitations of the decision maker (or decision support system). Therefore, bounded rationality is always associated with the notion of deficiency or insufficiency of a positive quality (of a rational player). Although bounded rationality as a research topic is not new, it was first proposed by Simon (1956), many modeling issues surrounding bounded rational decision making have not yet been fully addressed. Bounded rationality and learning in games are currently very active areas of research; however general and comprehensive models that integrate how agents (or humans) acquire, process, evaluate, search for information, and make decisions are still mostly open. As expressed by Aumman, "there is no unified theory of bounded rationality, and probably never will be." (Aumann, 1997 - page 4).

Rationality assumptions are very convenient from a modeling point of view. The self-referential nature of rationality (coupled with common knowledge in games) imposes astringent limitations on how a rational agent (player in a game) foresees his competitors' behavior and how the competitors foresee other players' behavior. Bounded rationality come with an embarrassment of riches in terms of the number of possible deviations from a fully "rational" model. When bounded rational behavior appears, it may take on many different forms. Bounded rational decision makers do not necessarily choose equally, even when having the same knowledge or information. Furthermore, there may be many "plausible" bounded rational models that can explain a given social or economic phenomenon. Correspondingly, the many possible ways each bounded rational bidder can model his competition adds a class of uncertainty not found where all players are perfectly rational.

Determining the bounded rationality of a carrier is crucial since it is equivalent to determining how the carrier bids (i.e. his bidding function) in a TM. Similarly, determining that all carriers are rational is equivalent to determining how the carriers bid (i.e. their bidding function) in a SIPV setting. A bidding function, as understood in this research, is a process, whose inputs are a carrier's private information and his knowledge about the auction and competitors, and whose output is a bid. Given the plethora of games and decision problems, bounded rational behavior is hard to define, classify, and model in general terms. When the restrictions of rationality are lifted, any general assumption about the behavior of the bidders that is not properly justified, introduces a strong sense of arbitrariness. In order to avoid this kind of arbitrariness, the discussion of bounded rationality is limited to the TM context. Any departure from the rationality model is connected to carriers' cognitive and problem solving processes.

Bounded rationality can stem from different cognitive and computational/physical limitations, in the TM context, the following classification of sources is proposed:

- Bounded Recall and Memory: a carrier has limited memory (physical capacity) to:
    - record and keep past data/information

- simulate and record data of all future possible paths in the decision tree

- Processing Speed: time is valuable in a dynamic setting. Most practical problems have a limited response time that may limit the solution quality or decrease the effectiveness of a delayed response.
- Data Acquisition and Transmission: data acquisition and processing is usually costly. Furthermore, the transmission of data among agents can be noisy. In a world with bounded resources (budget/memory/attention), deciding how, how much, and what type of information should be acquired, kept, transmitted, or analyzed can lead to complex decision problems.
- Knowledge Acquisition: in a dynamic strategic situation, as data is being revealed or obtained, carriers have the potential to acquire knowledge (truths about competitors or the environment) from logical and sound inferences. In particular, the decision maker may have limited ability to discover competitors' behavior, which may involve modeling and solving complex logical and econometrics problems.
- Problem Solving: as a carrier participates in a TM market, it is required to make decisions (bidding or fleet management decisions). These decisions may lead the carrier to formulate and solve complex optimization problems. In particular, the decision maker may have limited ability to predict or model the impact of his own actions on future fleet operational costs or on his competitors' behavior.

Although the five aspects of bounded rationality are somewhat interrelated, this research focuses on the knowledge acquisition and problem solving aspects. Memory and processing speed are physical limitations. It is assumed carriers have enough material resources and response time/speed to implement bidding and fleet management strategies with different degrees of sophistication. Carriers have limitations to formulate and elucidate knowledge acquisition problems. Similarly, carriers have limitations to formulate and solve complex optimization problems. The data available to carriers is only limited to data publicly and freely disclosed after each auction, which renders the data acquisition problem trivial. No transmission losses or alterations are considered.

The focus of this research is on the knowledge acquisition and problem solving aspects, as they capture how carriers can frame and solve TM problems. Therefore, the emphasis is on the more "mental" processes that determine behavior rather than on the "physical" limitations. Knowledge acquisition and problem solving in a TM are analyzed in the next two sections respectively.

5 Knowledge Acquisition in a Transportation Marketplace

In a TM, each carrier is aware that his actions have significant impact upon his rival's profits, and vice-versa. In the perfect rational model, common knowledge and logical inferences allow the estimation of the impact of a carrier's actions on competitors' profits and vice-versa. It is implicit that a rational bidder bids as a rational bidder. In a bounded rational model, a

carrier faces two basic types of uncertainties regarding the competition: (a) an uncertainty relative to the private information of his opponents, and (b) a strategic uncertainty relative to bounded rationality type of the others players.
The first type of uncertainty is about $\theta_j^{-i} = \{z_j^{-i}, a^{-i}, c^{-i}\}$ for a carrier $i \in \Im$ at time t_j, the private information regarding competitors' fleet status, assignment, and cost functions respectively. This type of uncertainty is also present in most game theoretic auction models (games of incomplete information). The second type of uncertainty is about the bidding strategies that the competitors use, $b^{-i} = \{b^1, ..., b^{i-1}, b^{i+1}, ..., b^n\}$ the set of bidding functions of all carriers but carrier i. It is implicit that a bounded rational bidder bids accordingly, i.e. as a bounded rational bidder. However, it is not evident for the competition to determine what "type" of bounded rationality a carrier has. This type of uncertainty is not present in game theoretic auction models.
Depending on a carrier's ability to elucidate uncertainties (a) and (b), two extreme cases may take place:

1. No knowledge acquisition. The carrier cannot form a useful model of competitors' behavior that links their private information and their bids. In this situation, the "best" a carrier can do is to observe market prices and estimate them as the result of a random process. This is similar to assuming that competitors are playing $b^{-i}(\xi) = f(\xi)$ or simply $b^{-i}(\xi) = \xi$, where ξ is a random process that is not linked in any way to carrier $i's$ bidding, capacity/deployment, and history of play or to the competitors' private information $\theta_j^{-i} = \{z_j^{-i}, a^{-i}, c^{-i}\}$.
2. Full knowledge acquisition. The carrier knows $\theta_j^{-i} = \{z_j^{-i}, a^{-i}, c^{-i}\}$ and also $b^{-i} = \{b^1, ..., b^{i-1}, b^{i+1}, ..., b^n\}$, therefore carrier i is able to precisely foresee what the competition is going to bid for shipment s_j. However, carrier i still has uncertainties about the future bids, simply because carrier i does not know the future realizations of the demand. Nevertheless, carrier i can estimate future prices not just as a stationary random process but as a function of shipment arrival distribution, shipment characteristics distribution, competitors' behavior, and competitors' private information. This is $\xi = f(\Omega, \theta_j^{-i}, b^{-i})$.

In game theoretic terms the former case is not possible since there is no "strategic" game if players cannot speculate about the competitors' actions. The latter case corresponds to a game of perfect and complete information *if all the players are rational* and the private information is common knowledge. Knowledge states in between the two extreme cases correspond to games of imperfect information, *if all the players are rational* and there is uncertainty about the players' private information.

The uncertainty about the players' private information can be expressed as $\mathrm{p}(\theta_j^{-i} \mid \theta_j^i, h_{j-1})$. In a game of incomplete information each player (bidder) has expectations (beliefs) about the competitors' private values. Following Harsanyi's (1967) modeling of games of incomplete information, players' types $\theta_j^{\Im} = \{\theta_j^i\}_{i=1}^n$ are drawn from some probability density function $\mathrm{p}(\theta_j^1, ..., \theta_j^n)$ where types θ_j^i belong to a space Θ^i. The conditional probability about his opponents' types $\theta_j^{-i} = \{\theta_j^1, ..., \theta_j^{i-1}, \theta_j^{i+1}, ..., \theta_j^n)$ given his own type θ_j^i is denoted $\mathrm{p}(\theta_j^{-i} \mid \theta_j^i, h_{j-1})$. This is what characterizes and complicates the solution of a dynamic game of incomplete information. Since the players do not know the competitors' types at the start of each auction, they have to update these conditional probabilities (beliefs about the competitors' status) as public information is revealed and the game evolves.

Acquiring knowledge about the competitors' private information and bounded rationality type poses a potentially highly complicated econometric/logical problem. A carrier's behavior is likely to be affected by his own history and how the carrier perceives and models the strategic situation. From the public information (revealed after each auction) and its own private information a carrier needs to build a model of the private information and bounded rationality type of his competitors.

Even in simple auctions, the econometric models can quickly become extremely complex and data are usually not rich enough to successfully estimate those structurally complex models (Laffont, 1997). Furthermore, the complexity of the underlying DVRP adds hurdles to the problem. However, the most challenging obstacle may come from the competitors, which may be "sophisticated" enough to realize that they are bidding against other bidders who are also learning and may adjust their behavior accordingly, in order to obstruct the process of knowledge acquisition. This type of sophistication is particularly important when the fact that the same carriers interact repeatedly is common knowledge.

In most game theoretic models, a simple private value probability distribution, symmetry, rationality, and common knowledge assumptions permit a closed analytical solution. In equilibrium bidders know the competitors' bidding function, however, they do not know the realization of the competitors' private value, therefore they do not know the competitors' actual bid. Conversely, in a TM, private values are not random but correlated, the status of a carrier at time t_j provide useful information to estimate the status of the carrier at time t_{j+1}. A bidder may potentially obtain information about competitors' private values and bidding functions if the bidder invests resources to infer them. Market settings, such as auction data disclosed and number of competitors, strongly affect the difficulty of the inference process.

Summarizing, repeated interaction can lead to learning and knowledge acquisition. This research distinguishes among the two. Learning takes place in the no-knowledge case; the carrier does not get to know the competitors' behavioral processes just the price function as a random process. Learning is superficial, it is merely phenomenological. In the full knowledge case, the carrier acquires knowledge about the competitors' behavioral processes. Knowledge acquisition is deeper; it is causal. Learning in a TM environment is discussed in Section 7.

