

Dynamic threshold policy for delaying and breaking commitments in transportation auctions

Martijn Mes*, Matthieu van der Heijden, Peter Schuur

May 15, 2007

Abstract

In this paper we consider a transportation procurement auction consisting of shippers and carriers. Shippers offer time sensitive pickup and delivery jobs and carriers bid on these jobs. We focus on revenue maximizing strategies for shippers in sequential auctions. For this purpose we propose two strategies, namely delaying and breaking commitments. The idea of delaying commitments is that a shipper will not agree with the best bid whenever it is above a certain reserve price. The idea of breaking commitments is that the shipper allows the carriers to break commitments against certain penalties. The benefits of both strategies are evaluated with simulation. In addition we provide insight in the distribution of the lowest bid, which is estimated by the shippers.

1 Introduction

The procurement of transportation is an important task for shippers (typically large manufacturers and retailers) because it greatly effects their costs and service levels. In practice, a procurement process includes carrier (trucking company) screening, carrier assignment, load tendering and performance review. During the last few years, this procurement has moved from telephone to web based services (Song and Regan 2001). In this paper we consider an automated transportation procurement auction where shippers offer time sensitive pickup and delivery jobs and carriers bid on these jobs. Each auction is initiated by a shipper and ends with commitment between the shipper and a carrier.

*Corresponding author. Address: Department of Operational Methods for Production and Logistics, School of Management and Governance, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands; phone +31-53-489-4062; fax: +31-53-489-2159; e-mail: m.r.k.mes@utwente.nl.

Auctions are often considered as appropriate means for dynamic job allocation in distributed environments. Desirable properties include limited required information exchange and, in some cases, Pareto efficient outcomes (McAfee and Mc Millan 1987; Wellman and Walsh 2001). However, when multiple jobs are auctioned at different points in time (sequential auctions), such an allocation would be less appropriate if we do not take into account the future consequences of an allocation. Also when jobs are complementary (e.g. transportation jobs that can be served sequentially by the same vehicle) or substitutable (e.g. transportation jobs that are available at the same time), a certain allocation may become unfavorable when new orders appear. To overcome this, several authors (e.g. Caplice and Sheffi (2003)) suggested the use of *combinatorial auctions* for transportation procurement in which a carrier can bid on multiple jobs. However, combinatorial auctions involve many inherently difficult problems. As mentioned by Song and Regan (2005), we face the bid construction problem, where bidders have to compute bids over different job combinations, and the winner determination problem, where jobs have to be allocated among a group of bidders. In addition (1) it may be unrealistic to bundle jobs which belong to different shippers, and (2) these procedures are not directly applicable in situations where jobs arrive at different points in time.

To improve the allocation of jobs, we take the sequential transportation procurement auction as given, and focus on strategies for the participants. In a previous paper (Mes et al. 2006) we focused on profit maximizing strategies for the carriers. We proposed a bid pricing strategy where the arrivals of future jobs are taken into account through the use of opportunity costs. In this paper we focus on strategies for the shipper. We propose two options, namely delaying and breaking commitments.

The idea of delaying commitments is that shippers postpone commitments for which they expect to make a better commitment in the future. So, if a shipper has plenty of time to auction a certain job, he will not agree with a relatively high bid. When the time for dispatch becomes nearer, the price he is willing to accept will raise. We denote this mechanism by *dynamic threshold policy*. The idea of breaking commitments is that the shipper allows a carrier to decommit from an agreement against a certain penalty. These penalties are chosen such, that whenever a carrier decommits a job, the shipper expects to find a new carrier for a lower price. We denote this mechanism by *decommitment policy*.

There is some strategic equivalence between the dynamic threshold policy and the decommitment policy. Suppose a carrier wins a certain job. After two days the carrier decommits from the

contract and pays a certain penalty. One can imagine that this penalty equals the expected costs for the shipper for auctioning this job two days later, which in turn has some strong connections with the dynamic threshold policy. We develop a dynamic programming algorithm that can be used for both, the dynamic threshold and the decommitment policy.

The goal of this paper is threefold. First to derive a dynamic threshold policy and decommitment policy. Second to evaluate their benefits (will they increase the overall system performance). Third to study their relation, that is, are they complementary or substitutable?

2 Literature

Our problem is related to several research areas, such as operations management (transportation), economic theory (optimal auctions), and mathematical theory (optimal stopping). In the next sections we describe the relation of this paper with these research areas and describe our contribution.

2.1 Transportation

The dynamic allocation of transportation jobs belongs to the large class of dynamic fleet management problems. A few representative examples of this stream include (Carvalho and Powell 2000; Godfrey and Powell 2002; Yang et al. 2004). This research mostly focus on real-time vehicle routing strategies for a single carrier. Here we focuses on the shipper, and more specifically on shippers that procure transportation services using auctions. A transportation procurement auction consists of three steps: (1) bid preparation by the shipper (who may bid on what), (2) bid pricing by the carriers (3) bid analysis by the shipper. The majority of research on transportation procurement auctions focuses on the bid preparation step (Caplice and Sheffi 2003). Others focus on bid pricing, scheduling, and routing decisions of the carriers, see (Song and Regan ; Figliozzi et al. 2003). Less attention has been paid to the shipper.

Traditionally, a shipper allocates transportation jobs to carriers one-by-one, that is through sequential auctions. Such a system ignores the interdependencies between subsequent jobs. A significant portion of the trucking industry costs is due to the repositioning of empty vehicles from the destination of one load to the origin of a subsequent load (Song and Regan 2002). Interdependencies occur because serving one job is greatly affected by the opportunity to serve another job. To cope with these dependencies, Caplice and Sheffi (2003) suggested to use combinatorial auctions.

As demonstrated by Ledyard et al. (2002), the benefits of combinatorial auctions to shippers can be significant. A survey on combinatorial auctions for the procurement of transportation services can be found in (Sheffi 2004). For reasons as mentioned in the introduction, we choose here for sequential one-shot auction procedures. In order to cope with the interdependencies among jobs we apply the tricks of delaying and breaking commitments.

2.2 Optimal auctions

The design of auction mechanisms that maximize the seller's expected revenue, called optimal auctions, received a great deal of attention. For an extensive literature survey on this topic we refer to (McAfee and Mc Millan 1987). Part of this work focuses on reserve prices in sequential auctions. As shown by Myerson (1981), the reserve price increases the expected revenue of the seller by preventing the object from being sold at a low price. Closely related is the work of (McAfee and Vincent 1997) who study the optimal reserve-price path in a sequence of first- and second- price auctions. In particular, the auctioneer puts the same object for sale repeatedly, until it is sold. At each round he chooses a reserve price according to his (increasingly pessimistic) beliefs about the buyers' valuations.

A crucial assumption in the optimal auction literature is that each bidder's valuation is known to be drawn from a common distribution (Bose et al. 2006). We do not make any assumptions on the valuation functions of the individual bidders, but instead estimate the distribution of the lowest bids. In addition we incorporate correlations in bids, together with a regression on time and several job characteristics. Another difference is that we consider reverse (procurement) auctions.

