Risk-Averse Auction Planning and its Integration into Supply Chain Management Systems

Yaxin Liu
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
yxliu@cc.gatech.edu

Richard Goodwin
IBM T.J. Watson Research Center
P.O. Box 218, Route 134
Yorktown Heights, NY 10598
rgoodwin@us.ibm.com

Sven Koenig
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
skoenig@cc.gatech.edu

Abstract

Auctions are an important means for purchasing materials. Consequently, a supply chain management system has to be able to decide whether to participate in auctions and how much to bid. To address these issues, we generalize results from auction theory in several ways. First, auction theory often assumes that decision makers want to maximize their expected profits. However, decision makers are often risk-averse when faced with the possibility of losing the auction and then incurring huge penalties for not being able to satisfy some orders. We take their risk attitudes into account by changing the objective from maximizing the expected profit to maximizing the expected utility of the profit according to a risk-averse utility function. Second, auction theory often assumes that decision makers know their valuations for the auctioned item. However, their valuations depend on how they can use the item in the production process. We therefore integrate auction strategies into a production planning system to derive the valuation automatically. Third, auction theory often assumes that the probability distribution over the competitors' valuations is known. We use simulations of the production planning system with integrated auction strategies to approximate these probability distributions automatically. The combination of these three research contributions results in a prototype of a supply chain management system with integrated auction strategies for the paper industry.

Introduction

Auctions are an important part of e-commerce, and a large number of websites are dedicated to them. For example, more than thirty websites are devoted to business-to-business auctions for the paper industry alone (including Colmart, PaperExchange, and PaperLoop), and it is expected that the use of auctions will continue to increase in the future. Auctions are, for example, run to sell paper pulp and paper rolls. Consequently, a supply chain management system for the paper industry has to be able to decide whether to participate in auctions to purchase material and how to bid. To build the auction-planning component of such a system, one can make use of existing research results on auctions (Klemperer 2000). Unfortunately, research on auctions often studies auctions in isolation and consequently makes simplifying assumptions about the planning objective and knowledge of the decision maker. In practice, however, production decisions and results of previous auctions can affect the behavior of the decision maker in subsequent auctions. In this paper, we therefore generalize results from auction theory in three different ways, both for first- and second-price sealed-bid auctions in the symmetric independent private values model:

- First, auction theory often assumes that decision makers want to maximize their expected profits (Friedman 1956; Vickrey 1961). However, decision makers are often risk-averse when faced with the possibility of losing the auction and then incurring huge penalties for not being able to satisfy some orders. A decision maker is risk-averse if he or she is willing to accept a decrease in the expected profit to reduce the variance and thus the possibility of a large loss. We take the risk attitudes into account by changing the objective function from maximizing the expected profit to maximizing the expected utility of the profit for a risk-averse utility function.
- Second, auction theory often assumes that decision makers know their valuations for the auctioned item (Friedman 1956; Vickrey 1961). The valuation reflects the difference in profit between owning and not owning the item, and thus determines the value of the item for the bidder. However, their valuations depend on how they can use the item in the production process. We therefore integrate auction strategies into a production system to derive the valuation automatically.
- Third, auction theory often assumes that the probability distributions over each competitor's valuation of the auctioned item are known (Friedman 1956; Vickrey 1961). We use simulations of the production planning system to approximate these probability distributions automatically.

Our approach seamlessly incorporates these three research contributions by combining ideas from auction

theory, utility theory, and dynamic programming, resulting in a prototype of a supply chain management system with integrated auction strategies for the paper industry.

The paper is organized as follows: We first discuss an example of how auctions are used in the paper industry and then review the theory of auctions. In the next sections, we then describe the three shortcomings of the theory and develop solutions that address them. First, we derive the bidding strategy for a risk-averse decision maker in a single auction, using results from utility theory. Second, we integrate the resulting theory with a production planning system, using dynamic programming. Third, we simulate the resulting system to obtain both the valuation of an item for the decision maker and an approximation of the competitors' valuation distribution of the item, two parameters that are needed by the production planning system with integrated auction strategies. Finally, we show some experimental results obtained with the system and discuss future work.

An Example

We integrate risk-averse auction planning into a simplified version of a scheduling system for the paper industry that was co-developed by one of the authors and reduced the production cost of papermills by millions of dollars per year (Murthy *et al.* 1999). We study a papermill that uses both paper machines and cutting machines. Paper machines produce different kinds of paper rolls out of paper pulp. The cost of switching from producing one kind of paper to another kind of paper is substantial, and there are lower and upper limits on the length of paper that the machines can produce. Cutting machines then cut wide paper rolls to produce narrower ones.

Such a papermill could use the existing scheduling system to manage the production process, that is, determine how to use the machines. The production process is driven by orders. Whenever a new order is placed, the scheduling system updates the production plan by solving a finite horizon multi-stage optimization problem, where the horizon is either fixed or set to the latest delivery date. If an order cannot be delivered completely by the negotiated date, the papermill has to pay a penalty to the customer.

We are interested in integrating auctions into the scheduling system because it is sometimes advantageous for a papermill to buy paper rather than produce it, and the current growth of e-commerce makes it easy to do so via auctions. Consider, for example, a publisher that places an order for the paper needed to produce paperback books. Such orders usually contain two different kinds of paper, namely the paper needed for the covers and the paper needed for the pages. It can be advantageous for the papermill to purchase the paper for the covers, because the amount of paper needed for the covers is much smaller than the amount of paper needed for the pages, and both the minimum length constraints and the switching costs can make it very expensive to produce the paper for the covers. Combined orders of this kind and the resulting participation in auctions are a common scenario in the paper industry and a supply chain management system must be able to handle them.

The Auction Model

A suitable auction model for our purposes is the symmetric independent private values (SIPV) model (Klemperer 1999; McAfee & McMillan 1987; Wolfstetter 1996). In this model,

- only one item is for sale, and the seller is willing to sell it to the highest bidder for any positive price;
- the number of bidders N is known to all bidders;
- each bidder i knows their own valuation v_i for the auctioned item (that is, the difference in profit between owning and not owning the item);
- no bidder knows the other bidders' valuations for the auctioned item, but these valuations are independent random variables drawn from a given continuous distribution $F(v)$ with density $f(v)$ over the nonnegative real-valued bids, and this distribution is known to all bidders; and
- the bidders are indistinguishable.