6 Problem Solving in a Transportation Marketplace

The previous section focused on "what can be learnt or known" about the competition. This section specifically contemplates "how carriers come up" with a bid or decision given what has been learnt or what knowledge has been acquired about a problem. Usually, models in which decision makers are assumed rational do not explain the procedures by which decisions are taken, rational procedures are implicitly embedded in the answer or approach. Furthermore, economic models pay no or little attention to how hard it is to make decisions. Conversely, bounded rational decision maker models detail the procedural aspects of decision making. Those detail procedures are the essence of a bounded rational decision making model. The degree of intricacy of the decision making procedure is used in the last part of this chapter to classify bounded rational behaviors. As a carrier participates in a TM market, it is required to make decisions, to choose among alternative future paths. Each decision poses a problem that the carrier has to solve (not necessarily optimally). The rest of this section analyzes, in this order, the type of decision a carrier faces in a TM and how bounded rationality can appear in the steps of a decision making process.

From the carriers' point of view, the choice problems that take place in a TM are either bidding or operational (fleet management) decisions. Bidding decisions may carry a strategic value since they directly affect competitors' profits. Bidding decisions are also the result of a bounded rational decision process, a carrier's choice and therefore can reveal or transmit information about a carrier's decision making process or intentions. Operational (fleet management) decisions mostly affect a carriers' own fleet status (private information). Therefore, operational decisions are considered non-strategic and take place as new information arrives: auctions are won or shipments are served. This type of decision, for example, includes the estimation of a shipment value or service cost, the rerouting of the fleet after a successful bid, the reaction to unexpected increase in travel times, etc. A detailed formulation of the value or service cost problem is found in Figliozzi et al. (2006, 2007).

There are several factors that contribute to the complexity of biding in a TM. These factors are: competitors bounded rationality, knowledge about the competitors, look-ahead depth, and the type of auction utilized. This section analyzes the first three factors. The auction characteristics that significantly affect bidding complexity, from the bidder's perspective, are: (a) the use of incentive compatible mechanisms and (b) the number of item being auctioned, e.g. combinatorial auctions are more demanding computationally than single item auctions. Incentive compatible mechanism simplify considerable the bidder's problem because the optimal bid is the cost or reservation value, regardless of the actions of the competitors (Figliozzi, 2006).

Sophisticated bounded rational players have a “model” of the other players. For example, in the work of Stahl and Wilson (1995) and Vidal and Durfee (1995), players model other players’ cognitive process and decision rules up to a finite number of steps of iterated thinking. The number of iterations that a player can perform is a measure of the sophistication of a player. A zero level player does not model his opponents, it simply ignores the fact that other agents exit. Reinforcement learning is an example of this type of agent sophistication. A one level agent models only the frequency or another statistic that represents other players’ actions. Fictitious play is an example of this type of agent sophistication. A two level agent can simulate the other agents’ internal reasoning process (i.e. a model of level zero or level one agents) and take an action by taking into account how the other players (of level zero or one) are going to play. A level three agent can build models, simulate them, and act in response to the behavior of players up to level two. Recursively, a level four agent can model the actions of level three agents and so on. Perfectly rational agents can follow the recursion to an infinite level. Then, if the level of rationality of a player is denoted by L^i, then that player can model the most sophisticated of his competitors up to a level $L^{-i} = L^i - 1$.

Section 5 dealt with the level of knowledge about the competition. A player with no-knowledge about the competition can only implement a level zero or level one type of player since it cannot link his actions (bids) to the consequences that his actions have. A player with full knowledge could possibly foresee (if it could only solve the corresponding problems) the behavior of any player type. However, the complexity increases as the level type to be implemented increases, i.e. as the competitors bounded rationality sophistication increases. The carrier with full-knowledge knows $\theta_j^{-i} = \{z_j^{-i}, a^{-i}, c^{-i}\}$ about the competition and also $b^{-i} = \{b^1, ..., b^{i-1}, b^{i+1}, ..., b^n\}$. Therefore, carrier i can compute precisely what the competition is going to bid for shipment s_j. However, carrier i still has uncertainties about the future bids, simply because carrier i does not know the future realizations of the demand. Nevertheless, carrier i can estimate future prices, not just as a random process but as a function of shipment arrival and characteristics distribution, competitors’ behavior and competitors’ private information. When the knowledge is imperfect, complexity further increases since there is a probability distribution over the competitors’ private information space. Furthermore, the probability distribution is a function of the history of play and the competitors’ fleet management strategies. In mathematical notation, the probability distribution of competitors’ future private information is $p(\theta_N^{-i} \mid h_N)$.

The third factor is the look-ahead depth. In a sequential auction setting like a TM, bids affect future auctions profits. The look-ahead depth is the number of future auctions that are taken into account when estimating how a bid may affect future auctions profits. A zero step look-ahead (or myopic) analysis does not consider future auction profits, just the profit for the current auction. A one-step look-ahead analysis considers one future auction, current plus the following auction profits. Similarly, a m-step look-ahead analysis considers m future auctions, current plus the following m auction profits. When the analysis is myopic, shipment s_j

is known and the uncertainties are reduced to a minimum. Projecting one step into the future, the arrival time (t_{j+1}) and characteristics of shipment s_{j+1} are uncertain. Furthermore, if the link between bidding and future prices ξ_{j+1} is incorporated, the optimal bid for shipment s_j takes into account its impact on competitors' bids (prices) in the next auction. Then, for shipment s_{j+1} the price function at time t_{j+1} is a function of the previous bids and the unknown previous arrival $\xi_{j+1}(s_j, b_j^{*i})$. In the one-step problem, the arrival and characteristics of s_{j+1} are uncertain, but the future history h_{j+1} is a function of the already known s_j. Projecting two steps into the future, the estimation of the future price function ξ_{j+2} becomes more complex. The price function ξ_{j+2} for shipment s_{j+2} is a function of the yet unknown s_{j+1} and the two previous bids $\{b_j^{*i}, b_{j+1}^{*i} \mid h_{j+1})$. Moving one extra step into the future increases the problem complexity significantly. For shipment s_{j+2} the price function at time t_{j+2} is a function of the previous bids and the unknown previous arrival $\xi_{j+2}(s_j, s_{j+1}(t_j, \Omega), b_j^{*i}, b_{j+1}^{*i} \mid h_{j+1})$. Calculation of future price functions is increasingly difficult as uncertainties and dependencies on earlier (but not yet realized) bids and shipments accumulate. When the look ahead is up to shipment s_N, the number of decision variables $B^{*i} = \{b_j^{*i}, ..., b_N^{*i} \mid h_N\}$ to be estimated is:

$$\sum_{k=0}^{N-j} p^k \ .$$

When the number of players (bidders) is n, after each auction there are n possible outcomes and future histories. If backward induction is used, for each possible history it is necessary to estimate an optimal bid, the total number of decision variables increases exponentially with the number of future look-ahead steps. Let denote by $\Sigma = \{s_j, s_{j+1}(t_j, \Omega) ..., s_N(t_{N-1}, \Omega)\}$ the set of shipments to be analyzed. Then, the future price function when earlier bids affect future prices and the carrier has imperfect information is a function of $\xi_N = \mathrm{f}(b_j^{*i}, ..., b_{N-1}^{*i}, \Sigma, \mathrm{p}(\theta_N^{-i} \mid h_N))$.

Table 1 puts the three factors together. The table is set up in such a way that the complexity of the price function ξ increases, moving downward or rightward. With higher levels of competitors' bounded rationality, the complexity of the problem increases exponentially with the number of iterations and players to be simulated.

The symbol $\langle \cdot \rangle^{nL^{-i}}$ is used to denote the number of iterations as a function of the number of players and the highest level of iterations that the competition can sustain. A "fully rational" equilibrium, is a special case of the imperfect knowledge case when all players are rational and $L^{-i} \to \infty$. In the game theoretic case, it is common knowledge that all the bidders are simultaneously foreseeing and simulating each other's bids and decisions at infinitum. Each cell of table 2 is a different decision theory problem that can potentially be expressed as a mathematical program or algorithm. It was mentioned that the complexity increases moving downward or rightward.

Table 1: Bidding Complexity as a function of price function (ξ) complexity

Knowledge Level	**Own Type** L^i	**Comp. Type** L^{-i}	**Look-Ahead Depth** **Myopic** $\Sigma = \{s_j\}$ $B^{*i} = \{b_j^{*i}\}$	**1-step** $\Sigma = \{s_j, s_{j+1}(t_j, \Omega)\}$ $B^{*i} = \{b_j^{*i}, b_{j+1}^{*i} \mid h_{j+1}\}$	**Multi-step** $\Sigma = \{s_j, ..., s_N(t_{N-1}, \Omega)\}$ $B^{*i} = \{b_j^{*i}, ..., b_N^{*i} \mid h_N\}$
NO	$L^i = 0$	-	Reinforc. Learning	-	-
	$L^i = 1$	-	Fictitious Play Stationary ξ	Fictitious Play Stationary ξ	Fictitious Play Stationary ξ
	$L^i = 1$	-	Acceptance Rejection	Acceptance Rejection Stationary ξ	Acceptance Rejection Stationary ξ
FULL	$L^i \geq 2$	$L^{-i} \leq 1$	Acceptance Rejection	Optimal Pricing Non-stationary $\xi_{j+1} = \mathrm{f}(b_j^{*i}, \Sigma)$	Optimal Pricing Non-stationary $\xi_N = \mathrm{f}(b_j^{*i}, ..., b_{N-1}^{*i}, \Sigma)$
	$L^i = m$	$2 \leq L^{-i} < m$	Iterated Acceptance Rejection	Iterated Optimal Pricing Non-stationary $\xi_{j+1} = \left\langle \mathrm{f}(b_j^{*i}, \Sigma) \right\rangle^{nL^-}$	Iterated Optimal Pricing Non-stationary $\xi_N = \left\langle \mathrm{f}(b_j^{*i}, ..., b_{N-1}^{*i}, \Sigma) \right\rangle$
IMPERFECT	$L^i = 1$	$L^{-i} \leq 1$	Fictitious Play $\xi_j = \mathrm{f}(\mathrm{p}(\theta_j^{-i} \mid h_j))$	Acceptance Rejection Stationary $\xi_{j+1} = \xi_j$	Acceptance Rejection Stationary $\xi_N = ... = \xi_{j+1} = \xi_j$
	$L^i \geq 2$	$L^{-i} \leq 1$	Fictitious Play $\xi_j = \mathrm{f}(\mathrm{p}(\theta_j^{-i} \mid h_j))$	Optimal Pricing Non-stationary $\xi_{j+1} = \mathrm{f}(b_j^{*i}, \Sigma, \mathrm{p}(\theta_{j+1}^{-i} \mid h_{j+1})$	Optimal Pricing Non-stationary $\xi_N = \mathrm{f}(b_j^{*i}, ..., b_{N-1}^{*i}, \Sigma, \mathrm{p}(\theta_N^{-i} \mid h_N))$
	$L^i = m$	$2 \leq L^{-i} < m$	Iterated Fictitious Play $\xi_j = \left\langle \mathrm{f}(\mathrm{p}(\theta_j^{-i} \mid h_j)) \right\rangle^{nL^{-i}}$	Iterated Optimal Pricing Non-stationary $\xi_{j+1} = \left\langle \mathrm{f}(b_j^{*i}, \Sigma, \mathrm{p}(\theta_{j+1}^{-i} \mid h_{j+1}) \right\rangle$	Iterated Optimal Pricing Non-stationary $\xi_N = \left\langle b_j^{*i}, ..., b_{N-1}^{*i}, \Sigma, \mathrm{p}(\theta_N^{-i} \mid h_N) \right\rangle^{nL^{-i}}$

The problem solving capabilities of the carrier determines the type of problem the carrier solves. For example, a carrier may have imperfect information about the competitors; however, problem solving limitation may force him to solve a myopic problem assuming no-knowledge about the competition. When cost or time limitations are added to the problems, carriers can choose to ignore part of his knowledge in order to get a reasonable answer in a reasonable time, in the spirit of the "satisfying" rule as proposed by Simon (1955). According to Simon, economic agents do not always optimize fully, they optimize up to a satisfying level. Level that depends on personal characteristics and circumstances.