Another line of research within the economic auction literature focuses on decommitment. In automated negotiation systems, contracts have traditionally been binding. Such contracts do not allow agents to efficiently deal with future events when contracts might become unfavorable. To overcome this, a leveled commitment protocol is introduced in (Sandholm and Lesser 2001). Here an agent can decommit (for whatever reason) simply by paying a decommitment fee to the other agent. It is shown through game-theoretic analysis that this leveled commitment feature increases the Pareto efficiency of contracts and can make contracts more beneficial for both parties. The efficiency of such protocols depends heavily on how the penalties are decided. Therefore (Sandholm et al. 1999) developed algorithms for optimizing contracts in terms of prices and penalties. Penalty functions for sequences of multiple leveled commitment contracts are studied in (Andersson and Sandholm 2001). They conclude that penalties as a percentage of the contract price, increasing

in decommitment time, perform best. Another example can be found in (Hoen and Poutré 2004) who apply the decommitment concept to a multi-agent transportation setting. Their setting differs from ours because they only allow exchange of jobs between vehicles of the same carrier, so carriers are not allowed to decommit from a contract with a shipper. In this paper, we provide a formal expression for time-dependent decommitment penalties by drawing a parallel with reserve prices.

2.3 Optimal stopping

The choice of whether or not to accept the lowest bid in a sequential auction is related to the so-called optimal stopping problems. The most famous example of an optimal stopping problem is the secretary problem. In the classical secretary problem, a decision-maker has to hire one applicant out of a pool of n applicants who appear sequentially. The decision-maker must decide immediately upon seeing an applicant whether to hire him or not. For a historical overview of the classical secretary problem we refer to (Ferguson 1989). The classical secretary problem is called a no-information problem in which the distribution of offers is unknown. The stopping decision is only based on the relative ranks of the observations and not on their actual values. Which is obviously different from our case because shippers are able to learn pricing behavior.

A number of authors have established the existence and properties of optimal stopping policies when the offer distribution is known in advance, and the offer rate is periodic. For example (Karlin 1962), who studied the problem of selling an asset. Our approach differs from this line of research in the sense that we consider (1) historic auction information to update the offer distribution, (2) time-dependent offers, (3) correlation between subsequent offers, and (4) the finite horizon problem as a special case. For more information on optimal stopping problems we refer to (Chow et al. 1971).

2.4 Contributions

To summarize the previous sections, our contribution consists of the following:

1. We develop profit maximizing strategies for shippers in transportation procurement auctions. To avoid combinatorial complexities we propose a dynamic threshold policy that enables the shippers to strategically delay commitments or set decommitment penalties for the carriers.
2. We determine optimal reserve price paths for reverse auctions using probability distribution functions for the lowest bid based on historical data. These distributions do not depend on

the iid assumption of bids or distributions of individual bidders.

3. We provide a formal expression for time-dependent decommitment penalties by drawing a parallel with the reserve price paths.
4. We aim at a wider applicability than competitive procurement auctions, in particular for closed environments, i.e., allocation to a closed group of trusted carriers, or for auction procedures that are commonly used in multi-agent systems for resource allocation. Therefore we provide a performance evaluation, using simulation, not only in terms of individual benefits for the shippers, but also in terms of total system performance.

3 Model

We consider a transportation market consisting of shippers and carriers. Shippers offer jobs by starting a transportation procurement auction and carriers bid on these jobs. One job involves the transport of a unit load (full truckload). We define a job by an announcement time a , an origin i , a destination j , and as a soft restriction a latest pickup time l of the load at the origin. We introduce a soft time-window length $\sigma = l - a$ within which transportation should be started. We introduce d as the distance between the origin and destination of a job. Note that for notational convenience we omit job indices.

Objective of a shipper is to minimize the price paid to a carrier for transporting a certain load. Shippers consider the prices as random variables with probability distributions which can be estimated based on historic data. These distributions are characterized by a time-dependent mean and standard deviation, and by correlations in prices between subsequent auction rounds.

In a transportation procurement auction, shippers typically put out a request for quotes from a set of carriers (Song and Regan 2002). This process is similar to a simple sealed-bid auction in which each bidder submits a sealed bid for a single item. Without loss of generality, we choose here for a reverse first-price sealed-bid auction in which the lowest bidder receives his bid amount, given the shipper does not reject all bids.

We implement the market mechanism as follows. When an job arrives at some shipper he starts an auction by sending an announcement to all carriers. In return, each carrier responds with a bid. Without a dynamic threshold policy, the shipper sends a grant message to the carrier with the lowest bid while the others receive a reject message. When using a dynamic threshold policy,

a shipper might expect to receive a better bid in the future. After all, prices fluctuate over time due to changes in the available transportation capacity and in the transportation schedules. So if the best bid is relatively high (which can be learned from history) it might be better to wait for more attractive prices. In this case, the shipper only selects a winner whenever the lowest bid is below a certain threshold level. Otherwise the auction stays open (Section 4.1) or the shipper will start a new auction some period later (Section 4.2). As the deadline for dispatch comes nearer (or is already reached), the shipper increases his threshold to get transportation. Basic assumption here is that we consider only one job at a time. So threshold prices of jobs are independent of the current set of open jobs.

In case of decommitment a vehicle is allowed to decommit from a job by paying some penalty to the shipper (Section 5). In this case the shipper immediately starts a new auction for this job.

4 Dynamic threshold

First we develop a continuous time threshold function (Section 4.1). This function can be used in a continuous auction where a shipper waits until a bid drops below the threshold price (bidders' take-it-or-leave-it prices). Next we develop a basic policy for repeated auctions in discrete time (Section 4.2). In this structure, the shipper will start a new auction a fixed auction period later when the best bid is above his threshold price. This structure can be considered as an approximation for the continuous case if the auction period is small. In Section 4.3 we incorporate possible correlation in bid prices between subsequent auction rounds.

4.1 Basic threshold policy for continuous auctions

To illustrate the theoretical benefits of a dynamic threshold policy we consider a continuous time model under some simplifying assumptions.

Whenever a shipper becomes aware of a new job, he will start an auction for this job. All vehicles immediately bid on this job, but may update their bid at any time. The shipper, of course, is only interested in the lowest bid. We assume independent and identically distributed lowest bids which can be described by a continuous distribution function $F(b)$. The time between subsequent updates of the lowest bid is exponentially distributed with rate λ . At each update of the lowest bid, the shipper has to decide whether or not to accept the current lowest bid. We only consider the period before the latest pickup time. If the job is not sold before this time, the

shipper will face costs $Z(\tau) = \beta + c\tau$, where β is a constant, c the penalty costs per unit time, and τ the tardiness with respect to the latest pickup time (see Appendix for a formal derivation of this cost function).