The SIPV model has been studied in the context of two standard types of auctions, namely the first- and second-price sealed-bid auctions (Klemperer 1999; McAfee & McMillan 1987; Wolfstetter 1996). In these auctions, the bidders submit secret bids. The highest bid wins the item. The winner pays the highest bid ("first price") in the first-price sealed-bid auction but only the highest losing bid ("second price") in the second-price sealed-bid auction.

Both kinds of auctions are modeled as noncooperative games under incomplete information, often under the assumption that the behavior of a bidder is determined by a differentiable bidding function. A bidding function $b(\cdot)$ maps the valuation $v \geq 0$ of the bidder for the auctioned item to their nonnegative bid $b(v) \geq 0$, resulting in a deterministic bidding strategy. Obviously, $b(0) = 0$. The optimal bidding functions $b^*(\cdot)$ must be in an equilibrium, that is, no bidder can do better by changing their bidding function provided that the other bidders do not change their bidding functions either. The optimal bidding functions of all bidders are the same due to symmetry in the SIPV model. Furthermore, they are monotonically increasing in the valuation, for both first- and second-price sealed-bid auctions (McAfee & McMillan 1987; Wolfstetter 1996). It follows immediately that the strategy that maximizes the expected profit for a decision maker in a second-price sealed bid auction is to bid his or her valuation, simply because this strategy dominates all other strategies. Consequently, the optimal bidding function is the identity function. Determining the strategy that maximizes the expected profit for a decision maker in a first-price sealed bid

auction is more complicated. The optimal bidding function is known to be (McAfee & McMillan 1987; Wolfstetter 1996)

$$b^*(v) = v - \frac{\int_0^v F(t)^{N-1} dt}{F(v)^{N-1}}. \quad (1)$$

Note that the bids have to be smaller than the valuation for first-price sealed-bid auctions since bidding the valuation results in zero profit.

Risk and Utility Functions

We have already seen that buying paper at auctions can sometimes be cheaper for papermills than producing it. However, the downside to auctions is that they may not be won. In this case, the papermill faces the possibility of incurring huge penalties for not being able to satisfy some orders. Furthermore, a large number of unsatisfied orders can ruin the reputation of the papermill in the long run. Thus, decision makers for papermills often make conservative (risk-averse) decisions when deciding whether to participate in auctions and how much to bid. For example, they bid high to increase their chances of winning the auction.

We suggest that one can use utility theory to model the bidding strategies of risk-averse decision makers in a principled way. Utility theory is a subfield of decision theory that provides a normative framework for rational decision making under uncertainty (von Neumann & Morgenstern 1947; Bernoulli 1738). Its main claim is that every rational decision maker has a monotonically increasing utility function $u(\cdot)$ that maps profits v into real-valued utilities $u(v)$ so that the decision maker always prefers the alternative with the highest expected utility. The form of the utility function determines the risk attitude of the decision maker. A decision maker is risk-averse if he or she is willing to accept a decrease in the expected profit to reduce the variance and thus the possibility of a large loss. The risk-averse utility functions most often used in utility theory are probably the concave exponential utility functions

$$u(v) = -\gamma^v,$$

where the parameter $0 < \gamma < 1$ determines the level of risk aversion (Watson & Buede 1987). If γ approaches one, the decision maker is less and less risk-averse and, in the limit, maximizes the expected profit (under appropriate assumptions) (Koenig & Simmons 1994a; Koenig 1998). As γ approaches zero, on the other hand, the decision maker becomes more and more risk-averse. Thus, concave exponential utility functions can model a continuous spectrum of risk aversion.

Risk-Averse Bidding

Research performed in the context of the SIPV model often assumes that decision makers want to maximize their expected profits (McAfee & McMillan 1987; Wolfstetter 1996). We have argued that we can take their risk aversion into account by assuming that they maximize the expected utility for a concave exponential utility function. Previous work has shown that the optimal bidding function for risk-averse decision makers in either first- and second-price sealed-bid auctions remains monotonically increasing in the valuation (Maskin & Riley 1984; Matthews 1987). Again, the strategy that maximizes the expected utility of the profit for a decision maker in a second-price sealed-bid auction is to bid his or her valuation, simply because this strategy dominates all other strategies. Consequently, the optimal bidding function is the identity function. Thus, the bids of risk-averse decision makers and decision makers that maximize the expected profit are identical for second-price sealed-bid auctions. However, this is not the case for first-price sealed-bid auctions. It is known that the bids of risk-averse decision makers are higher than the bids of decision makers that maximize their expected profits in first-price sealed-bid auctions (Maskin & Riley 1984; Matthews 1987). While this result explains why risk-averse decision makers bid high to increase their chances of winning the auction, it is qualitative in nature and does not specify the optimal bidding function. In the following, we thus derive the optimal bidding function for risk-averse decision makers with concave exponential utility functions in the context of first-price sealed-bid auctions in the SIPV model. We extend the symmetry assumptions of the SIPV model to include the assumption that all bidders have the same concave exponential utility function, and continue to assume that the optimal bidding function $b^*(\cdot)$ is differentiable and thus continuous. Under this assumption, the optimal bidding function is also strictly monotonically increasing (Maskin & Riley 1984; Matthews 1987). Thus, it has a differentiable inverse that we denote by $\sigma(\cdot)$, and it holds that $\sigma(b) = v$ iff $b = b^*(v)$.

Derivation of the Optimal Bidding Function: To derive the optimal bidding function for decision makers with concave exponential utility functions in the context of first-price sealed-bid auctions in the SIPV model, we first determine the probability $\rho(b)$ with which the decision maker wins the auction when he or she bids b. Without loss of generality, we consider bidder 1. Any other bidder $2 \leq i \leq N$ has valuation v_i and thus bids $b^*(v_i)$. Thus,

$$\begin{aligned} \rho(b) &= P(\max_{2 \leq i \leq N} b^*(v_i) < b) \\ &= P(b^*(v_2) < b \wedge \cdots \wedge b^*(v_N) < b) \\ &= \prod_{i=2}^{N} P(b^*(v_i) < b) = \prod_{i=2}^{N} P(v_i < \sigma(b)) \\ &= F(\sigma(b))^{N-1}. \end{aligned} \quad (2)$$