Simplifying (downgrading complexity) the problem due to bounded rational limitations is always possible. It can be interpreted that each problem type (each cell) of table 1 is a different way of measuring how desirable each possible bid is, for a given DVRP technology. If the value of knowledge can be defined as the profit difference that a carrier can obtain going from the no to full knowledge assumption, likewise, the value of computational power is the profit difference that a carrier can obtain from solving a more complex problem due to the increased performance of his computational resources. Summarizing, based on their knowledge level and problem solving capabilities, agents differ in the type of problem they can solve.

7 Learning in a Transportation Marketplace

The reminder of this chapter studies the bidding behavior of carriers in a no-knowledge and no-strategic environments of Table 1. Henceforth, it is assumed that carriers bid trying to maximize their profits but limited by their bounded rational limitations. In this competitive setting, three different auction formats are compared using computational experiments. These auction formats are second price auctions, first price auction with minimum information disclosure, and first price auctions with maximum information disclosure.

The high complexity of acquiring and using knowledge about competitors' behaviors was discussed in Section 5, even in TM market/model that has been streamlined to the indispensable elements. Knowledge acquisition and its use can be considerably more complex in a more complete model where other critical constraints and variables are added (for example, getting drivers home, variation in travel times, delays incurred while unloading the truck, etc). Furthermore, noisy information transmission, as reported by Powell et al. (2002), even among agents that respond to the same carrier (i.e. drivers, dispatchers, decisions support systems), seem to sustain the notion that perfect knowledge about competitors' private information and behavior could only be possible in a flight of the imagination. Imperfect knowledge is possible, but at the cost of even higher modeling complexity. Given the high level of complexity of full or imperfect knowledge assumptions, it is methodologically sensible to first focus on behaviors and settings which are more plausible for implementation in real-life TM marketplaces. The first tool that bounded-rational agents use to cope with insurmountable complexity is simplification. Henceforward, it is assumed that carriers can acquire a limited knowledge of the competitors' behavior. Carriers' knowledge is limited to learn about the distribution of past market prices or the relationships between realized profits and bids.

In an auction context, learning methods seek good bidding strategies by approximating the behavior of competitors. Most learning methods assume that competitors' bidding behavior is stable. This assumed bidding stability is akin to believing that all competitors are in a strategic equilibrium. Walliser (1998) distinguishes four distinct dynamic processes to play games. In a decreasing order of cognitive capacities they are: eductive processes, epistematic learning (fictitious play), behavioral learning (reinforcement learning), and evolutionary processes. An eductive process requires knowledge about competitors' behavior (agents simulate competitors' behavior). Epistemic and behavioral learning are similar to fictitious play and reinforcement learning respectively (fully described in the next section). In the evolutionary process, a player has (is born with) a given strategy; after playing that strategy the player dies and reproduces in proportion to the utilities obtained (usually in a game where it has been randomly matched to another player).

This reminder of this chapter studies the two intermediate types of learning. Eductive-like type of play requires carriers to have almost unbounded computational power and expertise. On the other hand, evolutionary model players seem too simplistic: they have no memory, and simply react in response to the last result. Furthermore, the notion that a company is born, dies, and reproduces with each auction does not fit well behaviorally in the defined TM. Ultimately, neither extreme approach is practically or theoretically compelling in the TM context. Carriers that survive competition in a competitive market like truck-load (TL) procurement cannot be inefficient or unskilled. They are merely limited in the strategies they can implement. It is assumed that carriers would like to implement the strategy (regardless of its complexity) that ensured higher profits, but they are restricted by their cognitive and informational (which give rise to bounded rationality).

In practical and theoretical applications, the process of setting initial beliefs has always been a thorny issue. Implemented learning models must specify what agents initially know. Ideally, how or why these initial assumptions were built should always be reasonable justified or explained. In this respect, restricting the research to the TM context has clear advantages. Normal operating ratios in the TL industry range from 0.90 to 0.95 (TCA, 2003). It is expected that operating ratios in a TM would not radically differ from that range. If prices are too high shippers can always opt out, abandon the marketplace and find an external carrier. Prices cannot be substantially lower because carriers would run continuously in the red, which does not lead to a self-sustainable marketplace. Obviously, operating ratios fluctuations in a competitive market are expected, in response to natural changes in demand and supply. However, these fluctuations should be in the neighborhood of historical long term operating ratios unless the market structure is substantially changed. Another practical consideration is the usage of ratios or factors in the trucking industry. Traditionally, the trucking industry has used numerous factors and indicators to analyze a carrier's performance, costs, and profits. It seems natural that some carriers would obtain a bid after multiplying the estimated cost by a bidding coefficient or factor. Actually, experimental data show that the use of multiplicative bidding factors is quite common (Paarsh, 1991).

7.1 Learning Mechanisms

In reinforcement learning the required knowledge about the game payoff structure and competitors behavior is extremely limited or null. From a single carrier's perspective the situation is modeled as a game against nature; each action (bid) has some random payoff about which the carrier has no prior 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.

Let M be the ordered set of real numbers that are multiplicative coefficients $\mathrm{M} = \{mc_0, ..., mc_K\}$, such that if $mc_k \in \mathrm{M}$ and $mc_{k+1} \in \mathrm{M}$, then $mc_k < mc_{k+1}$. Using multiplicative coefficients the profit obtained for any shipment s_j, when using the multiplicative coefficient mc_k is equal to:

$$\pi_j^i(mc_k) = (mc_k\, c_j^i - c_j^i) I_j^i = c_j^i\, I_j^i\, (mc_k - 1) \qquad (1)$$

$$\pi_j^i(mc_k) = (b_j^{(2)} - c_j^i) I_j^i \qquad (2)$$

The first equation applies to first price auctions while the second equation applies to second price auctions. In the second price auction the payment depends on the value of the second best bid which is represent by the term $b_j^{(2)}$. Adapting the reinforcement model to TM bidding, the probability $\varphi_j^i(mc_k)$ of carrier i using a multiplicative coefficient mc_k in the auction for shipment s_j is equal to:

$$\varphi_j^i(mc_k) = (1 - \lambda\, \pi_{j-1}^i(mc_k))\varphi_{j-1}^i(mc_k) + I_{j-1}^i(mc_k)\, \lambda\, \pi_{j-1}^i(mc_k) \qquad (3)$$

Narenda and Tatcher (1974) showed that a players' time average utility, when confronting an opponent playing a random but stationary strategy, converges to the maximum payoff level obtainable against the distribution of opponents' play. The convergence is obtained as the reinforcement parameter λ goes to zero. To use equation (3), each bidder only needs information about his bids and the result of the auction. To use this model the profits $\pi_{j-1}^i(mc_k)$ must be normalized to lie between zero and one so that they may be interpreted as probabilities. The indicator variable $I_j^i(mc_k)$ is equal to one if carrier i used the multiplicative coefficient mc_k when bidding for shipment s_j, the indicator is equal to zero otherwise. The parameter λ is called the reinforcement learning parameter, it usually varies between $0 < \lambda < 1$.

The reinforcement is proportional to the realized payoff, which is always positive by assumption. Any action played with these assumptions, even those with low performance, receives positive reinforcement as long as it is played (Fudenberg and Levine, 1998). Therefore, a "mediocre" action can be reinforced while at the same time "better" actions are negatively reinforced. Furthermore, in an auction context there is no learning when the auction is lost since $\pi_{j-1}^i(mc_k) = 0 \;\; \forall\, mc_k \in \mathrm{M}$ if $I_{j-1}^i = 0$. Borgers and Sarin (1996) propose a model that deals with the aforementioned problems. In this model the stimulus can be positive or negative depending on whether the realized profit is greater or less than the agent's "aspiration level". If the agent's aspiration level for shipment s_j is denoted ρ_j^i and the effective profit is denoted:

$$\tilde{\pi}_{j-1}^i(mc_k) = \pi_{j-1}^i(mc_k) - \rho_j^i, \tag{4}$$

and the probability becomes:

$$\varphi_j^i(mc_k) = (1 - \lambda\,\tilde{\pi}_{j-1}^i(mc_k))\varphi_{j-1}^i(mc_k) + I_{j-1}^i(mc_k)\,\lambda\,\tilde{\pi}_{j-1}^i(mc_k) \tag{5}$$

When $\rho_j^i = 0$, the equation (5) provides the same probability updating equation as (3). Borgers and Sarin explore the implications of different policies to set the level of the aspiration level. These implications are clearly game dependent. A general observation applies for aspiration levels that are unreachable. In this case equation (4) is always negative; therefore the learning algorithm can never settle on a given strategy, even if the opponent plays a stationary strategy.

