Impacts of Auction Settings on the Performance of Truckload Transportation Marketplaces

Miguel Andres Figliozzi, Hani S. Mahmassani, and Patrick Jaillet

The performance of different sequential auction settings for the procurement of truckload services is compared. In this environment, demands arrive randomly over time and are described by pickup and delivery locations and hard time windows. When loads arrive, carriers compete for their transport. Different auction and information disclosure settings are studied. Learning methodologies are discussed and analyzed. Simulation results are presented.

Recent research in auction and marketplace design highlights the importance of auction rules on bidders and market performance (*1, 2*). The main focus of this research is to compare and evaluate the effect of distinct auction rules on the performance of a transportation marketplace. This investigation focuses on the dynamic procurement of truckload pickup and delivery services in a sequential auction transportation marketplace; this marketplace is known as the truckload procurement marketplace (TLPM).

This work was motivated by the growth of network business-to-business forecasts (*3*). This growth is partly supported by the increasing use of private exchanges, whereby a company or group of companies invites selected suppliers to interact in a real-time marketplace, compete, and provide the required services. In the logistics sector, shippers have also set up private exchanges, which they use for confidential communications with their carriers. For example, DuPont has a logistics web portal to manage all inbound and outbound freight movements across all transportation modes (*4*). These exchanges enable freight visibility and also consolidation and optimization opportunities (*5*). On the supply side, carriers have begun to offer more Internet-based services, particularly the larger motor carriers (*6*).

Carriers participating in a TLPM face complex, interrelated decision problems. Two distinct problems stand out: (*a*) profit maximization problem (choose best pricing or bidding policy) and (*b*) cost minimization problem (operate the fleet in the most efficient way). Sequential auctions are notoriously complex problems; further, no equilibrium solution exists if there are several auctions (three or more) and multiunit demand bidders (*7*). Therefore, in this work, carriers are assumed boundedly rational. In addition, because of the inherent complexity, TLPM carriers are assumed not to attempt to acquire or use any knowledge about competitors' explicit decision (bidding) processes. Carriers solely learn about the distribution of past market prices or the relationships between realized profits and bids. Previous work has already dealt with the importance of dynamic vehicle routing technology and cost estimation in a TLPM (*8*). This work focuses on the previously mentioned profit maximization problem. Different learning approaches are adapted and evaluated in a freight transportation context.

The goal is not to find the optimal rules or procedures that lead to the best possible bidding. Rather, the goal is to define and simulate plausible bounded-rational procedures and behaviors of carriers competing in a TLPM. Three different auction formats are compared using simulation experiments: second-price auctions, first-price auctions with minimum information disclosure, and first-price auctions with maximum information disclosure.

MARKET DESCRIPTION

A TLPM enables the sale of cargo capacity mainly on the basis of price, yet still satisfies customer level-of-service demands. The specific focus of the study is the reverse auction, by which shippers post loads and carriers compete for them (bidding). The auctions operate in real time, and transaction volumes and prices reflect the status of demand and supply. A framework to study transportation marketplaces is presented by Figliozzi et al. (*9*).

The market consists of shippers, which independently call for shipment procurement auctions; and carriers, which participate in the auctions (assuming 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. 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.

In the TLPM, there are n carriers competing, and a carrier is denoted by $i \in \Im$, where $\Im = \{1, 2, \ldots, n\}$ is the set of all carriers. Let the shipment, auction arrival, and announcement epochs be $\{t_1, t_2, \ldots, t_N\}$ such that $t_i < t_{i+1}$. Let $\{s_1, s_2, \ldots, s_N\}$ 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

M. A. Figliozzi, Faculty of Economics and Business, Institute of Transport Studies (C37), University of Sydney, Sydney, New South Wales 2006, Australia. H. S. Mahmassani, Department of Civil and Environmental Engineering, University of Maryland, Martin Hall, College Park, MD 20742. P. Jaillet, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139-4307.

Transportation Research Record: Journal of the Transportation Research Board, No. 1906, Transportation Research Board of the National Academies, Washington, D.C., 2005, pp. 89–96.