We indicate the time until the latest pickup time by a time-to-go t . We introduce the value function $V(t)$ as the minimum expected price a shipper has to pay eventually, given a time-to-go t . For $t < 0$, the value $V(t)$ is given by the value $Z(-t)$ of auctioning the job after its latest pickup time. The recursive relationship which characterizes the finite horizon minimum expected price, as a function of the time-to-go t , is given by:

$$V(t) = E[\min(B, V(t - Y))] \quad (1)$$

with B the stochastic variable for the lowest bid and Y the exponentially distributed time until the next bid update. Karlin (1962) showed that there exist an optimal policy that accepts the first bid B whose value satisfied $B > E[V(t - Y)]$. So, we only accept a bid whenever it is lower than the minimum expected price at the expected time of the next update of the lowest bid. We rewrite this policy by using a threshold function $\alpha(t)$ which is given by:

$$\begin{aligned} \alpha(t) &= E[V(t - Y)] \\ &= E[\min(B, \alpha(t - Y))] \end{aligned} \quad (2)$$

Integration over all bid prices B and time between bid updates Y gives the following:

$$\alpha(t) = \int_0^t \left(\int_0^{\alpha(t-y)} b dF(b) + \alpha(t-y) \int_{\alpha(t-y)}^{\infty} dF(b) \right) \lambda e^{-\lambda y} dy + \int_t^{\infty} Z(y-t) \lambda e^{-\lambda y} dy \quad (3)$$

Following an approach similar to (Karlin 1962) we rewrite this function in the form of a differential equation. Partial integration of the first integral yields:

$$\alpha(t) = \int_0^t \left(\alpha(t-\xi) - \int_0^{\alpha(t-\xi)} F(b) db \right) \lambda e^{-\lambda \xi} d\xi + \left(\beta + \frac{c}{\lambda} \right) e^{-\lambda t} \quad (4)$$

Next we replace ξ by $t - x$, and multiply both sides by $e^{\lambda t}$:

$$\alpha(t) e^{\lambda t} = \lambda \int_0^t \left(\alpha(x) - \int_0^{\alpha(x)} F(b) db \right) e^{\lambda x} dx + \left(\beta + \frac{c}{\lambda} \right) \quad (5)$$

Differentiation by t , together with the boundary condition $\alpha(0)$, yields:

$$\begin{aligned}\alpha'(t) &= -\lambda \int_0^{\alpha(t)} F(b) db, \text{ with } t \geq 0 \\ \alpha(0) &= \beta + \frac{c}{\lambda}\end{aligned}\tag{6}$$

We can solve this differential equation numerically for different distributions $F(b)$ of the lowest bid. Since a uniform distribution allows an analytical derivation, let us assume that $F(b) = \frac{b}{\omega}$, for $0 \leq b \leq \omega$. We further take $\omega = \beta + \frac{c}{\lambda}$, so at the latest pickup time the shipper expects to pay the highest possible bid. Then we have the following threshold function:

$$\alpha(t) = \frac{2\omega}{\lambda t + 2}\tag{7}$$

If we get only one chance to auction a job, the expected price equals $\omega/2$. The relative savings $s(t)$ depending on the time-to-go t , by using the threshold policy are therefore given by:

$$s(t) = \left(1 - \frac{4}{\lambda t + 2}\right) \cdot 100\%\tag{8}$$

To illustrate the savings take $\lambda = 1$, then for $t = 2, 8, 18, 38$ the savings are respectively given by 0%, 60%, 80%, 90%.

Although these results are promising, they only hold under the assumption of independent and identically distributed lowest bids. In many cases, this assumption is not realistic. In the next two subsections we extend our results to time-dependent bid prices, and dependency among bids.

4.2 Basic threshold policy for repeated auctions

In order to incorporate a more realistic pricing behavior let us discretize time. The resulting dynamic threshold function may then be used as an approximation for the continuous case. However discretization may also be part of the market structure because basically the continuous auction is replaced by a repeated auction. In a repeated auction, we have a series of auctions used to sell the same object in subsequent periods.

We assume that the time between subsequent auction rounds is R . In contrast with the previous section, we express the timing of an auction in terms of auction round numbers instead of the time-to-go. Numbering starts at the first auction round at the announcement time of a job.

In each auction round, the shipper has to decide whether or not to accept the lowest bid.

Before rejecting all bids, the shipper has to calculate the probability of receiving a better bid in the future. In order to do so, we assume that shippers are able to estimate the distribution $F_n(b)$ of the lowest bid b as a function of the auction round n and some job characteristics (see Section 6). For notational convenience we omit the job characteristics from the distribution function.

We only consider auction rounds before the latest pickup time. Therefore the maximum number of auction rounds is given by $N = \lfloor (l - a) / R \rfloor + 1$. Again we use a cost function $Z(\tau)$ (see Appendix) for the expected price after the latest pickup time. We introduce the shorthand notation Z to denote the expected price of the first auction round after the latest pickup time.

Depending on the auction period R there is a probability that the lowest bid remains the same in the next auction round. Therefore, we introduce a probability q that the lowest bid is updated between two successive auction rounds. So, for a Poisson updating process with rate λ (see previous section) this probability is given by $q = 1 - e^{-\lambda R}$. We assume that the shippers are able to estimate this change probability, independent of the underlying updating process.

We introduce the value function $V_n(b_n)$ as the expected price a shipper has to pay eventually (in this or one of the remaining auction rounds) given a lowest bid b_n in the current auction round n . To get an optimal strategy, we deduce the optimum decision numbers by working backward from the last possible auction round. We get the following:

$$\begin{aligned} V_N(b_N) &= \min \{b_N, Z\} \\ V_n(b_n) &= \min \{b_n, E[V_{n+1}(B_{n+1}|B_n = b_n)]\} \end{aligned} \tag{9}$$

In the last auction round we have to choose between the current bid b_N and the expected price Z for auctioning the job after the latest pickup time. In all other rounds we have to choose between accepting the current bid b_n or reject it and expect a price $E[V_{n+1}(B_{n+1}|B_n = b_n)]$ later on. The price we expect to accept later on depends on the current lowest bid, because there is a probability of no change. Therefore we have to take into account the current lowest bid in the threshold prices. The threshold price in auction round n equals the expected price we accept in auction round $n + 1$ or later, given current lowest bid b_n :

$$\begin{aligned} \alpha_n(b_n) &= E[V_{n+1}(B_{n+1}|B_n = b_n)] \\ &= (1 - q) \min \{b_n, \alpha_{n+1}(b_n)\} + q \int_0^\infty \min \{b, \alpha_{n+1}(b)\} dF_{n+1}(b) \end{aligned} \tag{10}$$

We solve this equation by backwards dynamic programming. To be able to do so, we have to

include the bids in the state space. Therefore we discretize the bids, add the current lowest bid to the state space, and iterate on the auction round n :

$$\alpha_n(b_n) = (1 - q) \min \{b_n, \alpha_{n+1}(b_n)\} + q \sum_{b=0}^L P_{n+1}(b) \min \{b, \alpha_{n+1}(b)\} \quad (11)$$

where $P_n(b)$ is a discretization of $F_n(b)$, and L is chosen large enough.

To illustrate the behavior of the threshold prices, we calculate Equation 11 for some parameter settings. We use time-to-go $\sigma = 10$, a Poisson updating rate $\lambda = 1$, and we take $Z = \infty$, so in the last auction round we always accept the lowest bid. For the probability $P_n(b)$ of the lowest bid, we discretize a normal distribution with bins of width 1. We calculate the expected price $E[V_n(b)] = E[\min \{b, \alpha_n(b)\}]$ as a function of the auction round n .