These learning mechanisms were originally designed for games with a finite number of actions and without private values (or at a minimum for players with a constant private value). In the TM context, the cost of serving shipments may vary significantly. Furthermore, even the "best" or optimal multiplier coefficient can get a negative reinforcement when an auction is lost simply because the cost of serving a shipment is too high. This negative reinforcement for the "good" coefficient creates instability and tends to equalize the attractiveness of the different multiplicative coefficients. This problem worsens as the number of competitors is increased causing a higher proportion of lost auctions, i.e. negative reinforcement. This chapter utilizes a modified version of the stimulus response model with reinforcement learning that better adapts to TM bidding (Figliozzi, 2004, 2005). Each multiplicative coefficient m_k has an associated average profit value $\overline{\pi}_j^i(m_k)$ that is equal to:

$$\overline{\pi}_j^i(m_k) = \frac{\sum_{t \in \{1,\dots,j\}} \pi_t^i(s_t)\; I_t^i(m_k)}{\sum_{t \in \{1,\dots,j\}} I_t^i(m_k)}$$

The aspiration level is defined as the average profit over all past auctions:

$$\overline{\rho}_j^i = \frac{\sum_{t \in \{1,\dots,j\}} \pi_t^i(s_t)\; I_t^i}{j}$$

Therefore the average effective profit is defined as $\overline{\pi}_{j-1}^i(mc_k) = \overline{\pi}_{j-1}^i(mc_k) - \overline{\rho}_j^i$. Probabilities are therefore updated using equation (6).

$$\varphi_j^i(mc_k) = (1 - \lambda\,\overline{\pi}_{j-1}^i(mc_k))\varphi_{j-1}^i(mc_k) + I_{j-1}^i(mc_k)\,\lambda\,\overline{\pi}_{j-1}^i(mc_k) \tag{6}$$

With the latter formulation (6), a "good" multiplicative coefficient does not get a negative reinforcement unless its average profit falls below the general profit average. At the same time, there is learning even if the auction is lost. The learning mechanism that uses equation (6) is named as Average Reinforcement Learning (ARL) henceforth.

Stimulus-response learning requires the least information and can be applied to both first and second price auctions. The probability updating equations (3) and (6) are the same for first and second price auctions. Therefore the application of the reinforcement learning model does not change with the auction format that is being utilized in the TM. Using this learning method, a carrier does not need to model neither the behavior nor the actions of competitors. The learning method is essentially myopic since it does not attempt to measure the effect of the current auction on future auctions. The method clearly fits in the category of no-knowledge/myopic carrier bounded rationality. Since the method is myopic, for the first price auction the multiplicative coefficients must be equal or bigger than one, i.e. $mc_0 \geq 1$. A coefficient smaller than one, generates only zero or negative profits. In a second price auction the multiplicative coefficients can be smaller than one and still generate positive profits since the payment is dependent on the competitors' bids. In both types of auctions it is necessary to specify not just the set of multiplicative coefficients but the initial probabilities. If equation (5) is used it is also necessary to set the aspiration level. If equation (6) is used it is necessary to set the level of the initial profits but not the aspiration level. A uniform probability distribution is the classical assumption and indicates a complete lack of knowledge about the competitive environment.

Summarizing, in reinforcement learning, the agent does not consider strategic interaction. The agent is unable to model an agent play or behavior but his own. This agent is informed only by his past experiences and is content with observing the sequence of their own past actions and the corresponding payoffs. Using only his action-reward experience, he reinforces strategies that succeeded and inhibit strategies which failed. He does not maximize but moves in a utility-increasing direction, by choosing a strategy or by switching to a strategy with a probability positively related to the utility index. Reinforcement learning (and its variants) is a strategy that is designed to operate in an environment where the player (carrier) is unable to see the competitors' actions. Therefore, it is able to strongly reinforce (positively or negatively) only one action: the last action played. Unlike reinforcement learning, fictitious play requires the observation of competitors' actions. A good introduction to types of learning employed in this chapter (reinforcement learning and fictitious play) can be found in the work of Fudenberg and Levine (1998).

Fictitious play came about as an algorithm to look for Nash equilibrium in finite games of complete information (Brown, 1951). It is assumed that the carrier observes the whole sequence of competitors' actions and draws a probabilistic behavioral model of the opponents' actions. The agent does not try to reveal his or her opponents' bounded rationality from their actions although the agent may eventually know that opponents learn and modified their strategies too. The agent models not behavior but simply a distribution of opponents' actions. Players do not try to influence the future play of their opponents. Players behave as if they think they are facing a stationary, but unknown, distribution of the opponents' strategies. Players ignore any dynamic links between their play today and their opponents' play tomorrow. A player that uses a generalized fictitious play learning scheme assumes that his opponents' next bid vector is distributed according to a weighted empirical distribution of their past bid vectors. The method cannot be straightforwardly adapted to games with an infinite set of strategies (for example the real numbers in an auction). Two ways of tackling this problem are: a) the player divides the set of real numbers into a finite number of subsets, which are then associated with a strategy or b) the player uses a probability distribution, defined over the set of real number to approximate the probabilities of competitors play. In either case, the carrier must come up with a estimated stationary price function ξ (in our experiments carriers estimate a normal distribution using on competitors' past bids). If a second price auction format is used in the TM, the carrier bids using:

$$b_j^{*i} \in \underset{b \in R}{\arg\max}\ E_{(\xi)}\{[\xi - c^i(s_j, z_j^i)]I_j^i\} \tag{7}$$

If a first price auction format is used in the TM, the carriers bid using:

$$b_j^{*i} \in \underset{b \in R}{\arg\max}\ E_{(\xi)}\{\ [(b - c^i(s_j, z_j^i)]I_j^i\} \tag{8}$$

In the second price auction (equation 7) the best price is simply the corresponding cost $c^i(s_j, z_j^i)$ due to the special properties of one-item second price auctions (8) (independence between the winners bid and the corresponding payment). Equation 8 has to be solved numerically or analytically.

7.2 Automaton Interpretation

The previous sections have described reinforcement learning and fictitious play models of learning. Reinforcement learning and fictitious play were originally conceived as human methods of learning. However, they can also be used by machines or computerized systems. This section tries to link both views. An automaton is a self operating machine or mechanism. In a game context, an automaton is meant to be an abstraction of the process by which a player *implements* a given bounded rationality behavior. Rubenstein (1998) replaces the notion of a strategy with the notion of a machine called finite automaton. In Rubenstein's model a finite automaton that represents player i is a four-tuple $(Z^i, z_0^i, \mathrm{b}^i, \mathrm{a}^i)$, where Z^i is a finite set of machine states (from this constraint the adjective "finite"), z_0^i is the initial state for carrier i, $\mathrm{b}^i : Z^i \to A$ is an output function that produces an action (given the state of the automaton), and $\mathrm{a}^i : Z^i \times A^{-i} \to Z^i$ is a transition function that updates the state of the automaton (given the actions taken by the competitors in the previous period). The set of possible actions is denoted by A. Adapting these concepts to this research, a TM automaton can be defined as an abstraction of the process by which a carrier implements a given bounded rational behavior in a TM. A TM automaton can be defined by the eight-tuple $(Z^i, z_0^i, \Xi, \xi_0^i, S, \mathrm{b}^i, \mathrm{u}^i, \mathrm{a}^i)$ comprised by:

Z^i the set of possible states (private information states) ;

z_0^i the initial state for carrier i;

Ξ the set of possible price functions;

ξ_0^i the initial price function for carrier i;

$s_j \in S$ the stimulus sent by marketplace;
$\mathrm{b}^i : Z^i \times \Xi \times S \to R$ the bidding (output) function;
$\mathrm{u}^i : h \times \Xi \to \Xi$ the update function (updates the price function $\xi \in \Xi$) ; and
$\mathrm{a}^i : Z^i \times S \to Z^i$ the assignment function (assignment if an auction is won).

A TM automaton would work in the following way: the initial state and price function are z_0^i and ξ_0^i respectively, the automaton chooses a bid $\mathrm{b}^i(z_0^i, \xi_0^i, s_1)$ when the first shipment arrives. If carrier i wins, the assignment function updates the carrier's status $\mathrm{a}^i(z_0^i, s_1)$. The price function is updated based on the information revealed after the auction $\mathrm{u}^i(h_1, \xi_0^i)$. When the second shipment arrives the same process is repeated but starting with the new state and price function z_1^i and ξ_1^i respectively. Once the initial conditions are set, the transitions, bidding, and updating are set by the arrival of shipments. A TM automata game takes place when a player cannot change the working of his machine during the course of the game. The two learning approaches described in this section, reinforcement learning and fictitious play, can be interpreted as the work of an automaton (which is valid in general for any learning strategy that seeks or uses no knowledge about the competitors' behavior). Therefore, the simulation results presented in the next sections can also be interpreted as the interaction or competition of TM automata (which may represent the behavior of human, computerized, or hybrid dispatchers). It is assumed in this research that for a given status, price function, and stimulus, an action has the same probability of being played; as if the decision process is *wired-up* and cannot change (data and information can change over time, but not the decision-making process). This is consistent (in the short-medium term) with the industry experience (Powell et al., 2002).

8 Experimental Results

Closed analytical solutions for the complex carriers' decision problem in a TM setting would require many simplifications that could compromise the validity of the results. Therefore, computational experiments and simulation are used as needed to enhance and extend simpler theoretical models. Furthermore, simulation is used to study the dynamics of carriers' behaviors and interactions in controlled and replicable experiments.

8.1 Simulation Framework

The following sections study truck-load (TL) carriers that compete over a square area; the sides' lengths are equal to 1 unit of distance. For convenience, trucks travel at constant speed equal to one unit of distance per unit of time. Demands for truckload pickup-and-delivery arise over this area and over time. Origins and destinations of demands are uniformly distributed over the square area, so the average loaded distance for a request is 0.52 units of distance. All the arrivals are random; the arrival process follows a time Poisson process. The expected inter-arrival time is E [T] = 1/ (Kλ), where λ is the demand request rate per vehicle and K is the total market fleet size. The total market fleet size that was used in the results is 4 (though similar trends were obtained with larger fleets – 8 vehicles – as long as the same arrival rate/fleet size ratio is used). Roughly, the average service time for a shipment is 0.77 units of time (approximately $\lambda = 1.3$). The service time is broken down into 0.52 units of time corresponding to the average loaded distance, plus 0.25 units of time that approximate the average empty distance (average empty distance vary with arrival rates and time windows considered). Different Poisson arrival rates per truck per unit of time are simulated (ranging from 0.5 to 1.5). As a general guideline, these values correspond to situations where the carriers are:

- $\lambda = 0.5$ (uncongested)
- $\lambda = 1.0$ (congested)
- $\lambda = 1.5$ (extremely congested)

The shipments have hard time windows. In all cases, it is assumed that the earliest pickup time is the arriving time of the demand to the marketplace. The latest delivery time (LDT) is assumed to be:

LDT = arrival time + **2** x (shipment loaded distance + 0.25) + **2** x uniform (0.0, 1.0).