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 $\{t_1, t_2, \ldots, t_N\}$ follow some general arrival process. Further, arrival times and shipments are assumed to come from a probability space (Ω, F, P), with outcomes $\{\omega_1, \omega_2, \ldots, \omega_N\}$. 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_i, 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 $b_j^{\Im} = \{b_j^1, \ldots, b_j^n\}$ generates publicly observed information y_j. Under maximum information disclosure, all bids are revealed after the auction; this is $y_j = b_j^{\Im}$. Under minimum information disclosure, no bids are revealed after the auction; this is $y_j = \{\}$. Each carrier is solely informed about bidding outcome: successful or unsuccessful.

The fleet status of carrier i when shipment s_j arrives is denoted as z_j^i, which consists of 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 $c^i(s_j, z_j^i)$. 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 $I_j^{\Im} = \{I_j^1, \ldots, I_j^n\}$ and $\Sigma_{i\in\Im} I_j^i \leq 1$. Let π_j^i be the profit obtained by carrier i for shipment s_j, then $\pi_j^i = [b_j^i - c^i(s_j, \theta_j^i)]I_j^i$.

LEARNING IN A TLPM

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 like believing that all competitors are in a strategic equilibrium.

Walliser distinguishes four distinct dynamic processes to play games: in a decreasing order of cognitive capacities, they are eductive processes, epistematic learning (fictitious play, or FP), behavioral learning (reinforcement learning, or RL), and evolutionary processes (*10*). An eductive process requires knowledge about competitors' behavior (agents simulate competitors' behavior). Epistematic and behavioral learning are similar to FP and RL, respectively. 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 in which it has been randomly matched to another player).

This work studies the two intermediate types of learning. On the one hand, 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. Further, the notion that a company is born, dies, and reproduces with each auction does not fit well behaviorally in the defined TLPM. Ultimately, neither extreme approach is practically or theoretically compelling in the TLPM context. Carriers that survive competition in a competitive market like truckload procurement cannot be inefficient or unskilled. They are merely limited in the strategies that 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 abilities (which give rise to bounded rationality).

In practical and theoretical applications, the setting of 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 reasonably justified or explained. Thus, restricting research to the TLPM context has clear advantages.

Normal operating ratios in the trucking industry range from 0.90 to 0.95 (*11*). It is expected that operating ratios in a TLPM 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 ratio fluctuations in a competitive market are expected, in response to natural changes in demand and supply. However, these fluctuations should be in the range 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 (*12*).

LEARNING MECHANISMS

In RL, 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 previous 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, with $M = \{\text{mc}_0, \ldots, \text{mc}_K\}$, such that if $\text{mc}_k \in M$ and $\text{mc}_{k+1} \in M$, then $\text{mc}_k < \text{mc}_{k+1}$. Using multiplicative coefficients, the profit obtained for any shipment s_j, when using the multiplicative coefficient mc_k, equals

$$\pi_j^i(\text{mc}_k) = (\text{mc}_k c_j^i - c_j^i) I_j^i = c_j^i I_j^i (\text{mc}_k - 1) \quad (1)$$

$$\pi_j^i(\text{mc}_k) = (b_j^{(2)} - c_j^i) I_j^i \quad (2)$$

The first equation applies to first-price auctions, and 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 represented by the term $b_j^{(2)}$. A general introduction to auctions is found in the comprehensive book by Krishna (*7*).

Adapting the reinforcement model to TLPM bidding, the probability $\varphi_j^i(\text{mc}_k)$ of carrier i using a multiplicative coefficient mc_k in the auction for shipment s_j equals

$$\varphi_j^i(\text{mc}_k) = [1 - \lambda\pi_{j-1}^i(\text{mc}_k)]\varphi_{j-1}^i(\text{mc}_k) + I_{j-1}^i(\text{mc}_k)\lambda\pi_{j-1}^i(\text{mc}_k) \quad (3)$$