First we use an auction period $R = 1$ and evaluate three scenarios for the mean and standard deviation of the lowest bid:

- Constant mean of 100 and standard deviation of 50, independent on the auction round n
- Increasing mean of 5 per auction round, starting at 50 at the first round, and constant standard deviation of 50
- Constant mean of 100, and increasing standard deviation of 3 per auction round, starting at 20 at the first round




Figure 1: Varying time between orders




Figure 2: Varying amplitude in time between orders

In Figure 1 we see that prices increase with increasing auction rounds. Starting with a low mean will result in lower threshold prices, which then increase relatively faster. Starting with

low standard deviation results in higher threshold prices because we are less able to profit from the variation in bid prices. But because variances increase, this threshold function also increases relatively slower. In the last auction round we always accept the lowest bid. Therefore in the last auction round, all three prices equal 100.

Next, we evaluate the threshold prices for varying lengths of the auction period R . For the distribution of the lowest bid we use a constant mean and standard deviation of respectively 100 and 50. Going from $R = 1$ to $R = 0.2$ will result in 5 times more auction rounds. Going from $R = 0.2$ to $R = 0.04$ again results in 5 times more auction rounds. From Figure 2 we conclude that the added value of extra auction rounds clearly decreases.

4.3 Incorporating correlation

Bid prices of carriers fluctuate due to new job arrivals. In the previous section we developed a basic dynamic threshold policy that takes advantage of these fluctuations. We have seen that, for given job characteristics, these prices depend on the auction round and the time between subsequent auction rounds. Another important aspect, which we did not consider so far, is correlation between subsequent lowest bid updates. Suppose the current lowest bid is relatively high, indicating that all carriers are busy at the moment. When the shipper re-auctions the job a short time period later, and he receives a new lowest bid, it is likely that this bid is also relatively high. In this section we deal with these correlations between subsequent auction rounds.

Because a part of the correlation is caused by the time dependency of bids, we consider correlation between the deviations from expected lowest bids. We only consider correlations between two successive auction rounds (no time-lags) and assume that the correlation between two successive auction rounds n and $n + 1$ is independent of the number n (so correlation between round 1 and 2 is similar to the correlation between rounds 5 and 6). We further assume that the deviations can be described by a linear trend with coefficient ϕ . The lowest bid in auction round $n + 1$ is then given by:

$$b_{n+1} = E[b_{n+1}] + \phi (b_n - E[b_n]) + \epsilon_{n+1} \quad (12)$$

where ϵ_{n+1} is an error term.

The derivation of ϕ , and the way we incorporate this correlation in the threshold function of Equation 11, is presented in Section 6 (Equation 19).

As mentioned in the beginning of Section 4.2, the threshold policy for repeated auctions can be

used as an approximation for continuous auctions. Therefore we use linear extrapolation between the threshold price in the last auction round before t , and the first auction round after t . If there are no remaining auction rounds after t , we use the threshold value Z for a time-to-go zero. The continuous threshold function is given by:

$$\alpha_t(b) = \begin{cases} \alpha_n(b) + \left(\frac{\sigma-t}{R} - \lfloor \frac{\sigma-t}{R} \rfloor\right) (\alpha_{n+1}(b) - \alpha_n(b)) & \text{for } t \geq \sigma - (N-1)R \\ Z - \frac{t}{\sigma - (N-1)R} (Z - \alpha_N) & \text{for } 0 \leq t < \sigma - (N-1)R \end{cases} \quad (13)$$

where α_n refers to the original threshold price in auction round n with n given by $n = \lfloor \frac{\sigma-t}{R} \rfloor + 1$. The auction period R should be small enough, such that the probability of more than one lowest bid update within this period is 'low'. For $t < 0$, that is auctioning the job after its latest pickup time, $\alpha_t(b)$ is given by $Z(-t)$.

5 Decommitment

A second trick that can be used by shippers is to allow carriers to break commitments against certain penalties. The decommitment penalties are set by the shipper and are publicly available to the carriers at all times. These penalties should cover the extra costs for a shipper to find a new carrier. A shipper will face extra costs because the time until latest pickup time is shorter now. These extra costs are given by the difference in threshold prices as presented in Section 4.

The decommitment penalty $D_{s,t}$ for a job committed at time s and decommitted at time t is given by:

$$D_{s,t} = E[\alpha_t(b)] - E[\alpha_s(b)] \quad (14)$$

A carrier decommits from a job whenever his expected revenue of inserting a new job and removing the decommitted job, is higher than the current expected revenue plus the decommitment penalty. Note that in case of decommitment, a carrier will not receive his bid price for the decommitted job.

With Equation 14 we derived a formal expression for the levelled commitment penalties, only by using the threshold prices. The applicability however is quite different because now the decision has to be made by the carriers.

6 Parameter estimation

In the previous sections we assumed that the lowest bid B can be described by a probability distribution $F_t(b)$ or $F_n(b)$. In this section we estimate the distribution $F_t(b)$. The distribution $F_n(b)$ can easily be derived from this by using $t = \sigma - (n - 1)R$.

The lowest bid depends mainly on two factors: time-to-go t , distance d . The distance d will have a direct effect on the transportation costs of the vehicles. The time-to-go t is the difference between the current time θ and the latest pickup time l . Note that other order characteristics can easily be added in the statistical model described in this section. For example the route ij (with i the origin and j the destination) of an order when popularity of regions differs.

Although we use a single distribution $F_t(b)$ to describe the lowest bid, bids consist of different cost components. Here we distinguish between transportation costs and penalty costs. These cost factors differ in their dependence on the factors time-to-go, distance and route, and in the correlation between subsequent auction rounds. We decided to perform a regression on both cost factors and combine them into a single distribution for the lowest bid. After all, a shipper should have knowledge on both cost factors independent on whether or not they are included in the bid prices.

We introduce w_t and p_t for respectively the transportation costs and the penalty costs of the lowest bid, given time-to-go t . We use the following multiple linear regression functions:

$$w_t(d) = \alpha_w + \beta_w t + \gamma_w d \quad (15)$$

$$p_t(d) = \alpha_p + \beta_p t + \gamma_p d \quad (16)$$

For the penalties we cannot directly apply linear regression because we have censored observations. To be precise, we have left-censored data because penalties are never negative. If a carrier is able to pickup the job before the latest pickup time, the penalties are zero, independent on the time between pickup and the latest pickup time. In order to perform linear regression with left-censored data we use Tobit regression, see (Tobin 1958).

We also have to deal with heteroscedasticity because the variance in bid prices decrease with increasing time-to-go. Also longer jobs have higher variances. To incorporate heteroscedasticity, we proceed as follows. First we ignore heteroscedasticity and perform the regression mentioned above. Next, we divide the time-to-go t and distance d in discrete blocks, and calculate the residual

variance in each of these blocks. From this we derive a linear trend for the standard deviation:

$$\sigma_{td} = \alpha_\sigma - \beta_\sigma t + \gamma_\sigma d \quad (17)$$

We incorporate this trend in the regression model using Weighted Least Squares. Here we give points with lower variance a greater statistical weight. To be more precise, we multiply each residual with a weight equal to the inverse of the variance σ_{td}^2 . For information on heteroscedasticity in Tobit models, which we use to describe the penalty costs, we refer to (Green 1997).