All the shipments have a reservation price distributed as uniform (1.42, 1.52). In all cases, reservation prices exceed the maximum marginal cost possible in the simulated area ($\approx$1.41 units of distance). It is also assumed that all the vehicles and loads are compatible; no special equipment is required for specific loads. In all the simulations, two carriers are competing for the demands. In all cases there is an initial warm up or learning period of 250 auctions.

Multiple performance measures are used. The first is total profits, which equal the sum of all payments received by won auctions minus the empty distance incurred to serve all won shipments (it was already mentioned that shipment loaded distances are not included in the bids, loaded distances cancel out when computing profits). The profit for a particular shipment is defined as the difference between the payment received and the increment of the empty distance cost necessary to serve this shipment. The second performance measure is number of auctions won or number of shipments served, an indicator of market share. The third is shippers' consumer surplus, which is the accumulated difference between reservation prices and prices paid. The fourth is total wealth generated that is equal to the accumulated difference between reservation price (of served shipments) and empty distance traveled.

The second price auction used in the TM operates as follows: (a) each carrier submits a single bid, (b) the winner is the carrier with the lowest bid (which must be below the reservation price; otherwise the auction is declared vacant), (c) the item (shipment) is awarded to the winner, (d) the winner is paid either the value of the second lowest bid or the reservation price, whichever is the lowest, and (e) the other carriers (not winners) do not win, pay, or receive anything. The same procedure applies to first price auctions but the winner is paid the value of the winning bid, only point (d) changes.

In real time situations, this is an increasingly difficult task when optimal decision-making involves the solution of larger NP hard problems and the necessity of taking into account the stochastic nature of future demands. Three levels of DVRP technologies were simulated. These technologies are presented in an order that shows an increasing level of sophistication.

1. Base or Naïve Technology: this type of carrier simply serves shipments in the order they arrive. If the carrier has just one truck, it estimates the marginal cost of an arriving shipment s_j simply as the additional empty distance incurred when appending s_j to the end of the current route. If the carrier has more than one truck, the marginal cost is the cost of the truck with the lowest appending cost. This technology does not take into account the stochastic or combinatorial aspect of the cost estimation problem and is considered one of the simplest possible. Nonetheless, it provides a useful benchmark against which to compare the performance of more complex and computationally demanding technologies.
2. Static Fleet Optimal (SFO): this carrier optimizes the static vehicle routing problem at the *fleet* level. If the carrier has just one truck, the technology is equivalent to the previous case. If the carrier has more than one truck, the marginal cost is the increment in empty distance that results from *adding* s_j to the *total pool of trucks and loads* yet to be serviced. If the problem where static, this technology would provide the optimal cost. Again, like the two previous technology, it does not take into account the stochastic nature of the problem. This technology roughly stands for "the best" a myopic (as ignoring the future but with real time information) fleet dispatcher can achieve. Carriers fleet assignment and cost estimation is based on the static optimization based approach proposed by Yang et al. (2004).
3. One step Look ahead Fleet Optimal (1SLA): as the previous carrier, this carrier optimizes the static vehicle routing problem the fleet level. This provides the static marginal cost for adding s_j. However, this carrier also knows the distribution of load arrivals over time and their spatial distribution. Hence, the carrier can simulate whether and how much winning s_j affects the marginal cost of serving the next arriving load. This technology roughly stands for what a fleet dispatcher with real time information and knowledge of future (yet unrealized probabilistic demands) can do. However, 1SLA is not an "optimal" technology, rather it is a heuristic that tries to estimate how serving s_j affects the cost of serving the next shipment.

8.2 Analysis of Experimental Results

A significant characteristic of one-item second price auction is also cost bidding, i.e. one-item second price auctions are incentive compatible mechanisms. That characteristic cannot be necessarily maintained in multiunit sequential auctions setting such as the TM marketplace. Of the two learning methods proposed, only reinforcement learning can be applied to second price auctions since fictitious play in a single-item second price auction coincides with marginal cost bidding. Regardless of the price distribution, the expected profit is always optimized with marginal cost bidding.. In the TM context, the objective of reinforcement learning is to "learn" what the best bidding coefficient is; the bidding coefficient that maximizes a carrier's profits. Which raises the question: in a TM second price auction environment can carriers be better off by using bidding factors? This question is answered using computational experiments. Two carriers using the same type of DVRP technology compete against each other. However, while one carrier bids the marginal cost (called MC carrier) the other bids the marginal cost multiplied by a bidding factor (called BF carrier). Eleven different bidding factors are utilized, ranging from 0.5 to 1.5. The impact of these factors on carrier BF's profits are depicted in figure 2. The profit levels of a BF carrier when the bidding factor is equal to 1.0 are used as the reference or base level – they correspond to 100% level. Both carriers are using the SFO technology.

8.3 Performance of Marginal Cost Bidding

The results depicted in figure 2 show that for low arrival rates the best bidding factor is 1.0, corresponding to simply bidding the marginal cost. For medium arrival rates the best bidding factor is 1.1. For high arrival rates the best bidding factor is 1.3. Regardless of the arrival rate level, the "curve" is quite flat around the "optimal". Furthermore, if the profits are connected the resulting curve is concave-shaped. A possible explanation to the results of figure 2 may be obtained by analyzing how profits are generated. Total profits can be expressed as the average profit obtained per shipment multiplied by the number of shipments served. Figure 3 and figure 4 show the impact of bidding factors on number of shipments served and average shipment served profit respectively. Again, the number of shipments served and average profit used as reference are those of a BF carrier when the bidding factor is equal to 1.0.

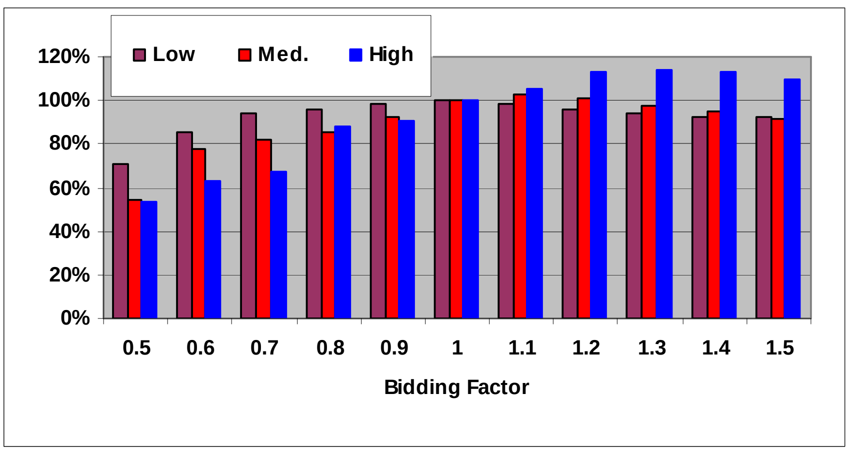

Figure 2: Profit Level for a BF Carrier

It is clear from figure 3 and figure 4 that, as expected, higher bidding factors increase the average profit per shipment won but decreases the number of shipments won. Vice versa, lower bidding factors decrease the average profit per shipment won but increases the number of shipments won. There are clearly two opposing forces at work when the bidding factor changes; this fact helps to explain the concave shape of the profit curve in figure 2.

At this point, it has not yet been explained why the low arrival rate "optimal" bidding factor is around 1.0 (marginal cost case), while the "optimal" bidding factors are shifted to the right for higher arrival rates. The answer to this matter lies in the relation between profit elasticity and shipment served volume elasticity. To understand why profit elasticity and shipment served volume elasticity changes with the arrival rate is necessary to introduce figure 5 and figure 6. They illustrate the different fleet utilization rates of carriers MC and BF respectively. Fleet utilization rate is defined as the average vehicle utilization. Vehicle utilization is defined as the percentage of the time a vehicle is moving (i.e. not idle).

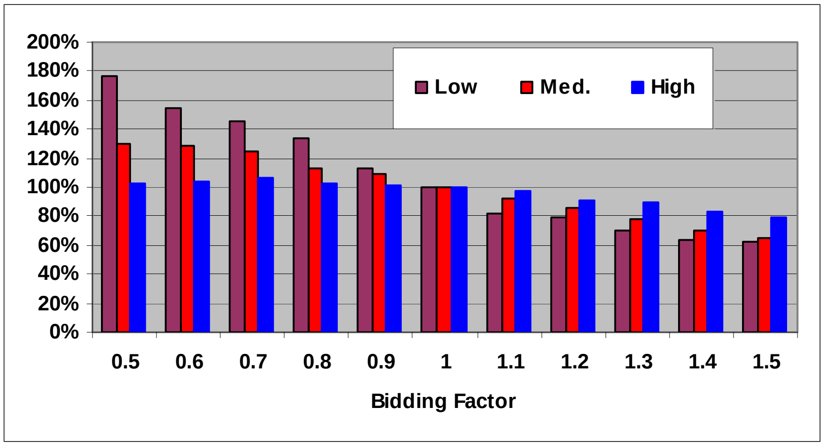

Figure 3: Shipments Served by BF carrier

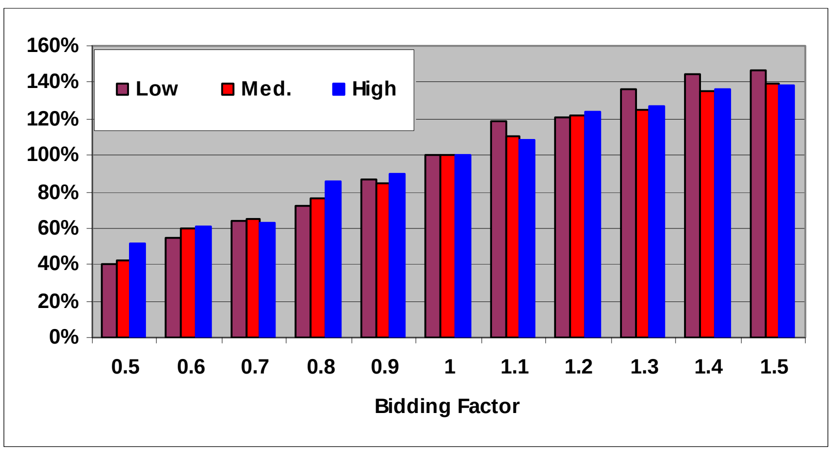

Figure 4: Average Profit per Shipment Won for a BF Carrier

With low arrival rates the utilization of the MC carrier is low (around 35% if the BF carrier uses a bidding factor equal to 1.0 - see figure 5). Therefore when carrier BF increases his prices (utilizing higher bidding factors) carrier MC gains a significant percentage of the demand. This explains why in figure 6 there is such an abrupt drop in demand (from 100 to 80%) when carrier BF moves from a bidding factor of 1.0 to 1.1. With higher arrival rates the fleet utilization of carrier MC is higher (at or over 70% - see figure 5) and at very high utilization rates it is more difficult to accommodate or to inexpensively add new shipments. As fleet utilization grows the capacity to serve new shipments decreases, therefore on average the opportunity costs of serving additional shipments starts to be significant. Figure 6 is the reverse mirror image of figure 5. With high arrival rates carrier BF can rise prices substantially and still have a high fleet utilization; the increase in profits prevails over the decrease in shipments served. The explanation provided is plausible but not definitive. However, similar phenomena as the ones observed in figure 2-6 have been widely recognized in the economics-industrial organization literature.