Narendra and Thathachar showed that a player's 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 (*13*). The convergence is obtained as the reinforcement parameter λ goes to 0. To use Equation 3, each bidder needs information only on what it has bid and the result of the auction. To use this model, the profits $\pi^i_{j-1}(\mathrm{mc}_k)$ must be normalized to lie between 0 and 1 so that they may be interpreted as probabilities. The indicator variable $I^i_j(\mathrm{mc}_k)$ equals 1 if carrier i used the multiplicative coefficient mc_k when bidding for shipment s_j; the indicator is equal to 0 otherwise. The parameter λ is called the RL 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 (*14*). Therefore, a mediocre action can be reinforced, while, at the same time, better actions are negatively reinforced. Further, in an auction context, there is no learning when the auction is lost, because $\pi^i_{j-1}(\mathrm{mc}_k) = 0\ \forall\ \mathrm{mc}_k \in \mathrm{M}$ if $I^i_{j-1} = 0$.

Borgers and Sarin propose a model that deals with the aforementioned problems (*15*). 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 as ρ^i_j and the effective profit is denoted as

$$\tilde{\pi}^i_{j-1}(\mathrm{mc}_k) = \pi^i_{j-1}(\mathrm{mc}_k) - \rho^i_j \tag{4}$$

then

$$\varphi^i_j(\mathrm{mc}_k) = \left[1 - \lambda\tilde{\pi}^i_{j-1}(\mathrm{mc}_k)\right]\varphi^i_{j-1}(\mathrm{mc}_k) + I^i_{j-1}(\mathrm{mc}_k)\lambda\tilde{\pi}^i_{j-1}(\mathrm{mc}_k) \tag{5}$$

When $\rho^i_j = 0$, Equation 5 provides the same probability as updating Equation 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 TLPM context, the cost of serving shipments may vary significantly. Further, 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 research proposes a modified version of the stimulus response model with RL that better adapts to TLPM bidding. Each multiplicative coefficient mc_k has an associated average profit value $\bar{\pi}^i_j(\mathrm{mc}_k)$, that equals

$$\bar{\pi}^i_j(\mathrm{mc}_k) = \frac{\sum_{t\in(1,\ldots,j)} \pi^i_t(s_t) I^i_t(\mathrm{mc}_k)}{\sum_{t\in(1,\ldots,j)} I^i_t(\mathrm{mc}_k)}$$

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

$$\bar{\rho}^i_j = \frac{\sum_{t\in(1,\ldots,j)} \pi^i_t(s_t) I^i_t}{j}$$

Therefore, the average effective profit is defined as $\bar{\pi}^i_{j-1}(\mathrm{mc}_k) = \bar{\pi}^i_{j-1}(\mathrm{mc}_k) - \bar{\rho}^i_j$. Probabilities are therefore updated using Equation 6.

$$\varphi^i_j(\mathrm{mc}_k) = \left[1 - \lambda\bar{\pi}^i_{j-1}(\mathrm{mc}_k)\right]\varphi^i_{j-1}(\mathrm{mc}_k) + I^i_{j-1}(\mathrm{mc}_k)\lambda\bar{\pi}^i_{j-1}(\mathrm{mc}_k) \tag{6}$$

With the latter formulation (Equation 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 average reinforcement learning (ARL).

Stimulus-response learning requires the least information and can be applied to first- and second-price auctions. The probabilities updating Equations 3 and 6 are the same for first- and second-price auctions. Therefore, the application of the RL model does not change with the auction format that is used in the TLPM. Using this learning method, a carrier does not need to model either the behavior or the actions of competitors. The learning method is essentially myopic because 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.

Because the method is myopic, for the first-price auction, the multiplicative coefficients must equal or be greater than 1 (i.e., $\mathrm{mc}_0 \geq 1$). A coefficient less than 1 generates only 0 or negative profits. In a second-price auction, the multiplicative coefficients can be less than 1 and still generate positive profits because the payment depends on the competitors' bids.

In both types of auctions, it is necessary to specify not only the set of multiplicative coefficients but also 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.

In RL, therefore, the agent does not consider strategic interaction. The agent is unable to model any other agent's play or behavior. This agent is informed only by past experience and is content with observing the sequence of own past actions and the corresponding payoffs. Using only this action-reward experience, successful strategies are reinforced and failed strategies are inhibited. Maximization does not occur but rather movement in a utility-increasing direction, by choosing a strategy or by switching to a strategy with a probability positively related to the utility index.