As mentioned before, we decided to use a single distribution function for the lowest bid in order. This is a major approximation because both cost components behave differently. For example, once there are penalty costs in the lowest bid, it is very unlikely that penalties are lower in the next auction round. It is also not possible to simply add up the variances in transportation- and penalty costs, because penalties are censored at zero. When adding up the variances, it might be the case that a price of zero (or even below zero) has a high probability of occurrence. As an approximation, we use a continuous distribution function (see Section 8) for $F_t(b)$ with a mean given by $\mu_t(d) = w_t(d) + \max(0, p_t(d))$ and standard deviation given by Equation 17. In the remainder we use the shorthand notation μ_t and μ_n to denote the mean price given respectively a time-to-go t or auction round n .

In order to incorporate the correlation in price deviation between subsequent auction rounds, we model the trend in price deviations by an AR(1)-process. Here the expected deviation in auction round n can be expressed as a linear function of the deviation in auction round $n - 1$. Therefore, the shipper stores the deviation $\delta_n = b_n - \mu_n$ for each auction round $n \geq 2$. Then we have:

$$\delta_n = \phi \delta_{n-1} + \epsilon_n \quad (18)$$

where ϕ is a scalar (see Equation 12), and ϵ_n the error term. Because we treat each pair δ_n, δ_{n-1} equally for all n , we also assume that the variance in price deviation remains the same for all auction rounds n . Then the coefficient ϕ equals the correlation coefficient:

$$\phi = \frac{Cov(\delta_n, \delta_{n-1})}{Var(\delta_n)} \quad (19)$$

In order to incorporate the impact of a price deviation δ_n in auction round n on the lowest bid in auction round $n + 1$, we simply increase the mean price μ_{n+1} in auction round $n + 1$ with

$\phi\delta_n$. So the error term ϵ_n from Equation 12 is no longer required. Using Equation 11, we derive the following function for the threshold prices with correlated bids:

$$\alpha_n(b_n) = (1 - q) \min \{b_n, \alpha_{n+1}(b_n)\} + q \sum_{b=0}^L P_{n+1}(b) \min \{b + \phi(b_n - \mu_n), \alpha_{n+1}(b + \phi(b_n - \mu_n))\} \quad (20)$$

This threshold function, together with the decommitment penalty function (Equation 14), are examined in the next sections.

7 Experimental settings

We simulate a transportation procurement market where shippers offer transportation jobs to a set of carriers. Selection of a carrier for a certain job is only based on price and delivery time. Transportation takes place in a square area of 100x100 km and all vehicles have a speed of 50 km/hour.

Because our focus is on profit maximizing strategies for a shipper, let us use a simple pricing and scheduling strategy for the carriers. In the remainder we speak in terms of individual vehicles and ignore the carriers. Each vehicle maintains a list of jobs. These jobs are carried out as soon as possible, and a job in process cannot be interrupted. Upon announcement of a new job, the vehicle evaluates the insertion of this job in his job lists, that is, without altering the relative ordering of jobs already in the list. The bid price for this job is given by the marginal costs of the cheapest insertion. These costs are 1 per hour travel time and 10 per hour tardiness. If a vehicle wins an auction, it uses the cheapest insertion position for the new job.

In case of decommitment, a vehicle also evaluates the impact of inserting the new job while decommitting a job already in his job list. To avoid combinatorial difficulties we only consider decommitment of a single job. For all jobs that allow decommitment, the vehicle temporarily removes the job and evaluates the insertion of the new job in his reduced job list. Again, the bid price of a vehicle is given by the cheapest insertion including possible decommitment penalties.

The simulation study consists of three parts. In the first part we examine the distribution of the lowest bid. In the second part we evaluate how well our statistical model of Section 6 can be used to describe the lowest bid. In the last part we simulate a transportation market where a single shipper uses the dynamic threshold policy and the decommitment policy.

7.1 Price distribution

Here we distinguish between learning jobs and regular jobs. We use the learning jobs to gain insight in the distribution of the lowest bid for various types of job and times to auction these jobs. To avoid influencing the market prices we do not award these jobs to vehicles. Regular jobs are auctioned once and are always awarded to a vehicle. All jobs have a time-window σ of 10 hours.

Regular jobs appear according to a Poisson process with arrival rate of 6 per hour. Origin and destination coordinates are chosen randomly from the square region. These jobs are auctioned as soon as possible, that is 10 hours before latest pickup time, and are always awarded to the vehicle with the lowest bid.

The arrival process of learning jobs is the same as for regular jobs. For these jobs we consider fixed length which we auction at several time instances. The length of these jobs is chosen randomly between 40, 60, 80, 100 and 120 km. Each job is centered around the middle of the square and lies on a diagonal. The direction and diagonal are also chosen randomly. The first auction moment of each order is 10 hours before the latest pickup time. The auction period between successive auction rounds is 2 hours and the latest auction round is 2 hours after the latest pickup time. So, we auction each job 8 times. After the latest auction round, we remove the job without rewarding it to a vehicle.

We use the lowest bid data for auctioning the learning jobs of different length in several auction rounds to evaluate its behavior and distribution characteristics.

7.2 Price evolution

Again we use learning jobs and regular jobs, and all jobs have a time-window σ of 10 hours. For the arrival rate, which is the same for both job types, we consider 5 per hour (quiet), 5.5 per hour (normal), and 6 jobs per hour (busy). For both job types, the origin and destination coordinates are chosen randomly from the square area. Again the regular jobs are auctioned as soon as possible, that is, 10 hours before latest pickup time, and are always awarded to a vehicle in this auction round. The learning jobs are auctioned in two rounds. The first auction round starts at a random time between zero and ten hours before the latest pickup time. The second round starts a random time between zero and 1 hour later. After the second round, the job is removed. We use the lowest bid data for the learning jobs to calculate the regression functions from Section 6.

7.3 Transportation market

Here we evaluate the impact of delaying and breaking commitments. We apply these policies both separately and in combination. When using only the decommitment policy, the threshold $\alpha_t(b)$ in the decommitment penalty function (see Equation 14) simply equals the expected price at time t . This because when a job is decommitted, the shipper immediately starts a new auction and always accepts the lowest bid in this auction. The threshold $\alpha_t(b)$ at a time t just before the latest auction round equals the expected price in this auction round.

In order to calculate the threshold functions, we use the regression functions found in the previous experiments. Therefore, part 2 can be seen as the learning phase for part 3. However, delaying and breaking commitments will have an effect on the market prices. To be precise, the allocation decisions for one job will have effect on the future bids for other jobs. To avoid focusing on learning issues, we decided to distinguish between one shipper and an external market. We evaluate the performance of the dynamic threshold policy and the decommitment policy, by enabling the individual shipper to use these strategies and compare it with the performance of the external market. The external market uses a standard policy where jobs are auctioned as soon as possible and may not be decommitted.

To describe the relative size of the individual shipper compared to the external market we introduce a market share. The market share of the individual company describes which portion of the incoming orders belongs to him. We consider a market share of 1% and 10%. For the arrival rate of all orders we use the same values as in part 2.

8 Simulation

Here we present the numerical results corresponding with the three experiments described in the previous section.