Later, we need the derivative of this probability with respect to b. It is

$$\rho'(b) = (N-1)F(\sigma(b))^{N-2} f(\sigma(b))\sigma'(b). \quad (3)$$

The expected utility $EU(v, b)$ of the profit obtained from participating in an auction with valuation v and

bid b is

$$EU(v,b) = \rho(b)u(v-b) + (1-\rho(b))u(0). \quad (4)$$

The optimal bid β maximizes $EU(v,b)$ and thus the derivative of $EU(v,b)$ with respect to b at β is zero. Thus,

$$\left.\frac{\partial EU(v,b)}{\partial b}\right|_{b=\beta} = 0$$

$$\rho'(\beta)u(v-\beta) - \rho(\beta)u'(v-\beta) - \rho'(\beta)u(0) = 0$$

$$\rho'(\beta)\left[u(v-\beta) - u(0)\right] = \rho(\beta)u'(v-\beta). \quad (5)$$

By substituting Equations (2) and (3) as well as $v = \sigma(\beta)$ into Equation (5) and then dividing both sides by $F(\sigma(\beta))^{N-2}$, we obtain

$$(N-1)f(\sigma(\beta))\sigma'(\beta)\left[u(\sigma(\beta)-\beta) - u(0)\right] = F(\sigma(\beta))u'(\sigma(\beta)-\beta). \quad (6)$$

This is a differential equation for $\sigma(\cdot)$, but we are interested in the differential equation for its inverse $b^*(\cdot)$. By substituting $\sigma(\beta) = v$, $\beta = b^*(v)$ and $\sigma'(\beta) = 1/b^{*\prime}(v)$, the relationship between the derivative of a function and the derivative of its inverse, into Equation (6), we obtain

$$(N-1)f(v)\left[u(v-b^*(v)) - u(0)\right] = F(v)u'(v-b^*(v))b^{*\prime}(v). \quad (7)$$

In general, this differential equation cannot be solved in closed form. However, we now show how one can obtain a closed form solution for concave exponential utility functions. In this case, Equation (4) simplifies to

$$EU(v,b) = -\left[\rho(b)\gamma^{v-b} + (1-\rho(b))\right]. \quad (8)$$

Similarly, since the derivative of the utility function with respect to v is $u'(v) = -(\ln\gamma)\gamma^v$, Equation (7) simplifies to

$$(N-1)f(v)\left[-\gamma^{v-b^*(v)} + 1\right] = F(v)\gamma^{v-b^*(v)}(-\ln\gamma)b^{*\prime}(v)$$

$$(N-1)f(v)\left[\gamma^v - \gamma^{b^*(v)}\right] = F(v)\gamma^v(\ln\gamma)b^{*\prime}(v). \quad (9)$$

We can solve Equation (9) with the substitution

$$B(v) = \gamma^{-b^*(v)}F(v)^{N-1}. \quad (10)$$

For Equation (10), it holds that

$$\gamma^{b^*(v)} = \frac{F(v)^{N-1}}{B(v)} \quad (11)$$

and also that

$$b^*(v) = (N-1)\log_\gamma F(v) - \log_\gamma B(v). \quad (12)$$

For Equation (10), it also holds that

$$\begin{aligned} B'(v) &= \gamma^{-b^*(v)}F(v)^{N-1}\left(-b^{*\prime}(v)\ln\gamma + (N-1)\frac{f(v)}{F(v)}\right) \\ &= B(v)\left(-b^{*\prime}(v)\ln\gamma + (N-1)\frac{f(v)}{F(v)}\right) \end{aligned}$$

and thus

$$b^{*\prime}(v)\ln\gamma = (N-1)\frac{f(v)}{F(v)} - \frac{B'(v)}{B(v)}. \quad (13)$$

By substituting Equations (11) and (13) into Equation (9), we obtain

$$(N-1)f(v)\left(\gamma^v - \frac{F(v)^{N-1}}{B(v)}\right) = F(v)\gamma^v\left((N-1)\frac{f(v)}{F(v)} - \frac{B'(v)}{B(v)}\right).$$

By expanding both sides of the equation and canceling the first terms and the common factors on both sides, we obtain

$$\begin{aligned} (N-1)f(v)F(v)^{N-2} &= \gamma^v B'(v) \\ B'(v) &= \gamma^{-v}\left[F(v)^{N-1}\right]' \\ B(v) &= C + \int_0^v \gamma^{-t}dF(t)^{N-1}. \end{aligned} \quad (14)$$

Since $F(0) = 0$ and thus $B(0) = \gamma^{-b^*(0)}F(0)^{N-1} = 0$, it follows that $C = 0$. By substituting Equation (14) with $C = 0$ into Equation (12), we obtain a solution of the differential equation for $b^*(v)$ of the form

$$b^*(v) = (N-1)\log_\gamma F(v) - \log_\gamma \int_0^v \gamma^{-t}dF(t)^{N-1}. \quad (15)$$

This is the optimal bidding function for risk-averse decision makers with concave exponential utility functions in the context of first-price sealed-bid auctions in the SIPV model, and is thus the equivalent of Equation (1) for risk-averse decision makers.

Properties of the Optimal Bidding Function: In the following we show that the optimal bidding function from Equation (15) has two desirable properties in the limit. First, the bidding function approaches Equation (1) as the decision maker becomes less and less risk-averse and thus becomes more and more interested in maximizing the expected profit. This was to be expected. Second, the optimal bidding function approaches the identity function as the decision maker becomes more and more risk-averse, and the profit is zero in the limit. This is consistent with the qualitative results in (Maskin & Riley 1984; Matthews 1987).

Case 1: As γ approaches one, the decision maker becomes less and less risk-averse. To determine

$$\lim_{\gamma\to 1} b^*(v) = \lim_{\gamma\to 1}\frac{(N-1)\ln F(v) - \ln\int_0^v \gamma^{-t}dF(t)^{N-1}}{\ln\gamma},$$

notice that both its numerator and denominator approach zero and we can thus apply L'Hôpital's rule to obtain

$$\begin{aligned} \lim_{\gamma\to 1} b^*(v) &= \lim_{\gamma\to 1}\frac{-\frac{\int_0^v(-t)\gamma^{-t-1}dF(t)^{N-1}}{\int_0^v \gamma^{-t}dF(t)^{N-1}}}{1/\gamma} \\ &= \frac{\int_0^v t\,dF(t)^{N-1}}{\int_0^v dF(t)^{N-1}} = v - \frac{\int_0^v F(t)^{N-1}dt}{F(v)^{N-1}}, \end{aligned}$$

which is the same as the optimal bidding function for a decision maker that maximizes the expected profit, as shown in Equation (1).