RL (and its variants) is a strategy designed to operate in an environment in which the player (carrier) is unable to see competitors' actions. Therefore, it is able to strongly reinforce (positively or negatively) only one action: the last action played. Unlike RL, FP requires the observation of competitors' actions. A good introduction to types of RL and FP can be found in the work of Fudenberg and Levine (*14*).

FP came about as an algorithm to look for Nash equilibrium in finite games of complete information (*16*). 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 opponents' bounded rationality from their actions, although the agent may eventually know that opponents learned 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 FP learning scheme assumes that the 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 (e.g., 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 numbers to approximate the probabilities of competitors' play. In either case, the carrier must determine an estimated stationary price function ξ (with carriers estimating a normal distribution using competitors' past bids). If a second-price auction format is used in the TLPM, the carrier bids using

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

If a first-price auction format is used in the TLPM, the carriers bid using

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

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

SIMULATION FRAMEWORK

In this study, truckload carriers compete over a square area, with each side's length equal to one 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 cover this area and also time. Origins and destinations of demands are uniformly distributed over the square area, so that the average loaded distance for a request is 0.52 unit of distance. All the arrivals are random; the arrival process follows a time Poisson process. The expected interarrival time is $E[T] = 1/(K\lambda)$, where λ is the demand request rate per vehicle and K is the total market fleet size.

The total market fleet size used is four vehicles (although similar trends were obtained with larger fleets—eight vehicles—as long as the same arrival rate/fleet size ratio is used). Roughly, the average service time for a shipment is 0.77 unit of time (approximately $\lambda = 1.3$). The service time is broken down into 0.52 unit of time corresponding to the average loaded distance, plus 0.25 unit of time that approximates the average empty distance (average empty distance varies with arrival rates and time windows). Different Poisson arrival rates/truck/per unit of time are simulated (ranging from 0.5 to 1.5). As a general guideline, these values correspond to situations in which the carriers are

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

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

$$\begin{aligned} \text{LDT} = {} & \text{arrival time} + 2 \times (\text{shipment loaded distance} + 0.25) \\ & + 2 \times \text{uniform}\ (0.0,\ 1.0) \end{aligned}$$

All the shipments have a reservation price that is uniformly distributed (1.42, 1.52). In all cases, reservation prices exceed the maximum marginal cost possible in the simulated area (≈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 demand. 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 (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 measure is shippers' consumer surplus, which is the accumulated difference between reservation prices and prices paid. The fourth measure is total wealth generated equal to the accumulated difference between reservation price (of served shipments) and empty distance traveled.

Carrier fleet assignment and cost estimation are based on the static optimization, which is based on an approach proposed by Yang et al. (*17*). This approach solves static snapshots of the dynamic vehicle routing problem with time windows using an exact mathematical programming formulation. As new load occurs, static snapshot problems are solved repeatedly, enabling a complete reassignment of trucks to loads at each arrival instance.

The second-price auction used in the TLPM operates as follows:

- Each carrier submits a single bid.
- The winner is the carrier with the lowest bid (which must be below the reservation price; otherwise, the auction is declared vacant).
- The item (shipment) is awarded to the winner.
- The winner is paid either the value of the second-lowest bid or the reservation price, whichever is the lowest.
- The other carriers (not winners) do not win, pay, or receive anything.
- The same procedure applies to first-price auctions, except that the winner is paid the value of the winning bid.

ANALYSIS OF EXPERIMENTAL RESULTS

Figure 1 illustrates the relative performance of ARL and 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 1 shows the relative performance of ARL and RL after 500 auctions. It is clear that ARL obtains higher profits as the arrival rate increases. RL has poorer performance because it cannot converge steadily to the optimal coefficient. RL speed can be quite slow in an auction setting like TLPM. The optimal bidding factor can be used, and there is still about 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 decrease considerably and hinder convergence to the optimal or even convergence at all. As discussed previously, this issue can be avoided using averages (ARL method). The carrier RL tends to price lower (it keeps probing low-bidding coefficients longer) and therefore serves a higher number of shipments.