8.1 Price distribution

To gain insight in the behavior of the lowest bid, we consider distribution parameters such as the mean, standard deviation and skewness of the lowest bid. In this section we are mainly interested in the skewness. The mean and standard deviation of the lowest bid as a function of the time-to-go and job length will be investigated in the next section.

To describe the skewness of the lowest bid, we perform a simulation experiment consisting of

10 replications with different seeds. In each replication we generate 100,000 orders, including a warm-up period of 100 orders. The number of replications corresponds with a confidence level of 95% with a maximum relative error of 5% for the skewness. The results can be found in Figure 3 for the transportation costs and in Figure 4 for the penalty costs.




Figure 3: Skewness in transportation costs




Figure 4: Skewness in penalty costs

From these figures we see that penalties always have a positive skew and skewness decreases with decreasing time-to-go. In case of a large time-to-go, the transportation costs for short jobs have negative skewness while long jobs have positive skewness. After the latest pickup time all jobs have negative skew.

These results can be explained using the extreme value theory (EVT). This theory describes the order statistics of a large set of random observations from the same (arbitrary) distribution. It was proven by Fisher and Tippett (1928) that the limiting distribution for the maximum is either the Gumbel, the Fréchet, or the Weibull distribution. The Fréchet distribution is used whenever the parent distribution has finite upper limit and is fat-tailed, for example Student-t, Cauchy or Pareto. The Weibull distribution is used whenever the parent distribution is bounded, for example uniform, Power law, Beta. The Gumbel distribution is used whenever the tails of the parent distribution are decreasing exponentially, for example Exponential, Weibull, Gamma, Normal.

In our case we are interested in the distribution of the lowest bid out of n bids (where n is 10, which is of course not a large set). Extreme value theory states that for $n \rightarrow \infty$, the distribution of the lowest bid is (1) Weibull whenever the parent distribution is bounded from below, and (2) Gumbel whenever the parent distribution declines exponentially. From our experiments we see

that small jobs have a higher probability of an insertion. Therefore prices of these jobs have a very long left tail, going to zero if the job can be nicely inserted. The limiting distribution for the lowest bid on these jobs is therefore the Gumbel distribution with a negative skew. Most of the time, long jobs have to be added to the end of a schedule. The price for appending a job consists of the costs for driving empty toward the origin, and the costs for the loaded move. These prices are therefore bounded from below at a value equal to the costs of the loaded move. The limiting distribution for the lowest bid on these jobs is a Weibull distribution with a positive skew.

Although extreme value theory can be used to explain the skewness of our observations, we can not use the extreme value distributions because we are not dealing with large samples (we only have 10 bidders). After fitting several distribution functions to the lowest bid data, it appears that only small jobs (40km and 60km) can be fitted nicely. Suitable distributions are the Logistic- and Normal distribution, and for a long time-to-go the Weibull distribution. Possible candidate distributions for the longer jobs are Gamma, Beta, Lognormal and the Gumbel distribution for the maximum. However none of these distributions provide a good fit because the data appears to have two peaks.

The peaks in the distribution of the lowest bid are caused by the shape of the transportation area. Therefore consider the bid price of a single vehicle for adding a new job at the end of his job list. This bid price depends on the loaded travel distance for the new order, but also on the costs of driving empty from the destination of the last order in his current job list, to the origin of the new order. One can image that the likelihood of a certain empty travel distance r can be described by the circumference of a circle with radius r around the origin of the new job. To be precise, on the part of the circumference that falls within the square area. Now consider the job of length 80km (on the diagonal, centered around the mean). After using standard trigonometry we derive the total length of the intersection of the circle with the square area, as a function of the radius r , see Figure 5.

When the distribution of individual bids exhibits two peaks, also the distribution of the minimum (given the limited number of bidders) will have two peaks. The only difference is that the peaks are moved to the left (corresponding with lower costs) and the first peak will be relatively higher.

Although we here present the results assuming an insertion scheduling method, similar results hold for (1) a scheduling method where every new job is appended to the end of a schedule and (2) a scheduling method which allows complete rescheduling. The results in case of append scheduling




Figure 5: Length of intersection of a circle with a square as a function of the radius

are similar to those of insertion scheduling for long jobs because long jobs have a low probability to be inserted. The results in case of complete rescheduling are similar to those of insertion scheduling with short jobs because there is more scheduling flexibility and the bounds are therefore less tight.

To summarize this section, it is impossible to come up with a single distribution function that can be used to describe the lowest bid for all types of jobs. For example, the skewness of the distribution ranges from high negative values to high positive values. To overcome this, one could use separate distribution functions, depending on the job length and time-to-go. However, in order to keep things simple, we decided to use the normal distribution in the remainder of this section.

8.2 Price evolution

For all three network settings (quiet, normal, busy) we perform regression on the transportation costs and penalty costs of the lowest bid.

First we perform regression on the transportation costs. We consider three functions: (1) regression $w'_t(d)$ on the mean transportation costs assuming constant variance, (2) regression σ_{td} on the standard deviation in transportation costs, and (3) the regression $w_t(d)$ on the mean transportation costs considering the heteroscedasticity described by σ_{td} . In order to perform regression on σ_{td} , we divide both, the time-to-go t and distance d , in 10 bins. We then use weighted least squares, with weights equal to the inverse of the variances. The results can be found in Table 1.

Clearly, prices increase with decreasing time-to-go or increasing distance. The reason for this is that with decreasing time-to-go, there is less flexibility to schedule a job, so the probability

Network	Regression	R^2
Quiet	$w'_t(d) = 1742.40 - 138.68t + 85.33d$	0.73
	$\sigma_{td} = 1266.53 - 68.34t + 8.24d$	0.71
	$w_t(d) = 1630.94 - 123.55t + 85.98d$	0.75
Normal	$w'_t(d) = 1874.48 - 164.89t + 87.32d$	0.71
	$\sigma_{td} = 1369.07 - 84.76t + 9.42d$	0.88
	$w_t(d) = 1745.44 - 145.15t + 87.83d$	0.75
Busy	$w'_t(d) = 2061.15 - 200.68t + 90.49d$	0.70
	$\sigma_{td} = 1590.51 - 105.92t + 9.07d$	0.84
	$w_t(d) = 1921.61 - 175.22t + 90.57d$	0.74

Table 1: Regression on the transportation costs

of a cheap insertion decreases. The same holds for longer transportation jobs. We also see that variances increase with decreasing time-to-go or increasing distance, just like the expected prices. The differences between $w'_t(d)$ and $w_t(d)$ are small, but taking into account heteroscedasticity certainly leads to a better fit.

Next we perform regression on the penalty costs. We consider two functions: (1) regression $p'_t(d)$ on the mean penalty costs, ignoring the fact that we are dealing with censored observations, and (2) the regression $p_t(d)$ on the mean penalty costs using Tobit regression. The results can be found in Table 2.