Case 2: As γ approaches zero, the decision maker becomes more and more risk-averse. This time, both the numerator and denominator approach minus infinity and we can thus continue to apply L'Hôpital's rule:

$$\begin{aligned}\lim_{\gamma\to 0+} b^*(v) &= \lim_{\gamma\to 0+} \frac{-\frac{\int_0^v (-t)\gamma^{-t-1}dF(t)^{N-1}}{\int_0^v \gamma^{-t}dF(t)^{N-1}}}{1/\gamma} \\ &= \lim_{\gamma\to 0+} \frac{\int_0^v t\gamma^{-t}dF(t)^{N-1}}{\int_0^v \gamma^{-t}dF(t)^{N-1}}.\end{aligned}$$

Notice that $F(v)^{N-1}$ is also a probability distribution function, with a density function $(N-1)F(v)^{N-2}f(v)$. The appendix shows that the last expression converges to v as γ goes to zero from the positive side (under reasonable assumptions). Therefore, it holds that

$$\lim_{\gamma\to 0+} b^*(v) = v.$$

Bidding and Production Planning

So far, we have derived the bidding strategy for a risk-averse decision maker in a single auction if the decision maker knows his or her valuation for the auctioned item. The valuation depends on how he or she can use the item in the production process. We now show how to integrate risk-averse bidding into a production planning system to derive the valuation automatically, using dynamic programming. To this end, we first show how a production planning system can, in principle, be based on dynamic programming and then how auctions can be integrated seamlessly into the resulting system.

Production Planning: A production planning problem can, in principle, be formulated as a discrete-time Markov decision problem (MDP) with a limited time horizon, that can then be solved with standard techniques from dynamic programming (Bertsekas 1987; Howard 1964). The MDP process starts at time 0. At time t, the process is in some state $s \in S$ and the decision maker can choose an action $a \in A_t(s)$ for execution. The execution then results in a transition to state $s' \in S_t(s,a)$ with probability $P_t(s'|s,a)$, the decision maker receives an immediate reward of r according to the reward distribution $r_t(r|s,a,s')$, and the time is increased by one. Execution stops at time $T+1$, the time horizon, and the decision maker can no longer obtain any immediate rewards. In the context of production planning for papermills, for example, the states correspond to the inventory levels and the status of the various machines, the actions correspond to decisions about how to use the machines, and rewards correspond to payments as well as production costs.

Suppose the state of the MDP process at time step t is s_t and the decision maker executes action a_t and receives immediate reward r_t. If the decision maker wants to maximize the expected sum of the immediate rewards, that is, the expected profit, the decision maker has to maximize the following expectation along all possible state sequences $\{s_0, s_1, \cdots, s_{T+1}\}$

$$E_{r_0,\cdots,r_T}\left[\sum_{t=0}^{T} r_t\right],$$

where the decision variables are the actions $a_t \in A_t(s_t)$ at time steps $t = 0, \cdots, T$. This can be done with techniques from dynamic programming because the immediate rewards are additive, resulting in the one-step decomposition

$$\begin{aligned}E_{r_0,\cdots,r_T}\left[\sum_{t=0}^{T} r_t\right] &= E_{r_0,\cdots,r_T}\left[r_0 + \sum_{t=1}^{T} r_t\right] \\ &= E_{r_0}\left[r_0 + E_{r_1,\cdots,r_T}\left[\sum_{t=1}^{T} r_t\right]\right].\end{aligned}$$

Based on this decomposition, dynamic programming calculates a value function V_t which maps states s at time step t to the largest expected profit that the decision maker can obtain if execution starts in state s at time step t. The value functions are recursively defined by the following equations, known as the Bellman Equations: (Bellman 1957)

$$\begin{aligned}V_t(s) &= \max_{a\in A_t(s)} E_{s',r}[r + V_{t+1}(s')] \\ V_{T+1}(s) &= 0\end{aligned}$$

for all states s, where $s' \in S_t(s,a)$. To maximize the expected profit, the decision maker should then execute the action $a \in A_t(s)$ at time step t that maximizes $E_{s',r}[r + V_{t+1}(s')]$.

Risk-Averse Production Planning: We have assumed that the decision maker is risk-averse with a concave exponential utility function. Thus, the decision maker does not maximize the expected profit but rather the expected utility of the profit and thus the expectation

$$E_{r_0,\cdots,r_T}\left[u\left(\sum_{t=0}^{T} r_t\right)\right].$$

In general, this objective cannot be decomposed. However, it is known that it can be decomposed for utility functions with the delta property (Howard & Matheson 1972) or, equivalently, constant local risk aversion (Pratt 1964). Concave exponential utility functions have this property. The objective can thus be decomposed for concave exponential utility functions, resulting in the one-step decomposition

$$\begin{aligned}E_{r_0,\cdots,r_T}\left[u\left(\sum_{t=0}^{T} r_t\right)\right] &= E_{r_0,\cdots,r_T}\left[-\gamma^{\sum_{t=0}^{T} r_t}\right] \\ &= E_{r_0,\cdots,r_T}\left[-\prod_{t=0}^{T}\gamma^{r_t}\right] = E_{r_0,\cdots,r_T}\left[\gamma^{r_0}\cdot\left(-\prod_{t=1}^{T}\gamma^{r_t}\right)\right]\end{aligned}$$

$$= E_{r_0}\left[\gamma^{r_0} \cdot E_{r_1,\cdots,r_T}\left[-\prod_{t=1}^{T}\gamma^{r_t}\right]\right]$$
$$= E_{r_0}\left[\gamma^{r_0} \cdot E_{r_1,\cdots,r_T}\left[u\left(\sum_{t=1}^{T} r_t\right)\right]\right].$$

Based on this decomposition, one can continue to apply dynamic programming, this time by calculating a value function V_t which maps states s at time step t to the largest expected utility of the profit if execution starts in state s at time step t (Koenig & Simmons 1994b). The value functions are recursively defined by the following equations:

$$V_t(s) = \max_{a \in A_t(s)} E_{s',r}\left[\gamma^r V_{t+1}(s')\right] \quad (16)$$
$$V_{T+1}(s) = u(0) = -1 \quad (17)$$

for all states s, where $s' \in S_t(s,a)$. To maximize the expected utility of the profit, the decision maker should then execute the action $a \in A_t(s)$ at time step t that maximizes $E_{s',r}\left[\gamma^r V_{t+1}(s')\right]$. In previous work, we have applied similar dynamic programming methods to sensor planning (Koenig & Liu 2000).

