

Auction-based Queue Disciplines

Thomas Kittsteiner and Benny Moldovanu[□]

February 14, 2003; Preliminary and Incomplete

Abstract

Each one of several impatient agents has a job that needs to be processed by a server. The server can process the jobs sequentially, one at a time. Agents are privately informed about the realization of a random variable representing processing time. If the cost of delay is represented by a concave function, the efficient shortest processing time \bar{r} st schedule arises in the equilibrium of a simple auction where agents bids for slots in the queue. If the cost function is convex, the equilibrium yields the anti efficient longest processing time \bar{r} st schedule. The performance of the auction in this case (both efficiency and revenue) can be improved by capping bids from above. Finally, we show that the mechanisms that minimize the expected total waiting cost cannot depend on the private information available to the agents.

1 Introduction

Imagine that you arrive with a bulky package at the department's copier machine together with one of your esteemed colleagues. He/She says to you: "I have only a few pages to copy. May I do it \bar{r} st?" Most of us usually agree to this request, and indeed this courtesy is well-founded in economic theory: total waiting cost is minimized if shorter jobs are processed before longer ones. Moreover, it seems intuitive that your colleague should be willing to pay more than you for the right to be \bar{r} st since he/she can thus avoid a longer delay than you. But sometimes the colleague's job turns out to be longer than the announced few pages ("I just noticed that I also need this chapter"), and in other situations everyone claims to have the shorter job... Are these announcements sincere? Should we base the queue discipline upon them? This paper is about the design of simple pricing mechanisms (e.g., auctions or lotteries) for allocating slots in a queue among impatient, privately informed agents.

[□]Kittsteiner: Nu±eld College, Oxford University; Moldovanu: Department of Economics, University of Bonn

Queueing theory¹ was developed in order to provide models that predict the behavior of systems providing service for randomly fluctuating customer demand. Besides a very large and valuable theoretical literature, there are many practical applications to traffic flow (e.g., vehicles at toll booths, aircraft at airport landing or take-off gates, network communication), scheduling (e.g., patients in hospitals, jobs on machines, programs on computers) and facility design (e.g., banks, post offices, amusement parks and fast-food restaurants). A standard task for the queueing analyst is to determine an appropriate measure for system performance (which depends of course on the customers' and system's characteristics) and to design an "optimal system" according to such a measure.

A queueing system is generally described by several basic characteristics pertaining to the stochastic arrival pattern of possibly impatient customers (e.g., interarrival times, the possibility to balk before or renege after entering the queue, the possibility to jockey for position, etc.) and to the pattern of service (e.g., system capacity, number of service channels, queue discipline, etc...).

Practically the entire queueing literature views customers as non-strategic agents endowed with some (possibly) random characteristics. The process generating these characteristics is assumed to be common knowledge. In particular, the theory does not consider the interplay between system design (e.g., queue discipline) and the strategic incentives to manipulate accessible information arising from the desire to improve one's position in the queue and thus to increase one's utility.

In this paper we consider the most basic scheduling problem with one server and with impatient customers all arriving at the same time. Each customer has a job that needs to be processed, and the server can sequentially process one job at a time². The design problem reduces to determining the allocation of slots in the waiting line. Each customer is privately informed about her needed processing time, and agents incur costs of delay. From the point of view of other agents, the processing time of a specific agent is a random variable governed by a common-knowledge distribution. The total waiting cost born by an agent i that is allocated the j^{th} slot depends both on i 's privately known processing time and on the processing times of the agents scheduled to slots $1; 2; \dots; j-1$: Thus agent i 's utility depends both on the allocation of slots to other agents and on information that is ex-ante available to other agents. In the language of mechanism design, we obtain a setting with both allocative and informational externalities. An important assumption that is implicit in our analysis is that processing times are too costly to monitor ex-post or, equivalently, that no additional fees can be imposed after processing times have been realized and service has been completed.

¹see for example the recent textbook by Gross and Harris (1998).

²We assume here that all agents can be served, but our results do not qualitatively change if the number of slots is less than the number of jobs to be processed.

In the above framework we first consider a natural auction procedure in which agents bid for slots: the highest bidder gets the first slot, the second-highest bidder gets the second slot, and so on till all slots are allocated. All agents pay their own bid. It turns out that the performance of this auction crucially depends on the form of the function describing the costs of delay. If this function is concave, agents with a shorter processing time bid in a Bayes-Nash equilibrium more than agents with a longer processing time and, for any realization of the stochastic processing times, the auction implements the efficient shortest processing time first (SEPT) schedule. In contrast, if the cost function is convex, agents with longer processing times bid in equilibrium more than agents with shorter times, and the auction implements the "anti-efficient" longest processing time first (LEPT) schedule.

Since the case of convex cost functions is the more pertinent one for most applications (e.g., consider the ubiquitous exponential cost functions or the presence of deadlines) it makes sense to inquire whether there are mechanisms that perform better than the auction in this case. We first analyze auctions with bid caps, i.e., auctions where agents are constrained to make bids that are lower than a pre-determined maximum. Since with convex cost functions the high bids come from agents with long processing time, constraining such bidders to a maximum bid implies that the allocation of slots among these bidders will contain a random element. Such a lottery necessarily improves upon the welfare attained by the LEPT schedule implemented by the unconstrained auction.