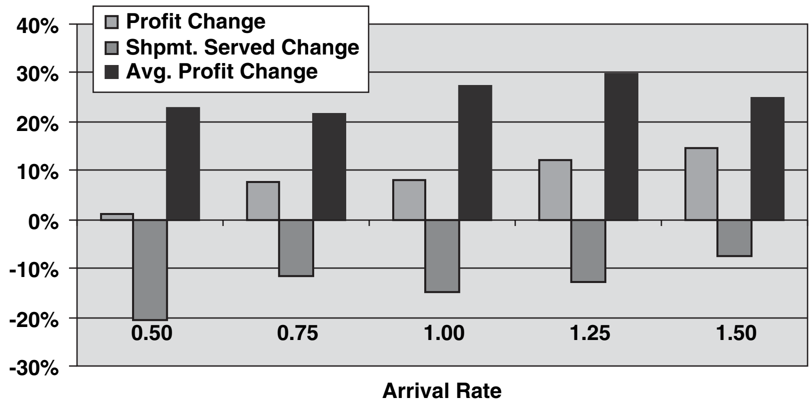

FIGURE 1 ARL versus RL (RL performance as a base).

Although convergence has not been proved, the simulation data suggest that the learning influence on profits levels out as the simulated number of auctions is increased. Figure 2 shows profit changes as the number of auctions is increased (ARL versus RL with an arrival rate of 1). In all cases, the same initial warm-up is used. To ease comparisons, the profit level after 500 auctions in Figure 1 (arrival rate = 1) corresponds to the profit level shown in Figure 2 for a simulation length of 500 auctions. In general, the length of the necessary warm-up will strongly depend on the learning parameters being used, initial conditions, and the simulated marketplace. In general, larger values of the learning factor λ will tend to need shorter warm-up periods (therefore shorter learning, because they tend to converge faster) but at the higher risk of adopting nonoptimal policies for longer periods of time.

The next experiment compares the performance of RL and FP in first-price auctions. RL uses more information than FP. Therefore, it is expected that a carrier using FP must outperform a carrier using RL. Figure 3 shows the relative performance of FP and ARL after 500 auctions. The ARL player has the same characteristics as the ARL player in Figure 1. The FP carrier divides possible competitors' bids into 15 intervals (from 0.0 to 1.5 in intervals of width 0.1) and starts with their uniform probability distribution.

Clearly, the FP carrier obtains higher profits across the board. Using 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 diminishes as the arrival rate increases. In the TLPM 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 use FP or a similar bidding strategy, particularly because FP (myopic) and ARL complexity differ little.

In second-price auctions, FP coincides with marginal cost bidding (regardless of the price distribution, the expected profit is always optimized with marginal cost bidding). RL and ARL do not perform better than FP in the simulated experiments.

The next experiment compares the performance of different sequential auction settings from carriers' and shippers' points of view. Four different measures are used to compare the auction environments: carrier profits, consumer surplus, number of shipments served, and total wealth generated. To facilitate comparisons among the four graphs shown in Figures 4 through 7, second-price auctions with marginal cost bidding are used as the standard to measure the two types of first-price auctions.

Figure 4 illustrates the profits obtained by carriers. FP carriers obtain higher profits than ARL carriers across the board. FP carriers use the obtained price information to their advantage. The highest carrier profit levels occur with second-price auctions. These results do not alter or contradict theoretical results. With asym-

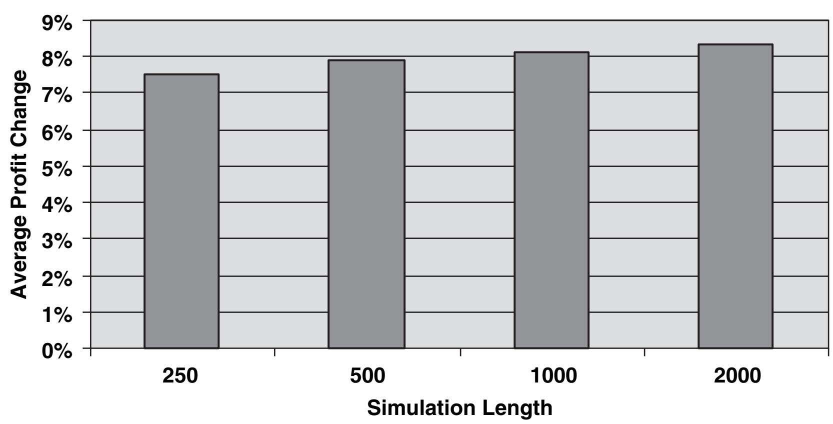

FIGURE 2 Effect of simulation length on profit change (ARL versus RL; arrival rate = 1).