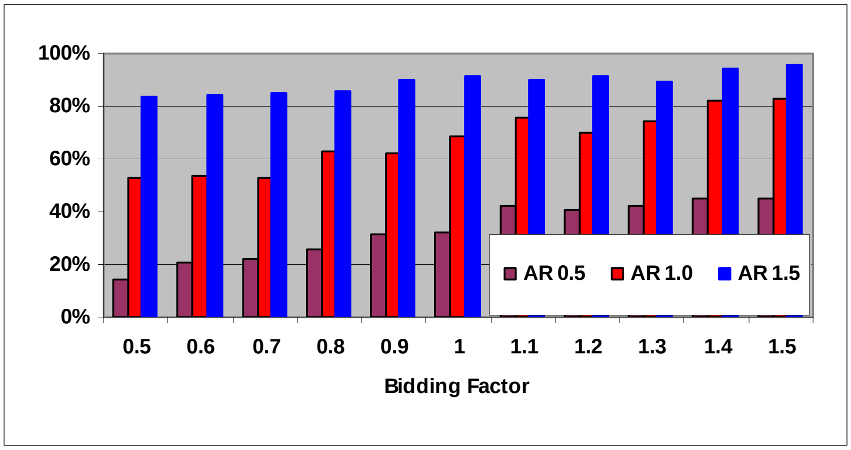

Figure 5: Fleet Utilization (MC Carrier)

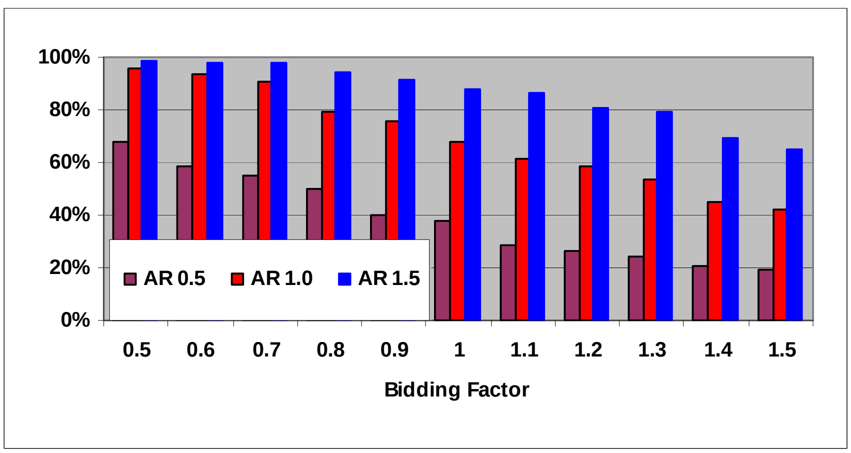

Figure 6: Fleet Utilization (BF Carrier)

The incentives to increase prices as remaining market capacity decreases are contemplated in price-capacity oligopoly models. For example, in the Edgeworth-Bertrand model of competition, pricing is at marginal cost levels when demand is low, however prices increase after a critical capacity utilization threshold is surpassed (Martin, 1993). Similar intuition is obtained from Benoit and Krishna (2001) model of capacity constrained auctions, with limited capacity it is advantageous to speculate. Even in fleet management, the idea of filtering out shipments or similarly increasing the "admission" price of shipments under very high arrival rate conditions has been previously used (though not in a competitive environment). The Kim et al. (2002) study indicates that a fleet dispatcher under very high arrival rates (over capacity) is better off filtering out some demands (not being too close to capacity). Similar results are also found when carriers use other technologies such as the naïve or 1SLA. Figure 7 shows the profit changes when both carriers use naive technologies. Even when carriers have different technologies, similar results can be expected. Figure 8 show the profit changes for the BF carrier using naïve technology against a MC carrier using SFO technology.

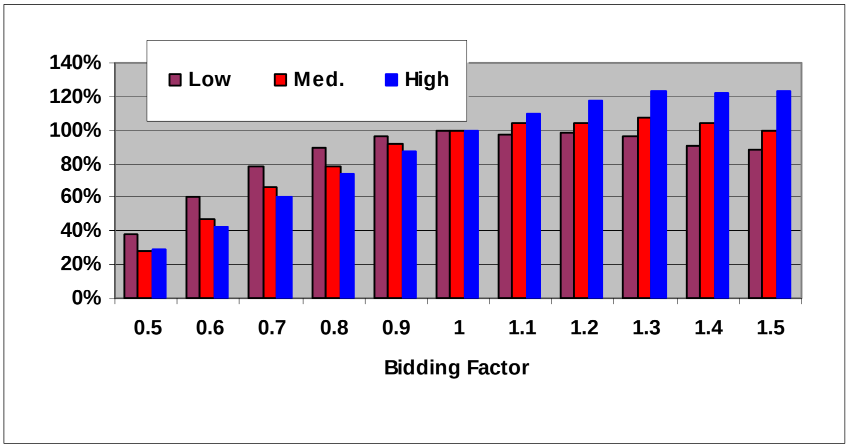

Figure 7: Profit Level for a BF Carrier (both carrier use naïve technology)

The question that motivated these simulations was: in a TM second price auction environment can carriers be better off by using bidding factors? The answer is yes, but only at high arrival rates. This answer provides additional insights into the applicability of auction analysis to online algorithms/technologies. The results confirm the notion that DVRP technological leadership can be better exploited under low to moderate arrival rate conditions, where there is no incentive to adopt bidding factors that are not one. If there is an incentive to adopt bidding factors that are higher than one, there is an incentive to restrain capacity or to increase prices (profits are increased without increasing fleet management efficiency). As the arrival rate grows the advantage of being more efficient decreases; in general, scarcity exposes the incompetent while abundance hides inefficiencies.

8.4 Learning Methods Performance

The following results address the issue of learning performance of the two learning methods presented in this chapter. The previous results show that bidding factors can be used to increase carriers' profits in TM second price auctions with high arrival rates. Reinforcement learning could be used to "learn" which bidding factors produce a higher profits on average; as the auction results accumulates the most profitable bidding factors continuously increase their probability of being used. With low arrival rates, there is nothing to learn but the fact that marginal cost bidding (bidding factor 1.0) is the best alternative. Learning can be expensive though. For example, in a second price auction the longer it takes a bidder to learn that underbidding (bidding below his marginal costs) is not a good strategy, the more the bidder loses potential profits. The importance of the right learning coefficient then becomes evident. If the learning coefficient λ is too small learning is too slow; if λ is too big it may lock the learning algorithm in an undesirable bidding factor too quickly. Another important element is the number of alternatives that the learning algorithm must choose from; as a general rule, the more the alternatives the smaller the λ.

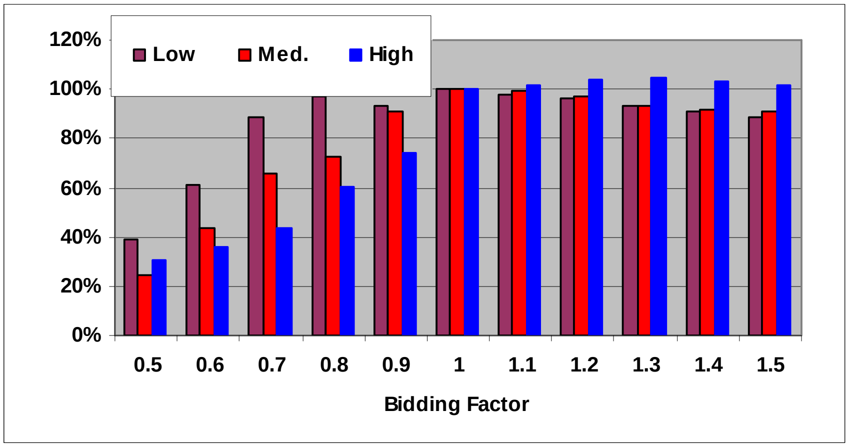

Figure 8: Profit Level for a BF Carrier (SFO vs. naïve technology)

The speed of reinforcement learning can be quite slow in an auction setting like TM. The "optimal" bidding factor can be used and there is still roughly a 50% chance of losing (assuming two bidders with equal fleets and technologies). If the "optimal" bidding factor loses two or three times its chances of being played again may reduce considerably which hinders convergence to the "optimal" or even convergence at all. As discussed previously in this section, this issue can be avoided using "averages" (ARL method). Figure 9 illustrates the relative performance of Average Reinforcement Learning (ARL) and Reinforcement Learning (RL) in a first price auction. Both learning methods select a bidding factor among 11 different possibilities, ranging from 1.0 to 2.0 in intervals of 0.1. The learning factor is $\lambda = 0.10$. Figure 9 shows the relative performance of ARL and RL after 500 auctions. It is clear that RLA obtains higher profits as the arrival rate increases. RL has a poorer performance because it cannot converge steadily to the "optimal" coefficient due to the reasons mentioned in the previous paragraph. The carrier RL tends to price lower (it keeps probing low bidding coefficients longer) and therefore serves a higher number of shipments. As shown in the previous section, as arrival rates increase after a critical point, a carrier can charge higher prices regardless of what the competitor is doing.

In first price auctions reinforcement learning and fictitious play can be used. The latter uses more information than the former. Therefore, it is expected that a carrier using fictitious play must outperform a carrier using reinforcement learning. Figure 10 shows the relative performance of Fictitious Play (FP) and ARL after 500 auctions. The ARL player is the same as in figure 9. The FP carrier divides the possible competitors' bids in 15 intervals (from 0.0 to 1.5 in intervals of width 0.1) and start with a uniform probability distribution over them.