Network	Linear regression	Tobit regression
Quiet	$p'_t(d) = 4679.95 - 735.17t + 6.42d$	$p_t(d) = 9518.49 - 6380.91t + 16.85d$
Normal	$p'_t(d) = 5386.53 - 857.49t + 9.20d$	$p_t(d) = 9583.33 - 6059.55t + 23.16d$
Busy	$p'_t(d) = 7231.12 - 1191.72t + 20.64d$	$p_t(d) = 6830.08 - 6473.49t + 92.37d$

Table 2: Regression on the penalty costs

Remarkable difference is that price depends more heavily on the time-to-go t . The reason for this is that, by using Tobit regression, penalties of zero are treated as uncertain values less than or equal to zero (because values are censored at zero).

Finally, we consider the correlation coefficient between subsequent auction rounds and the change probabilities $q(r)$ depending on the time r between successive auction rounds. To calculate the correlation coefficient we only consider subsequent auction rounds in which the lowest bid is updated. To determine the change probabilities we divide all data in 20 bins, based on the time between subsequent auction rounds. We then fit an exponential distribution to derive the change probability $q(r) = 1 - e^{-\lambda r}$ depending on the time r between subsequent auction rounds. The results can be found in Table 3.

The probability of an update of the lowest bid increases with increasing number of orders. The

Network	Correlation coefficient	Rate change probability
Quiet	0.50	0.911
Normal	0.48	0.934
Busy	0.46	0.982

Table 3: Correlation and change probability

correlation in deviation from expected bid prices decreases with increasing number of orders.

8.3 Transportation market

In this experiment we investigate the performance of the dynamic threshold policy (DT) and the decommitment policy (DC). We use two performance indicators. First, the relative total costs compared to the situation where no shipper uses the described policies. Second, the relative costs per order for the individual shipper compared to the costs per order for the external market. For the costs we only consider the costs for driving empty and the penalty costs. The loaded travel costs are omitted because they do not depend on the policy used.

Each experiment consists of 10 replications with different seeds. In each replication we generate 50,000, including 100 orders as a warm-up period. The number of replications corresponds with a confidence level of 95% with a maximum relative error of 5% for both performance indicators.

First we consider a market share of 1%. From the results of Table 4 we draw the following conclusions. First, the total costs are always lower when using one of the two strategies. Especially in busy networks and when we use both strategies in combination. Second, the relative costs advantage of the individual shipper is positive when using the dynamic threshold policy, especially in quiet networks. However when using the decommitment policy, the individual shipper has higher costs per order than the external market. There are two reasons for this. First reason is that we are only working with cost prices. For the shipper this means that he determines the decommitment penalties such that he expects to play even after decommitment. However, it appears that the expected costs after decommitment are higher than expected. This because the expected costs are based on the learning period where we auction each job as soon as possible. When allowing delaying or breaking commitments, some jobs are auctioned later. This results in less scheduling flexibility, and therefore higher costs, of the vehicles. Especially in busy networks this will be the case because vehicles have longer schedules. A second reason for the relative difference in case of decommitment is that the external market will benefit from the decommitment option for jobs of the individual shipper. A vehicle only decommits from a job whenever this results in savings for inserting the new job. With probability of 99%, this new job comes from the external market

which obviously benefits from the decommitment option for jobs of the individual shipper.

Network	Indicator	DT	DC	DT & DC
Quiet	Total costs	99.6	99.8	99.2
	Difference	73.4	119.8	88.1
Normal	Total costs	99.6	99.7	99.1
	Difference	76.7	124.2	89.2
Busy	Total costs	99.3	97.2	97.0
	Difference	80.0	202.5	107.1

Table 4: Regression on the transportation costs

Next we consider a market share of 10% (Table 5). Major difference with the previous results are (1) the total costs are always lower and (2) the costs advantage for the individual shipper when using the dynamic threshold policy is lower. The latter is caused by the fact that the difference in expected prices (based on the learning phase) and the actual market prices, increases with increasing market share. Despite the estimation error, the total costs are reduced. This provides an indication that with increasing market share, so towards the application of an internal allocation mechanism, the total costs will be reduced even further. Of course we then have to focus on updating procedures of the data from Section 6.

Network	Indicator	DT	DC	DT & DC
Quiet	Total costs	98.5	98.9	97.8
	Difference	76.7	119.9	88.0
Normal	Total costs	98.2	98.7	97.7
	Difference	79.1	125.4	95.9
Busy	Total costs	97.3	97.6	93.6
	Difference	82.5	169.1	126.8

Table 5: Regression on the transportation costs

9 Conclusions

In this paper, we presented a dynamic threshold policy. This policy enables shippers to postpone commitments for which they expect to make a better commitment in the future. In addition, we described a decommitment policy which allows carriers to decommit from an agreement with a shipper against a certain penalty.

Applications of these policies include both competitive transportation procurement auctions and closed environments. For the procurement auction we are mainly interested in the benefits of these policies for an individual shipper within a larger market. For the closed environments (group

of trusted carriers, internal allocation mechanism) we investigate the impact of these policies on the total costs.

From our simulation experiments, we conclude that both policies reduce the total costs. Also the costs per order for an individual shipper using the dynamic threshold policy are significantly lower than those who did not use such a threshold policy. The decommitment policy, as we derived it without profit margins, will result in relatively higher costs per order for the individual shipper, compared to the external market.

Further research is required on learning, profits, and the combination with 'smart' carriers. With learning we refer to updating procedures of the estimated data. By using the described policies, the market prices will change and possibly other players will react on that. This especially becomes an issue in networks that change over time, closed networks, and networks with multiple players with different strategies. A second aspect of further research is concerned with profits and auction types. For clarity of presentation we ignore profits in this paper and use a simple first-price auction. For practical applications we should also take into account profits, for example by using a second-price auction.

The last aspect of further research is concerned with the combination of profit maximizing strategies for both the shippers and carriers. In a previous paper (Mes et al. 2006) we focused on profit maximizing strategies for the carriers. We proposed a bid pricing strategy where the arrival of future jobs are taken into account through the use of opportunity costs. We evaluated this opportunity valuation method for the carriers by assuming a naïve strategy for the shipper. In this paper we evaluated profit maximizing strategies for the shipper by assuming naïve pricing- and scheduling strategies for the carriers. If both parties (carriers and shippers) are using intelligent strategies, it might be the case that they compete against each other and no party will be better off. It might even be the case that this has a negative effect on the total system performance. Therefore we aim to enhance the strategies for both players by incorporating the opponents' behavior. In addition we will perform an extensive simulation experiment to study the interrelation of the different policies of both players.

References

- Andersson, M. and T. Sandholm (2001). Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics and Control* 25, 615–640.