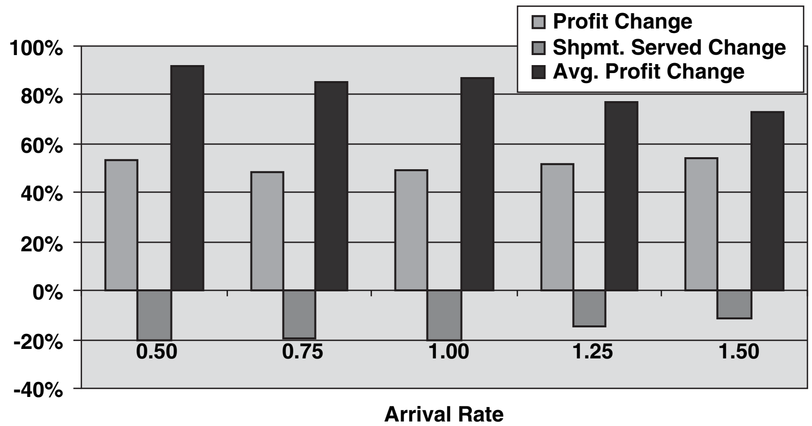

FIGURE 3 ARL versus FP (RL performance as base).

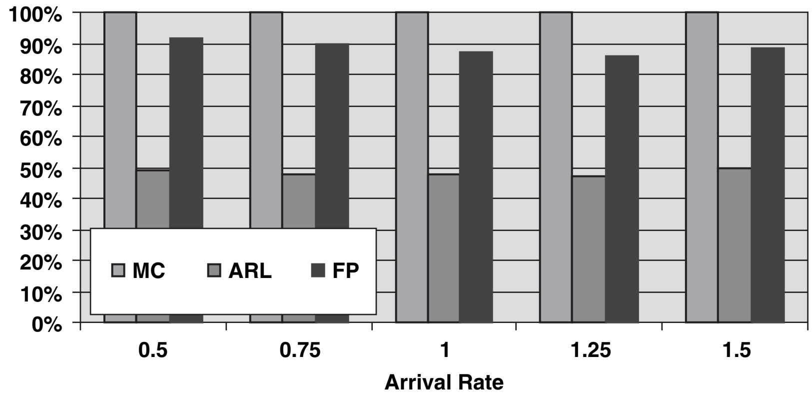

FIGURE 4 Carriers' profit level (second-price auction MC as base).

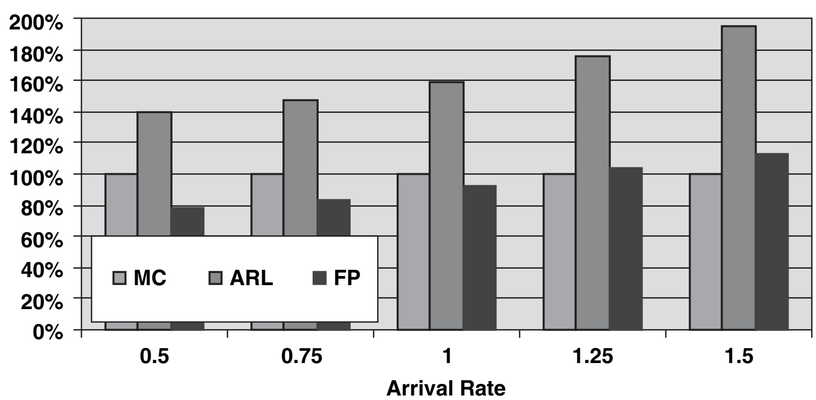

FIGURE 5 Consumer surplus level (second-price auction MC as base).

metric cost distribution functions, Maskin and Riley show that there is no revenue ordering between independent-value first-and-second price auctions (*18*).