Risk-Averse Bidding and Risk-Averse Production Planning: We now integrate risk-averse bidding into the risk-averse production planning system derived so far and show how to determine the valuation for an auctioned item automatically. We do this by deriving the optimal bid directly and then calculating the valuation based on the optimal bidding function. It turns out that we can reuse the results that we have derived in the context of single auctions, and that the resulting valuations have very intuitive interpretations.

We can easily integrate bidding into the risk-averse dynamic programming formulation of production planning, as given by Equations (16) and (17). Certain time steps t now correspond to auctions. Their actions correspond to the possible bids b, and a bid of zero corresponds to the decision not to participate in the auction. The actions have uncertain outcomes. With probability $\rho(b)$ the auction is won. In this case, the immediate reward is the negative of the paid price p and the inventory level increases by the auctioned item (resulting in state s_+). With the complementary probability $1-\rho(b)$ the auction is not won. In this case, the immediate reward is zero and the inventory level does not change (remaining in state s). Consequently, the value functions for the auction time steps are recursively defined by the following equations

$$\begin{aligned} V_t(s) &= \max_b E_{s',p}\left[\gamma^{-p} V_{t+1}(s')\right] \\ &= \max_b E_{s'}\left[E_p[\gamma^{-p}|s']V_{t+1}(s')\right] \\ &= \max_b \Big[\rho(b)E[\gamma^{-p}|s_+]V_{t+1}(s_+) \\ &\qquad +(1-\rho(b))E[\gamma^0|s]V_{t+1}(s)\Big] \\ &= \max_b \Big[\rho(b)E[\gamma^{-p}|s_+]V_{t+1}(s_+) \\ &\qquad +(1-\rho(b))V_{t+1}(s)\Big]. \end{aligned}$$

For simplicity, we define
$EU_t(s,b) = \rho(b)E[\gamma^{-p}|s_+]V_{t+1}(s_+) + (1-\rho(b))V_{t+1}(s)$
so that $V_t(s) = \max_b EU_t(s,b)$. The price p and thus $EU_t(s,b)$ depend on the kind of auction. We therefore study first- and second-price sealed-bid auctions separately.

First-Price Sealed-Bid Auctions: In this case, $p = b$ and thus $E[\gamma^{-p}|s_+] = \gamma^{-b}$. Therefore,

$$\begin{aligned} EU_t(s,b) &= \rho(b)\gamma^{-b}V_{t+1}(s_+) + (1-\rho(b))V_{t+1}(s) \\ &= \left[\rho(b)\frac{V_{t+1}(s_+)}{V_{t+1}(s)}\gamma^{-b} + (1-\rho(b))\right]V_{t+1}(s). \end{aligned}$$

Comparing the expression with Equation (8) yields

$$EU_t(s,b) = EU\left(\log_\gamma \frac{V_{t+1}(s_+)}{V_{t+1}(s)}, b\right)(-V_{t+1}(s)).$$

$V_{t+1}(s)$ is negative according to Equation (16) and (17). Thus, $-V_{t+1}(s)$ is positive, and the bid that maximizes the first factor also maximizes $EU_t(s,b)$. According to Equation (15), the optimal bid is

$$\beta = b^*\left(\log_\gamma \frac{V_{t+1}(s_+)}{V_{t+1}(s)}\right).$$

The valuation therefore is

$$\begin{aligned} v &= \log_\gamma \frac{V_{t+1}(s_+)}{V_{t+1}(s)} \\ &= \log_\gamma[-V_{t+1}(s_+)] - \log_\gamma[-V_{t+1}(s)] \\ &= u^{-1}(V_{t+1}(s_+)) - u^{-1}(V_{t+1}(s)). \end{aligned}$$

In utility theory, the expression $u^{-1}(x)$ is called the certainty equivalent of the expected utility x. The valuation therefore is simply the difference in certainty equivalents of the state that results when the item is won and the state that results when the item is not won.

Second-Price Sealed-Bid Auctions: In this case, assume that the bid of the decision maker is b, and the highest bid made by others is b'. Then $p = b'$ and

$$E[\gamma^{-p}|s+] = E[\gamma^{-b'}|b' < b].$$

So,

$$EU_t(s,b) = \rho(b)E[\gamma^{-b'}|b' < b]V_{t+1}(s_+) + (1-\rho(b))V_{t+1}(s)$$

Recall that, for a random variable x with the distribution $F(x)$, the conditional expectation of a function $g(x)$, given that event A happens is

$$E[g(x)|A] = \frac{\int_A g(x)dF(x)}{P(A)}.$$

Thus, the conditional expectation in $EU_t(s,b)$ can be calculated as

$$\begin{aligned} \rho(b)E[\gamma^{-b'}|b' < b] &= \int_0^b \gamma^{-b'}d\rho(b') \\ &= \gamma^{-b}\rho(b) - \int_0^b \rho(b')(-\ln\gamma)\gamma^{-b'}db' \\ &= \gamma^{-b}\rho(b) + \ln\gamma\int_0^b \rho(b')\gamma^{-b'}db'. \end{aligned}$$

Therefore,

$$EU_t(s,b) = \left[\gamma^{-b}\rho(b) + \ln\gamma \int_0^b \rho(b')\gamma^{-b'}db'\right] V_{t+1}(s_+) + (1-\rho(b))V_{t+1}(s).$$

The optimal bid β maximizes $EU_t(s,b)$ and thus the derivative of $EU_t(s,b)$ with respect to b is zero.

$$\begin{aligned}\left.\frac{\partial EU_t(s,b)}{\partial b}\right|_{b=\beta} &= [(-\ln\gamma)\gamma^{-\beta}\rho(\beta) + \gamma^{-\beta}\rho'(\beta) \\ &\quad +(\ln\gamma)\gamma^{-\beta}\rho(\beta)]V_{t+1}(s_+) - \rho'(\beta)V_{t+1}(s) \\ &= \rho'(\beta)(\gamma^{-\beta}V_{t+1}(s_+) - V_{t+1}(s)) = 0\end{aligned}$$

The solution to the above equation is

$$\begin{aligned}\beta &= \log_\gamma \frac{V_{t+1}(s_+)}{V_{t+1}(s)} \\ &= u^{-1}(V_{t+1}(s_+)) - u^{-1}(V_{t+1}(s)).\end{aligned}$$

Since the optimal bidding function for second-price sealed-bid auctions is the identity function, the valuation is again the difference in certainty equivalents of the state that results when the item is won and the state that results when the item is not won.