Clearly the FP carrier obtains higher profits across the board. The usage of a competitor past bidding data to obtain the bid that maximizes expected profits clearly pays off. In this case carrier ARL tends to bid less and serve more shipments, again, the difference diminished as the arrival rate increase. In the TM context even a simple static optimization provides better results than a search based on reinforcement learning. Not surprisingly, more information and optimization lead to better results. Therefore, if there is maximum information disclosure, carriers will choose to play fictitious play or a similar bidding strategy, especially since the complexity of FP (myopic) and ARL are not too different.

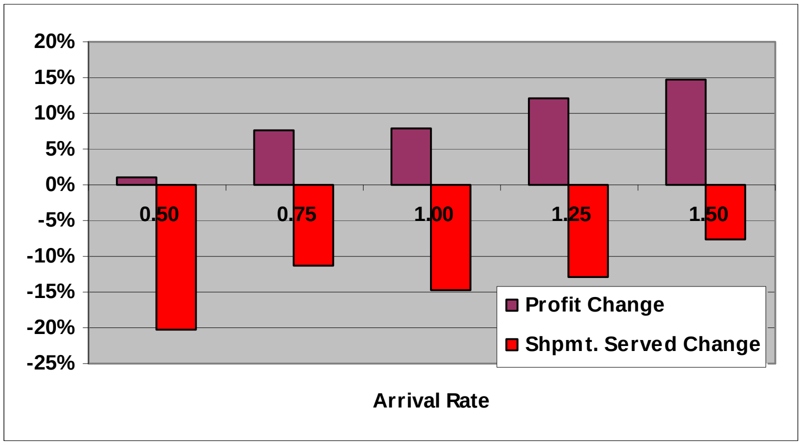

Figure 9: ARL vs. RL (RL performance base of comparison)

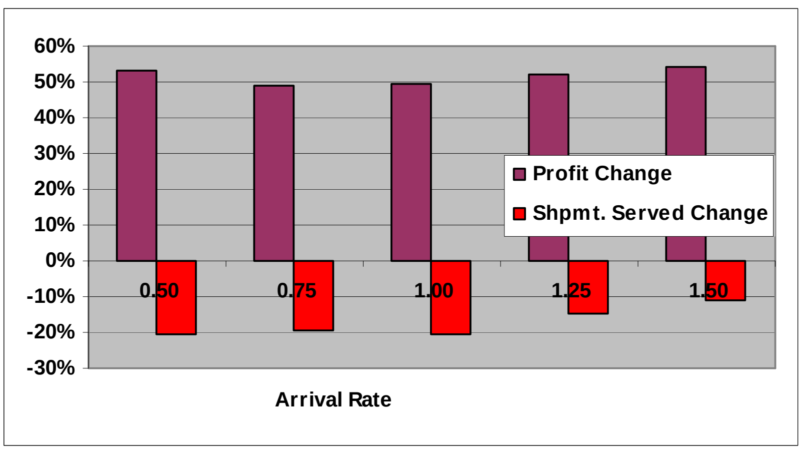

Figure 10: ARL vs. FP (RL performance base of comparison)

8.5 Comparing Auction Settings

The following results describe the outcomes of TM competition with different sequential auction settings. Within the competitive no-knowledge assumptions, three basic auction settings are compared: second price auction with marginal cost bidding, first price auction with reinforcement learning, and first price auction with fictitious play. Four different measures are used to compare the auction environments: carriers' profits, consumer surplus, number of shipments served, and total wealth generated.

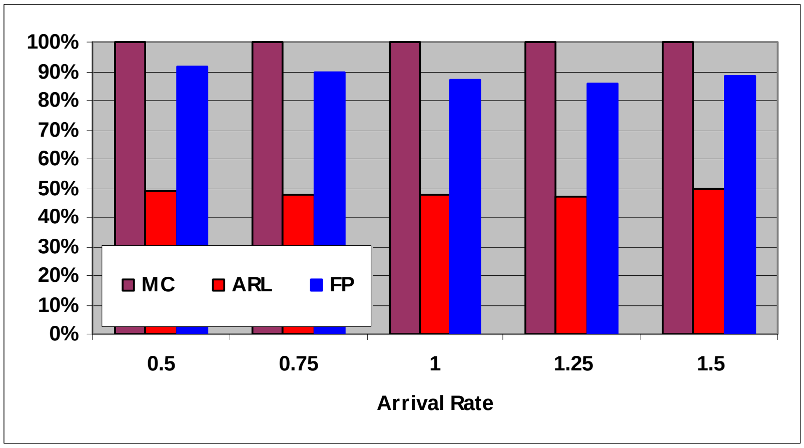

Figure 11: Carriers' Profit level (Second Price Auction MC as base)

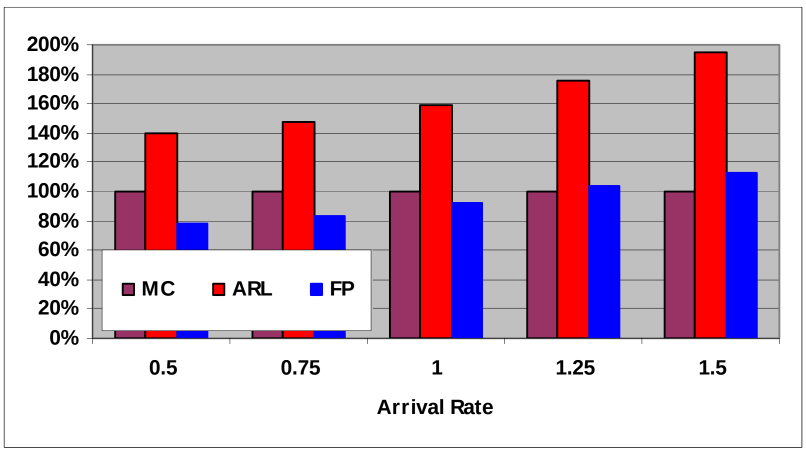

Figure 12: Consumer Surplus level (Second Price Auction MC as base)

To facilitate comparisons in all the four graphs that are presented subsequently, second price auctions with marginal cost bidding are used as the standard to measure up the two types of first price auction. All two carriers use SFO technologies. Figure 11 illustrates the profits obtained by carriers. After the results of the previous section, it is not surprising that FP carriers obtain higher profits than ARL carriers. FP carriers use the obtained price information to their advantage. The highest carrier profit levels takes place with the second price auctions. These results do not alter or contradict theoretical results. With asymmetric cost distribution functions, Maskin and Riley (2000) show that there is not revenue ordering between independent value first and second price auctions. Figure 12 illustrates the consumer surplus obtained with the three auction types. Clearly, first price auction with reinforcement learning (minimum information disclosed) benefit shippers. Unsurprisingly, figure 12 is almost the reverse image of figure 11.

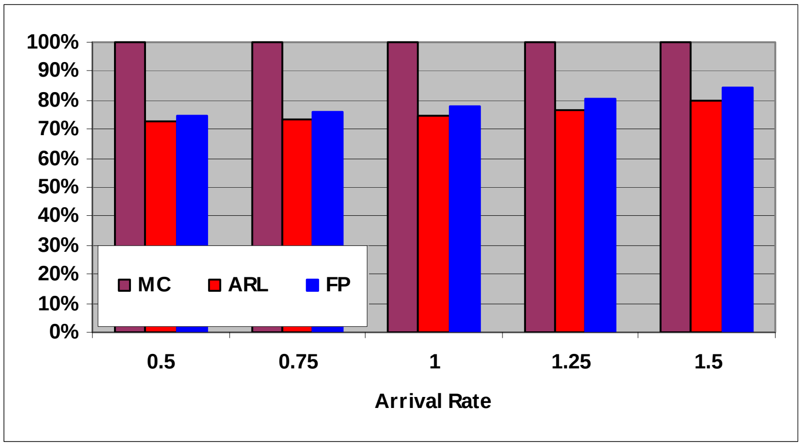

Figure 13: Number of Shipments Served (Second Price Auction MC as base)

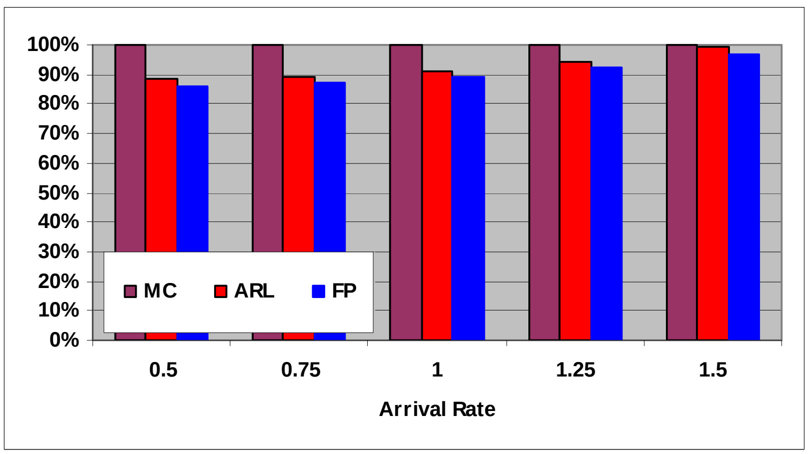

Figure 14: Total Wealth Generated (Second Price Auction MC as base)

Figure 13 shows the number of shipments served with each auction setting. As expected, with second price auctions more shipments get served. Even in asymmetric auctions, it is still a weakly dominant strategy for a bidder to bid his value in a second price auction – recall that this property of one-item second price auction is independent of the competitors' valuations. Therefore, in the second price auction the shipment goes to the carrier with the lowest cost. In contrast, with ARL there is a positive probability that there are inefficient assignments since a higher cost competitor can use a bidding coefficient that results in a lower bid. Similarly with FP carriers, if the price functions are different (which is very likely since each carrier models the competitors' prices), a lower cost carrier can be underbid by a higher cost carrier with a positive probability. The results of figure 12 and figure 13 are similar to the insights provided by the reverse auction model with elastic demand (Wolfstetter, 1999), where introducing higher price uncertainty decreases prices (carriers' profits) but also decreases the probability of completing a potentially feasible transaction (number of shipments served). Figure 14 shows the wealth generated with each auction setting. Predictably, with second price auctions more wealth is generated. This is not surprising since marginal cost bidding is a "price efficient" mechanism. As the arrival rate increases the gap in total wealth generated tends to close up (figure 14). Consistently, the lowest wealth generated corresponds to the case with FP bidders.

Summarizing, under the current TM setting, carriers, shippers, and a social planner would each select a different auction setting. Carriers would like to choose a second price auction. If first price auction are used, carriers would like to have maximum information disclosure. More information allows players to maximize profits, though total wealth generated is the lowest. Shippers would like to choose a first price auction with minimum information disclosure; more uncertainty about winning leads carriers to offer lower prices. However, the uncertainty leads to a reduction in the number of shipments served. Finally, from society viewpoint the most efficient system is the second price auction. More shipments are served and more wealth is generated.