Figure 5 illustrates the consumer surplus obtained with the three auction types. Clearly, first-price auctions with RL (minimum information disclosed) benefit shippers. Unsurprisingly, Figure 5 is almost the reverse image of Figure 4. Figure 6 shows the number of shipments served with each auction setting. As expected, second-price auctions serve more shipments. Even in asymmetric auctions, it is still a weakly dominant strategy for a bidder to bid its own value in a second-price auction (this property of one-item second-price auction is independent of the competitors' valuations). Therefore, in second-price auctions, the shipment goes to the carrier with the lowest cost.

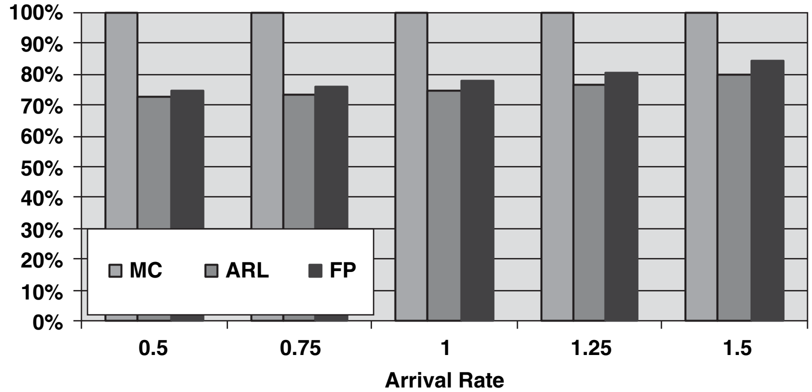

FIGURE 6 Number of shipments served (second-price auction MC as base).

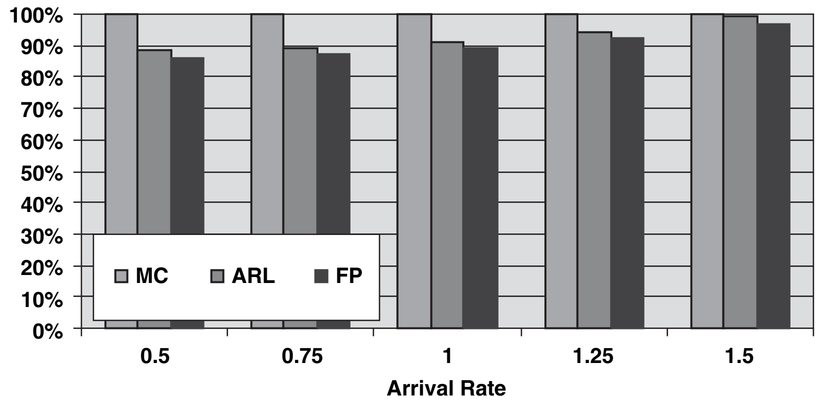

FIGURE 7 Total wealth generated (second-price auction MC as base).

In contrast, with ARL, there is a positive probability that inefficient assignments exist, because a higher-cost competitor can use a bidding coefficient that results in a lower bid. Similarly with FP carriers, if the price functions differ (which is likely because each carrier models the competitors' prices), a higher-cost carrier can underbid a lower-cost carrier with a positive probability.

Figure 7 shows the wealth generated with each auction setting. Second-price auctions generate more wealth across the board. Marginal cost bidding is the most price-efficient mechanism of the tested auction settings. As the arrival rate increases, the gap in total wealth generated tends to close (Figure 7). Consistent with previous results, the lowest wealth generated corresponds to the case with FP bidders.

In summary, under the current TLPM setting, carriers, shippers, and a social planner would each select a different auction setting. Carriers would like to choose a second-price auction. If a first-price auction were used, carriers would like to have maximum information disclosure. More information enables players to maximize profits, although 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 fewer shipments served. Finally, from a societal viewpoint, the most efficient system is the second-price auction. More shipments are served and more wealth is generated.