Estimating the Valuation Distributions

Auction theory often assumes that the competitors' valuation distributions of the auctioned item are known. We use simulations of the production planning process to estimate these probability distributions automatically. This can easily be done because we have already shown that the valuation of a decision maker for an item is the difference in certainty equivalents of the state that results when the item is won and the state that results when the item is not won, in the context of both first- and second-price sealed-bid auctions. We thus use simulations of the production planning process to sample these certainty equivalents and then approximate the distribution of their differences with the observed frequencies. To ensure that we sample states from an approximately correct distribution, we first sample a production problem and solve it. The optimal state sequence $\{s^0, s^1, \cdots, s^{T+1}\}$ is deterministic since we assume that machine failures can be fixed quickly enough not to impact the production schedule, which is a realistic assumption at our level of detail. Next, we randomly pick a time step t and assume that an item is auctioned off at that time. The valuation is the difference of the certainty equivalent of s_+^t, the hypothetical state of winning the item at time t, and the certainty equivalent of s^t. Thus, our approach is a model-based approach and differs from the more common data-based approaches that estimate the valuation distributions from the observed bids (Bajari 1998; Hong & Shum 1998). It can, however, only approximate the actual valuation distribution, since the simulation only allows for a single auction in each sample

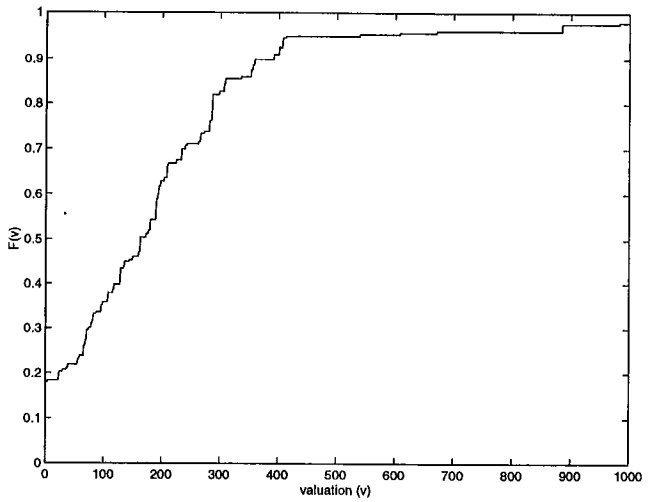

Figure 1: Estimated Valuation Distribution

of planning problems. Because we do not have a valuation distribution to start with, we do not know how to bid in auctions, nor the probability of winning. Therefore, if we allowed for more than one auction in each sample, we would not determine the optimal state sequence prior to the last auction, because the probability of winning the last auction affects the optimal schedule. Currently, we are investigating iterative methods to get more accurate estimates of the valuation distribution, where in each iteration the approximated valuation distribution obtained in the previous iteration is used for bidding. We are also investigating sampling methods which can be used for solving the stochastic version of the scheduling problem.

Experimental Results

We have implemented the system described so far, using a heuristic search approach instead of a dynamic programming approach to find optimal solutions to the scheduling problems. An advantage of heuristic search methods over dynamic programming methods is that they can use heuristics to focus their search, and thus are generally faster. To find an optimal solution, one needs to use a heuristic which underestimates the cost, or synonymously, overestimates the reward. We used a heuristic that was based on a linear programming relaxation, which overestimates the expected profit of any given state. We applied the resulting system to a simple problem involving a papermill that has one paper machine and one cutting machine. Figure 1 shows the valuation distribution that we obtained. It is interesting to note that the bidders may value the item for sale as low as 0 dollar if orders can be satisfied easily, and may value it as high as 1000 dollars if the papermill would not be able to satisfy some orders without the auction, thus resulting in high penalties. Figure 2 shows the resulting bidding function for four bidders in first-price sealed-bid auctions. The differ-

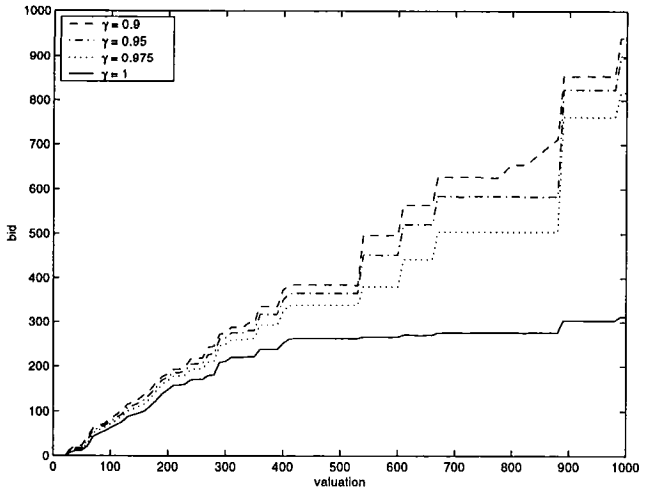

Figure 2: Optimal Bidding Functions for Different Risk Attitudes

ent lines correspond to different risk attitudes. The bids are always lower than the valuations, as expected. As γ increases, the decision makers become less and less risk-averse, which decreases their bids. If γ equals one, the decision makers maximize their expected profits. These results are consistent with the qualitative claims from auction theory (Maskin & Riley 1984; Matthews 1987).

Future Work

Our work so far has used simple (but reasonable) models of auctions and production processes. In the following, we describe some of our assumptions that we intend to relax in future work to widen the applicability of the production planning system with integrated auction strategies.

For example, we have assumed that only one item is auctioned off. Often, however, several items are auctioned off together and the bids are combinations of prices and quantities. It is still unknown which bidding strategy maximizes the expected profit for this setting, except in very simple cases (Tenorio 1997; 1999). Second, we have assumed that all bidders have the same degree of risk aversion. We are interested in relaxing this assumption, using previous results about bidders with asymmetric valuation distributions (Bajari 1997). Third, we have assumed that there are no interactions among auctions. However, this is often not the case. For example, if several auctions are held in parallel, it is often unnecessary to participate in all of them. Fourth, we have assumed that there are no constraints on the bids. However, the bids are often constrained by the available budget (Benoît & Krishna forthcoming). Fifth, we have solved the scheduling problem optimally. Since this is NP-hard in general, we could solve only relatively small production planning problems or problems with simple structures. To scale up, we intend to use approximation methods, resulting in less accurate valuation distributions. Sixth, we have used simulations of the production planning system to approximate the valuation distributions automatically. However, the simulations did not consider the effect of auctions on the valuation distributions, resulting in only approximately correct valuation distributions. We intend to investigate iterative methods that use simulations repeatedly to improve the valuation distributions.