- Bose, S., E. Ozdenoren, and A. Pape (2006). Optimal auctions with ambiguity. *Theoretical Economics* 1, 411–438.
- Caplice, C. and Y. Sheffi (2003). Optimization-based procurement for transportation services. *Journal of Business Logistics* 24(2), 109–128.
- Carvalho, T.A. and W.B. Powell (2000). A multiplier adjustment method for dynamic resource allocation problems. *Transportation Science* 34, 150–164.
- Chow, Y.S., H. Robbins, and D. Siegmund (1971). *Great Expectations: The Theory of Optimal Stopping*. Boston: Houghton Mifflin Company.
- Ferguson, T.S. (1989). Who solved the secretary problem? *Statistical Science* 4, 282–296.
- Figliozzi, M.A., H.S. Mahmassani, and P. Jaillet (2003). Framework for study of carrier strategies in auction-based transportation marketplace. *Transportation Research Record* 1854, 162–170.
- Fisher, R.A. and L.H.C. Tippett (1928). Limiting forms of the frequency distribution of the largest and smallest member of a sample. *Proceedings of the Cambridge Philosophical Society* 24, 180–190.
- Godfrey, G.A. and W.B. Powell (2002). An adaptive dynamic programming algorithm for dynamic fleet management, i: Single period travel times. *Transportation Science* 36(1), 21–39.
- Green, W.H. (1997). *Econometric Analysis* (3rd edition ed.). Upper Saddle River, New Jersey: Prentice Hall.
- Hoen, P.J. 't and J.A. la Poutré (2004). A decommitment strategy in a competitive multi-agent transportation setting. In D. P. P. Faratin and J. Rodriguez-Aguilar (Eds.), *Agent Mediated Electronic Commerce V (AMEC-V)*, pp. 56–72. Berlin: Springer - Verlag.
- Karlin, S. (1962). Stochastic models and optimal policy for selling an asset. In S. K. K. J. Arrow and H. Scarf (Eds.), *Studies in Applied Probability and Management Science*, pp. 148–158. Stanford University Press.
- Ledyard, J.O., M. Olson, D. Porter, J.A. Swanson, and D.P. Torma (2002). The first use of a combined value auction for transportation services. *Interfaces* 32(5), 4–12.
- McAfee, R.P. and J. Mc Millan (1987). Auctions and bidding. *Journal of Economic Literature* 25(2), 699–738.

- McAfee, R.P. and D. Vincent (1997). Sequentially optimal auctions. *Games and Economic Behavior* 18, 246–276.
- Mes, M.R.K., M.C. van der Heijden, and P.C. Schuur (2006). Opportunity costs calculation in agent-based vehicle routing and scheduling. *Beta Working Paper Series, WP-168*.
- Myerson, R.B. (1981). Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73.
- Sandholm, T. and V. Lesser (2001). Leveled commitment contracts and strategic breach. *Games and Economic Behavior* 35, 212–270.
- Sandholm, T.W., S. Sikka, and S. Norden (1999). Algorithms for optimizing leveled commitment contracts. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 535–540.
- Sheffi, Y. (2004). Combinatorial auctions in the procurement of transportation services. *Interfaces* 34(4), 245–252.
- Song, J. and A.C. Regan. An auction based collaborative carrier network. *Transportation Research, Part E under review*.
- Song, J. and A.C. Regan (2001). Transition or transformation? emerging freight transportation intermediaries. *Transportation Research Record* 1763, 1–5.
- Song, J. and A. Regan (2002). Combinatorial auctions for transportation service procurement: the carrier perspective. *Transportation Research Record. Journal of the Transportation Research Board* 1833, 40–46.
- Song, J. and A. Regan (2005). Approximation algorithms for the bid construction problem in combinatorial auctions for the procurement of freight transportation contracts. *Transportation Research Part B* 39, 914–933.
- Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica* 26, 24–36.
- Wellman, M. and W. Walsh (2001). Auction protocols for decentralized scheduling. *Games and Economic Behavior* 35, 271–303.
- Yang, J., P. Jaillet, and H.S. Mahmassani (2004). Real-time multivehicle truckload pickup and delivery problems. *Transportation Science* 38, 135–148.

10 Appendix

Here we derive the expected costs $Z(\tau)$ to auction the job a time τ after its latest pickup time. After the latest pickup time, this job can only be scheduled with penalties. These penalties are c per unit time. We model this by saying that if we do not accept the current lowest bid, and the time until the next bid update is Y , then the shipper faces extra penalties cY , which he has to pay immediately. Of course in our 'real' application, these penalties will be incorporated in future bid prices. Next we introduce a value function $V(\tau)$ which reflects the minimum expected price the shipper has to pay eventually, given a time τ after the latest pickup time. This price excludes penalties paid so far. So, we write:

$$Z(\tau) = V(\tau) + c\tau \quad (21)$$

The minimum expected price a shipper has to pay, as a function of the time τ is given by:

$$V(\tau) = \min(B, E[V(\tau + Y) + cY]) \quad (22)$$

with B the stochastic variable for the lowest bid and Y the exponentially distributed time until the next bid update. The threshold function in time τ equals the expected price at the next update of the lowest bid:

$$\begin{aligned} \beta(\tau) &= E[V(\tau + Y) + cY] \\ &= E[\min(B, \beta(\tau + Y)) + cY] \end{aligned} \quad (23)$$

Now we only accept an offer b at time τ if it is lower than $\beta(\tau)$. So, we get the following:

$$\beta(\tau) = \int_0^\infty \left(\int_0^{\beta(\tau+y)} b dF(b) + \beta(\tau + y) \int_{\beta(\tau+y)}^\infty dF(b) + cy \right) \lambda e^{-\lambda y} dy \quad (24)$$

Basically $\beta(\tau + y)$ does not differ from $\beta(\tau)$ because we consider an infinite horizon problem. Note that if you receive a bid update a time Y after the previous bid update, and you decide not to accept it, this bid 'disappears' and you already lost an amount cY . So, it is just like starting

the problem over again, that is, the problem is invariant of time. So we write:

$$\beta = \int_0^\infty \left(\int_0^\beta b dF(b) + \beta \int_\beta^\infty dF(b) + cy \right) \lambda e^{-\lambda y} dy \quad (25)$$

$$= \int_0^\beta b dF(b) + \beta \int_\beta^\infty dF(b) + \frac{c}{\lambda} \quad (26)$$

From this we get:

$$c = \lambda \int_0^\beta (\beta - b) dF(b) \quad (27)$$

Solving this equation for β yields a stationary threshold policy after the latest pickup time. We accept the first offer below β .

The expected costs $Z(\tau)$, given we receive a bid update τ time units after the latest pickup time, are given by:

$$Z(\tau) = \beta + c\tau \quad (28)$$

As an example, suppose $F(b)$ is $U[0, \omega]$, the uniform distribution on the interval $(0, \omega)$. Then we get the following:

$$\begin{aligned} c &= \lambda \int_0^\beta \frac{(\beta - b)}{\omega} db, \text{ if } \beta \leq \omega \\ c &= \lambda \int_0^\omega \frac{(\beta - b)}{\omega} db, \text{ if } \beta > \omega \end{aligned} \quad (29)$$

Equating to c , we find:

$$\begin{aligned} \beta &= \sqrt{\frac{2\omega c}{\lambda}}, \text{ if } c \leq \frac{\omega\lambda}{2} \\ \beta &= \frac{c}{\lambda} + \frac{\omega}{2}, \text{ if } c > \frac{\omega\lambda}{2} \end{aligned} \quad (30)$$

So, for penalty factors $c > \frac{\omega\lambda}{2}$, we have a threshold price larger than ω . This means that we accept the first bid after the latest pickup time, because it is always lower than ω . The expected price will then be $\omega/2$ and the expected penalties c/λ because we expect to receive this bid update a time $1/\lambda$ after the latest pickup time.

For the discrete case, we replace the expected time $1/\lambda$ between successive bid updates by the auction period R .