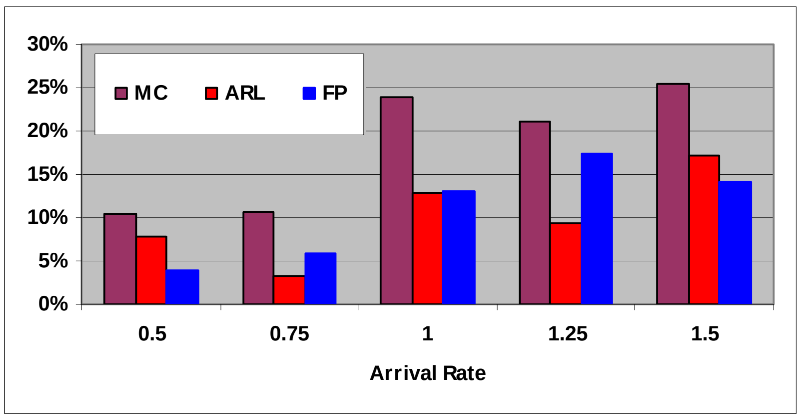

Figure 15: Impact of Auction Type and Technology upgrading on Profits

8.6 Auction Settings and DVRP Technology Benefits

The final set of experiments looks at how auction settings impact the competitive edge that a more sophisticated DVRP can provide. Figure 15 illustrates the profit improvement of a carrier using a SFO technology over a carrier using the naïve technology. As expected, the second price auction better rewards a lower cost carrier. Again, this can be attributed to the lack of speculation about prices, which removes unnecessary speculation about competitors.

9 Conclusions

A competitive TM setting was analyzed to determine the likely sources of bounded rationality and the context of carriers' decision making process. Given the complexity of the bidding/fleet management problem, carriers can tackle it with different levels of sophistication. The complexity of the different bidding problems that a bounded rational carrier can be faced with was analyzed and classified. In the framework presented, sequential auctions can be used to model an ongoing transportation market, where the effect of carrier competition, knowledge and information availability, dynamic vehicle routing technologies, computational power, and decision making processes can be studied. This is an alternative framework to traditional models of behavior, equilibrium, decision-making, and analysis for transportation carriers. Decision making and behavior is defined as an expression of the goals and bounded rationality of the carrier as the type of pricing/bidding/fleet management problem that the carrier is able to tackle. Table 1 coupled with the appropriate learning mechanisms (for example reinforcement learning and fictitious play when they suit) embody the approach to carrier behavior proposed in this research.

Reinforcement learning and fictitious play, two learning methodologies for this type auction setting and assumptions are introduced and analyzed, as well as carrier learning and behavioral assumptions. Carrier's behavior is compared with the behavior of a machine. Computational experiments indicate that auction setting and information disclosure matters. Maximum information disclosure allows carriers to maximize profits at the expense of shippers' consumer surplus; minimum information disclosure allows shippers to maximize consumer surplus but at the expense of lowering the number of shipments served. Marginal bidding in second price auctions remains the most efficient incentive compatible auction mechanism, producing more wealth and more shipments served than first price auctions. It is demonstrated that under critical arrival rate there is no incentive to use bidding factors (no deviations from static marginal cost bidding). Furthermore, second price auction TM is the mechanism that provides the highest reward to carriers with more sophisticated DVRP technology.

References

[1] AUMANN, R. (1997) Rationality and Bounded Rationality. Games and Economic Behavior, V21, pp 2-14.

[2] BENOIT, J. & KRISHNA, V. (2001) Multiple-Object Auctions with Budget Constrained Bidders. Review of Economic Studies, V 68, pp 155-179.

[3] BLUME, A. & HEIDHUES, P. (2006) Private monitoring in auctions. Journal Of Economic Theory, 131, 179-211.

[4] BORGES, T. & SARIN, T. (1996) Naïve reinforcement learning with endogenous aspirations. Mimeo, University College of London.

[5] BROWN, G. (1951) Iterative Solutions of games by fictitious play. In Activity Analysis of Production and Allocation. ed. by Koopmans T., Wiley. New York.

[6] CAMERER, C. (1995) "Individual Decision Making". in Kagel, John and Roth, Alvin eds., The Handbook of Experimental Economics. Princeton Princeton NJ, Princeton University Press.

[7] COLINSK, J. (1996) Why Bounded Rationality. Journal Economic Literature, N34, pp 669-700.

[8] FIGLIOZZI, M. (2004) Performance and Analysis of Spot Truck-Load Procurement Markets Using Sequential Auctions. Ph. D. Thesis, School of Engineering, University of Maryland College Park.

[9] FIGLIOZZI, M. (2006) Analysis and Evaluation of Incentive Compatible Dynamic Mechanisms for Carrier Collaboration. Transportation Research Record 1966**,** 34-40.
[10] FIGLIOZZI, M., MAHMASSANI, H. & JAILLET, P. (2003a) Framework for study of carrier strategies in auction-based transportation marketplace. Transportation Research Record 1854**,** 162-170.
[11] FIGLIOZZI, M., MAHMASSANI, H. & JAILLET, P. (2003b) Modeling Carrier Behavior in Sequential Auction Transportation Markets. 10th International Conference on Travel Behaviour Research (IATBR), August 2003.
[12] FIGLIOZZI, M., MAHMASSANI, H. & JAILLET, P. (2004) Competitive Performance Assessment of Dynamic Vehicle Routing Technologies Using Sequential Auctions. Transportation Research Record 1882**,** 10-18.
[13] FIGLIOZZI, M., MAHMASSANI, H. & JAILLET, P. (2005) Auction Settings and Performance of Electronic Marketplaces for Truckload Transportation Services. Transportation Research Record 1906**,** 89-97.
[14] FIGLIOZZI, M., MAHMASSANI, H. & JAILLET, P. (2006) Quantifying Opportunity Costs in Sequential Transportation Auctions for Truckload Acquisition. Transportation Research Record 1964**,** 247-252.
[15] FIGLIOZZI, M., MAHMASSANI, H. & JAILLET, P. (2007) Pricing in Dynamic Vehicle Routing Problems. Forthcoming Transportation Science.
[16] FUDENBERG, D. & LEVINE, D. (1998) The Theory of Learning in Games, MIT Press, Cambridge Massachusetts.
[17] HARSANYI, J. (1967) Games with incomplete information played by Bayesian players. Management Science, V14, pp 159-182 and 320-334.
[18] KAGEL, J. R. A. (1995) Handbook of Experimental Economics. Princeton NJ. Princeton University Press.
[19] KIM, Y., MAHMASSANI, H. S. & JAILLET, P. (2002) Dynamic truckload truck routing and scheduling in oversaturated demand situations. Transportation Network Modeling 2002. Washington, Transportation Research Board Natl Research Council.
[20] KRISHNA, V. (2002) Auction Theory. Academic Press, San Diego, USA.
[21] LAFFONT, J. J. (1997) Game theory and empirical economics: The case of auction data. European Economic Review, 41**,** 1-35.
[22] LUCKING-REILEY, D. & SPULBER, D. (2001) Business-to-Business Electronic Commerce. Journal of Economic Perspectives, vol. 15, no. 1, pp. 55-68.
[23] MAHMASSANI, H. (2001) Freight and Commercial Vehicle Applications. IN HENSHER, D. (Ed.) Travel Behaviour Research. Pergamon - Elsevier Science.
[24] MARSHALL, R. & MARX, L. (2002) Bidder Collusion. Working Paper, Penn State University - Duke University.
[25] MARTIN, S. (1993) Advanced Industrial Economics. Blackwell Publishers, Cambridge.
[26] MASKIN, E. & RILEY, J. (2000) Asymmetric Auctions. Review of Economic Studies, V67, pp 413-438.
[27] MCAFEE, R. & MCMILLAN, J. (1987) Auctions and Biddings. Journal of Economic Literature, V 25, pp 699-738.
[28] NANDIRAJU, S. & REGAN, A. (2005) Freight Transportation Electronic Marketplaces: A Survey of Market Clearing Mechanisms and Exploration of Important Research Issues. Proceedings 84th Annual Meeting of the Transportation Research Board, Washington D.C, January 2005.
[29] NARENDRA, K. & THATCHER, M. (1974) Learning Automata: a survey. IEEE Transactions on Systems, Man, and Cybernetics, N4, pp. 889-899.

[30] PAARSH, H. (1991) Deciding between Common and Private Value Paradigms in Empirical Models of Auctions. Journal of Econometrics, N15, pp 191-215.
[31] POWELL, W., MARAR, A., GELFAND, J. & BOWERS, S. (2002) Implementing Real-Time Optimization Models: A Case Application form the Motor Carrier Industry. Operations Research, Vol50, N4, pp. 571-581.
[32] RAIFFA, H., RICHARSON, J. & METCALFE, D. (2002) Negotiation Analysis. The Belknap Press of Hardvard University Press.
[33] RUBINSTEIN, A. (1998) Modeling Bounded Rationality. MIT Press, Cambridge, USA.
[34] SIMON, H. (1955) A behavioral model of rational choice. Quarterly Journal of Economics, V.69, pp. 99-118.
[35] SIMON, H. (1956) Rational choice and the structure of the environment. Psychological Review, V63, pp.129-138.
[36] STAHL, D. & WILSON, P. (1995) On Players Models of Other Players: theory and experimental evidence. Games and Economic Behavior, 10, pp 213-254.
[37] TCA (2003) Truckload Carrier Association Website - http://www.truckload.org/infocenter/TCAdocs/info_08_02_02.htm - Accessed December 16th, 2003.
[38] TIROLE, J. (1989) The Theory of Industrial Organization, The MIT Press.
[39] VIDAL, J. & DURFEE, E. (1995) Recursive agent modeling using limited rationality. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95). Menlo Park, California, AAAI Press, pp. 376-383.
[40] WALLISER, B. (1998) A spectrum of equilibration processes in game theory. Journal of Evolutionary Economics, V.8, N.1, pp. 67-87.
[41] WOLFSTETTER, E. (1999) Topics in microeconomics: industrial organization, auctions, and incentives. Cambridge University Press, USA.
[42] YANG, J., JAILLET, P. & MAHMASSANI, H. (2004) Real-time multivehicle truckload pickup and delivery problems. Transportation Science, 38, 135-148.