Conclusions

In this paper, we described our first prototype of an auction-planning system that is part of a complete supply chain management system. First, we derived the bidding strategy for risk-averse decision makers in single auctions, using results from utility theory. Second, we integrated the resulting theory with a production planning system, using dynamic programming. Third, we simulated the resulting system to obtain both the decision maker's valuation of an item and an approximation of the competitors' valuation distributions of the item, two parameters that are needed by the production planning system with integrated auction strategies. The resulting system seamlessly combines ideas from auction theory, utility theory, and dynamic programming.

References

Bajari, P. 1997. *The First Price Auction With Asymmetric Bidders: Theory and Applictions.* Ph.D. Dissertation, Department of Economics, University of Minnesota.

Bajari, P. 1998. Econometrics of sealed-bid auctions. In *Proceedings of the Business and Economic Statistics Section of the American Statistical Association.*

Bellman, R. 1957. *Dynamic Programming.* Princeton University Press.

Benoît, J.-P., and Krishna, V. forthcoming. Multiple-object auctions with budget constrained bidders. *Review of Economic Studies.*

Bernoulli, D. 1738. Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5. Translated by L. Sommer, *Econometrica*, 22: 23–36, 1954.

Bertsekas, D. 1987. *Dynamic Programming, Deterministic and Stochastic Models.* Prentice Hall.

Friedman, L. 1956. A compititive-bidding strategy. *Operations Research* 4:104–112.

Hong, H., and Shum, M. 1998. Structural estimation of auction models. Working Paper, University of Toronto.

Howard, R., and Matheson, J. 1972. Risk-sensitive Markov decision processes. *Management Science* 18(7):356–369.

Howard, R. 1964. *Dynamic Programming and Markov Processes.* MIT Press, third edition.

Klemperer, P. 1999. Auction theory: A guide to the literature. *Journal of Economic Surveys* 13(3):227–286.

Klemperer, P., ed. 2000. *The Economic Theory of Auctions.* Edward Elgar.

Koenig, S., and Liu, Y. 2000. Representations of decision-theoretic planning tasks. In *Proceedings of the International Conference on AI Planning and Scheduling*, 187–195.

Koenig, S., and Simmons, R. 1994a. How to make reactive planners risk-sensitive. In *Proceedings of the International Conference on Artificial Intelligence Planning Systems*, 293–298.

Koenig, S., and Simmons, R. 1994b. Risk-sensitive planning. In *Proceedings of the AAAI Spring Symposium on Decision-Theoretic Planning*, 141–147. Available as AAAI Technical Report SS-94-06.

Koenig, S. 1998. Representation changes for planning with exponential utility functions. In *Proceedings of the Symposium on Abstraction, Reformulation, and Approximation*, 79–84.

Maskin, E., and Riley, J. 1984. Optimal auctions with risk averse buyers. *Econometrica* 52(6):1473–1518.

Matthews, S. 1987. Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica* 55(3):633–646.

McAfee, R. P., and McMillan, J. 1987. Auctions and bidding. *Journal of Economic Literature* 25:699–738.

Murthy, S.; Akkiraju, R.; Goodwin, R.; Keskinocak, P.; Rachlin, J.; Wu, F.; Kumaran, S.; Yeh, J.; Fuhrer, R.; Aggarwal, A.; Sturzenbecker, M.; Jayaraman, R.; and Daigle, B. 1999. Cooperative multi-objective decision-support for the paper industry. *Interfaces* 29(5):5–30.

Pratt, J. 1964. Risk aversion in the small and in the large. *Econometrica* 32(1-2):122–136.

Tenorio, R. 1997. On strategic quantity bidding in multiple unit auctions. *Journal of Industrial Economics* 45:207–217.

Tenorio, R. 1999. Multiple unit auctions with strategic price-quantity decisions. *Economic Theory* 13:247–260.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.

von Neumann, J., and Morgenstern, O. 1947. *Theory of games and economic behavior.* Princeton University Press, second edition.

Watson, S., and Buede, D. 1987. *Decision Synthesis.* Cambridge University Press.

Wolfstetter, E. 1996. Auctions: An introduction. *Journal of Economic Surveys* 10:367–420.

Appendix

In this appendix, we prove that

$$\lim_{\gamma\to 0+} \frac{\int_0^v t\gamma^{-t} f(t)dt}{\int_0^v \gamma^{-t} f(t)dt} = v$$

where $f(t) \geq 0$ is a probability density function over $[0, +\infty)$, and $v > 0$ is a finite number, as long as $f(t) > 0$ holds somewhere in the interval $[0, v]$.

We first notice that $f(t)$ can be uniformly approximated from above by a step function over the interval $[0, v]$, that is, for any given $\varepsilon > 0$, there exists a step function $s_\varepsilon(t)$, such that $0 \leq f(t) \leq s_\varepsilon(t)$ and $s_\varepsilon(t) - f(t) < \varepsilon$, for all $t \in [0, v]$.

Since

$$\int_0^v t\gamma^{-t} f(t)dt \leq \int_0^v v\gamma^{-t} f(t)dt = v\int_0^v \gamma^{-t} f(t)dt,$$

we have that

$$\begin{aligned} v &\geq \frac{\int_0^v t\gamma^{-t} f(t)dt}{\int_0^v \gamma^{-t} f(t)dt} \geq \frac{\int_0^v t\gamma^{-t} f(t)dt}{\int_0^v \gamma^{-t} s_\varepsilon(t)dt} \\ &> \frac{\int_0^v t\gamma^{-t}(s_\varepsilon(t) - \varepsilon)dt}{\int_0^v \gamma^{-t} s_\varepsilon(t)dt} \\ &= \frac{\int_0^v t\gamma^{-t} s_\varepsilon(t)dt}{\int_0^v \gamma^{-t} s_\varepsilon(t)dt} - \varepsilon\frac{\int_0^v t\gamma^{-t}dt}{\int_0^v \gamma^{-t} s_\varepsilon(t)dt} \end{aligned}$$