CONCLUSIONS

A sequential auction framework was used to compare distinct sequential auction settings. RL and FP, two learning mechanisms adequate for TLPM settings, are introduced and analyzed.

Computational experiments indicate that auction setting and information disclosure affect TLPM performance. Maximum information disclosure enables carriers to maximize profits at the expense of shippers' consumer surplus; minimum information disclosure enables shippers to maximize consumer surplus but at the expense of fewer shipments served. Marginal bidding in second-price auctions generates more wealth and more shipments served than first-price auctions. The results illustrate that critical arrival rates provide no incentive to use bidding factors (no deviations from static marginal cost bidding).

ACKNOWLEDGMENT

This study was supported by a grant from the National Science Foundation.

REFERENCES

1. Klemperer, P. What Really Matters in Auction Design. *Journal of Economic Perspectives,* Vol. 16, No. 1, 2002.

2. Roth, A., and A. Ockenfels. Last-Minute Bidding and the Rules of Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet. *American Economic Review,* Vol. 92, No. 4, 2002.
3. Mullaney, T., H. Green, M. Arndt, R. Hof, and L. Himelstein. The E-BIZ Surprise. *Business Week,* May 2003.
4. Gilligan E. DuPont Takes a Global View. *Journal of Commerce,* Feb. 16–22, 2004. www.joc.com. Accessed June 2004.
5. Cooke, J. For Your Eyes Only. *Logistics Management,* July 1, 2002. www.manufacturing.net/lm/index.asp?layout=articlePrint&articleID=CA232464. Accessed June 2004.
6. Ellinger, A., D. Lynch, J. Andzulis, and R. Smith. B-to-B E-Commerce: A Content Analytical Assessment of Motor Carrier Websites. *Journal of Business Logistics,* Vol. 24, No. 1, 2003, pp. 199–220.
7. Krishna, V. *Auction Theory.* Academic Press, San Diego, Calif., 2002.
8. Figliozzi, M. A., H. S. Mahmassani, and P. Jaillet. Competitive Performance Assessment of Dynamic Vehicle Routing Technologies Using Sequential Auctions, In *Transportation Research Record: Journal of the Transportation Research Board, No. 1882,* Transportation Research Board of the National Academies, Washington, D.C., 2004, pp. 10–18.
9. Figliozzi, M. A., H. S. Mahmassani, and P. Jaillet. Framework for Study of Carrier Strategies in an Auction-Based Transportation Marketplace. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1854,* Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 162–170.
10. Walliser, B. Spectrum of Equilibration Processes in Game Theory. *Journal of Evolutionary Economics,* Vol. 8, No.1, 1998, pp. 67–87.
11. Truckload Carrier Association. 2002. www.truckload.org/infocenter/TCAdocs/info_08_02_02.htm. Accessed Dec. 16, 2004.
12. Paarsh, H. Deciding Between Common and Private Value Paradigms in Empirical Models of Auctions. *Journal of Econometrics,* No. 15, 1991, pp. 191–215.
13. Narendra, K. S., and M. A. L. Thathachar. Learning Automata: A Survey. *IEEE Transactions on Systems, Man, and Cybernetics,* No. 4, 1974, pp. 889–899.
14. Fudenberg, D., and D. Levine. *Theory of Learning in Games.* MIT Press, Cambridge Mass., 1998.
15. Borgers, T., and T. Sarin. Naïve Reinforcement Learning with Endogenous Aspirations (Mimeo). University College of London, 1996.
16. Brown, G. Iterative Solutions of Games by Fictitious Play. In *Activity Analysis of Production and Allocation* (T. Koopmans, ed.), Wiley, New York, 1951.
17. Yang, J., P. Jaillet, and H. S. Mahmassani. On-Line Algorithms for Truck Fleet Assignment and Scheduling Under Real-Time Information. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1667,* TRB, National Research Council, Washington, D.C., 1999, pp. 107–113.
18. Maskin, E., and J. Riley. Asymmetric Auctions. *Review of Economic Studies,* Vol. 67, 2000, pp. 413–438.

The Freight Transportation Planning and Logistics Committee sponsored publication of this paper.