Therefore,

$$\begin{aligned} v &\geq \lim_{\gamma\to 0+} \frac{\int_0^v t\gamma^{-t} f(t)dt}{\int_0^v \gamma^{-t} f(t)dt} \\ &\geq \lim_{\gamma\to 0+} \frac{\int_0^v t\gamma^{-t} s_\varepsilon(t)dt}{\int_0^v \gamma^{-t} s_\varepsilon(t)dt} - \varepsilon \lim_{\gamma\to 0+} \frac{\int_0^v t\gamma^{-t}dt}{\int_0^v \gamma^{-t} s_\varepsilon(t)dt} \end{aligned} \tag{18}$$

Assume that $s_\varepsilon(t)$ can be represented by $n \geq 1$ segments of constant values as follows:

$$s_\varepsilon(t) = a_k, \qquad t_{k-1} \leq t < t_k$$

where $0 = t_0 < t_1 < \cdots < t_n = v$. Since $f(t) \geq 0$, $a_k \geq 0$ for all k. We have

$$\int_0^v \gamma^{-t} s_\varepsilon(t)dt = \sum_{k=1}^n a_k \int_{t_{k-1}}^{t_k} \gamma^{-t}dt \tag{19}$$

and

$$\int_0^v t\gamma^{-t} s_\varepsilon(t)dt = \sum_{k=1}^n a_k \int_{t_{k-1}}^{t_k} t\gamma^{-t}dt \tag{20}$$

Without loss of generality, we can assume $a_n > 0$. Otherwise, the corresponding term in either integral is zero, then we can set $n' = n - 1$ without changing the values of the above integrals, and we still have $n' \geq 1$ since $f(t)$ has to be strictly greater than zero. Recall that

$$\int_a^b \gamma^{-t}dt = -\frac{\gamma^{-t}}{\ln\gamma}\bigg|_a^b = \frac{1}{\ln\gamma}(\gamma^{-a} - \gamma^{-b}) \tag{21}$$

and

$$\int_a^b t\gamma^{-t}dt = -\frac{1}{\ln\gamma}\int_a^b t d\gamma^{-t}$$
$$= -\frac{1}{\ln\gamma}\left(t\gamma^{-t}\Big|_a^b - \int_a^b \gamma^{-t}dt\right)$$
$$= \frac{1}{\ln\gamma}(a\gamma^{-a} - b\gamma^{-b}) + \frac{1}{\ln^2\gamma}(\gamma^{-a} - \gamma^{-b}) \quad (22)$$

Substituting Equation (21) back into Equation (19) yields

$$\int_0^v \gamma^{-t}s_\varepsilon(t)dt = \sum_{k=1}^n a_k \int_{t_{k-1}}^{t_k} \gamma^{-t}dt$$
$$= \sum_{k=1}^n a_k \frac{1}{\ln\gamma}(\gamma^{-t_{k-1}} - \gamma^{-t_k})$$
$$= \frac{1}{\ln\gamma}\left(\sum_{k=1}^n a_k\gamma^{-t_{k-1}} - \sum_{k=1}^n a_k\gamma^{-t_k}\right)$$
$$= \frac{1}{\ln\gamma}\left(a_1 + \sum_{k=1}^{n-1} a_{k+1}\gamma^{-t_k} - \sum_{k=1}^{n-1} a_k\gamma^{-t_k} - a_n\gamma^{-v}\right)$$
$$= \frac{1}{\ln\gamma}\left(a_1 - a_n\gamma^{-v} + \sum_{k=1}^{n-1}(a_{k+1} - a_k)\gamma^{-t_k}\right) \quad (23)$$

Similarly, we can obtain

$$\int_0^v t\gamma^{-t}s_\varepsilon(t)dt$$
$$= \frac{1}{\ln\gamma}\left(-a_n v\gamma^{-v} + \sum_{k=1}^{n-1}(a_{k+1} - a_k)t_k\gamma^{-t_k}\right)$$
$$+ \frac{1}{\ln^2\gamma}\left(a_1 - a_n\gamma^{-v} + \sum_{k=1}^{n-1}(a_{k+1} - a_k)\gamma^{-t_k}\right) \quad (24)$$

From Equations (23) and (24), we have

$$\lim_{\gamma\to 0+}\frac{\int_0^v t\gamma^{-t}s_\varepsilon(t)dt}{\int_0^v \gamma^{-t}s_\varepsilon(t)dt}$$
$$= \lim_{\gamma\to 0+}\left(\frac{-a_n v\gamma^{-v} + \sum_{k=1}^{n-1}(a_{k+1} - a_k)t_k\gamma^{-t_k}}{a_1 - a_n\gamma^{-v} + \sum_{k=1}^{n-1}(a_{k+1} - a_k)\gamma^{-t_k}} + \frac{1}{\ln\gamma}\right)$$
$$= \lim_{\gamma\to 0+}\frac{-a_n v + \sum_{k=1}^{n-1}(a_{k+1} - a_k)t_k\gamma^{v-t_k}}{-a_n + a_1\gamma^v + \sum_{k=1}^{n-1}(a_{k+1} - a_k)\gamma^{v-t_k}}$$
$$= \frac{-a_n v}{-a_n} = v, \quad (25)$$

and from Equations (22) and (24), we have

$$\lim_{\gamma\to 0+}\frac{\int_0^v t\gamma^{-t}dt}{\int_0^v \gamma^{-t}s_\varepsilon(t)dt}$$
$$= \lim_{\gamma\to 0+}\frac{-v\gamma^{-v} + \frac{1}{\ln\gamma}(1 - \gamma^{-v})}{-a_n\gamma^{-v} + a_1 + \sum_{k=1}^{n-1}(a_{k+1} - a_k)\gamma^{-t_k}}$$
$$= \lim_{\gamma\to 0+}\frac{-v + \frac{1}{\ln\gamma}(\gamma^v - 1)}{-a_n + a_1\gamma^v + \sum_{k=1}^{n-1}(a_{k+1} - a_k)\gamma^{v-t_k}}$$
$$= \frac{-v}{-a_n} = \frac{v}{a_n} > 0 \quad (26)$$

Substituting Equations (25) and (26) back into Equation (18), we have

$$v \geq \lim_{\gamma\to 0+}\frac{\int_0^v t\gamma^{-t}f(t)dt}{\int_0^v \gamma^{-t}f(t)dt} \geq v - \varepsilon\frac{v}{a_n}$$

Since ε is an arbitrary positive number, we have

$$\lim_{\gamma\to 0+}\frac{\int_0^v t\gamma^{-t}f(t)dt}{\int_0^v \gamma^{-t}f(t)dt} = v$$