In many situations of interest the server is owned by an agent that needs to raise revenue in order to maintain operation and/or to make a profit (think about a data processing firm or an airport operator). At first sight it seems that, whenever revenue raising is important, constraining the high bidders via a bid cap will lower expected revenue. But this intuition is misleading since in the constrained auction there are in fact agents with low processing times that bid higher than their respective bids in the unconstrained auction. While bid caps are shown never to be revenue enhancing in an auction with concave cost functions, we find that bid caps may raise revenue if the cost function is convex (in some cases it is even revenue maximizing to impose an extremely low cap such that the auction degenerates to a pure lottery where agents pay a fixed fee!) Thus, with convex cost functions, schedule auctions constrained by bid caps may perform better than unconstrained auctions on both efficiency and revenue measures.

Since we found that we can enhance performance by using procedures that do not, or only partially condition on privately available information (lotteries and capped auctions, respectively) we are next interested to characterize mechanisms that achieve the highest possible welfare for the customers (i.e., minimize total expected waiting costs) subject to the incentive compatibility constraint (i.e., subject to the requirement that their outcome arises in the equilibrium of a game played by strategic, privately informed agents). It turns out that, in a

large class of mechanisms, lotteries are indeed optimal if the agents have the same expected processing time. If more information about the respective distribution of processing time is available (e.g, if it is known that one agent's distribution stochastically dominates another), this information can be used to determine optimal schedules. Surprisingly, even if it is the case that the individual distributions of processing times are stochastically ordered in the usual sense, it is not necessarily the case that the random variable governing the total costs of delay associated with the SEPT schedule (based on the expectations of processing times) stochastically dominates all other schedules. But such a result holds if the individual processing times are ordered in the likelihood ratio sense.

The basic scheduling problem with interdependent costs has been introduced by Hain and Mitra (2002)³. These authors show that concavity of the cost function is a necessary condition for the implementability of the efficient schedule in ex-post equilibria (note that any ex-post equilibrium is Bayes-Nash). Their main result is that, for cost functions that are concave polynomials of degree less than or equal to $n - 2$ (where n is the number of agents and slots), a generalized Clarke-Groves-Vickrey mechanism can be constructed⁴ that is efficient and budget balanced. Note that in our framework the auction's designer (who has no private information) is a residual claimant and budget balancedness is satisfied per definition.

Wellmann et. al. (2001) study private values scheduling problems without waiting costs: agents derive the same utility if their jobs are completed early or late as long as this is done before a deadline. These authors apply insights gained from the theory of matching markets.

Gavious, Moldovanu and Sela (2002) analyze a private value all-pay auction for a single object (there are no externalities of any kind) where the seller can impose bid caps. While a bid-cap is disadvantageous if the function describing the bid cost is linear or concave, it is shown that a bid cap may increase revenue if this function is convex.

This paper is organized as follows: In Section 2 we describe the design problem arising from a scheduling problem with waiting costs. In Section 3 we derive Bayes-Nash equilibria of a multi-object auction that allocates slots in the queue based on the respective bids. In Section 4 we study auctions with bid caps and analyze the effects of these on auction efficiency and revenue. In Section 5 we focus on the case of convex cost functions and show that the welfare maximizing incentive compatible mechanism (which minimizes expected processing time) does not condition on the agents' private information. Finally, we connect this result to well-known insights about the application of stochastic orders to queueing problems.

³Mitra (2001) focuses on efficiency and budget-balancedness for simpler scheduling problems with private values.

⁴The construction of the CGV mechanism is based on general insights about efficient implementation for multi-object auctions with interdependent valuations due to Dasgupta and Maskin (2000) and Jehiel and Moldovanu (2001).

2 The Model

A processing unit (e.g., a computer, a landing/take-off gate, a repair tool) is available to n risk-neutral agents. Each agent needs to perform a task (or job), and the agents' jobs can be processed sequentially, one job at a time. Each agent i has information about her own needed processing time t_i , which is the realization of a random variable with support $[\underline{t}; \bar{t}]$: The distribution of t_i is given by a strictly increasing and continuous distribution function F with density f : F is common knowledge and the individual processing times are independently distributed.

A schedule is a permutation $\sigma \in S_n$; $\sigma = (\sigma(1); \dots; \sigma(n))$ where S_n is the set of all one to one mappings $\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$: In particular, $\sigma(1)$ is the index of the agent who is served first, $\sigma(2)$ the index of the agent served second, and so forth. All agents bear a cost of waiting given by a strictly increasing function $C: \mathbb{R}^+ \rightarrow \mathbb{R}^+$; i.e. an agent waiting for t time units for her job to be finished bears a cost $C(t)$: Given a schedule $\sigma \in S_n$, the cost borne by an agent who is served j th is given by $C\left(\sum_{k=1}^j t_{\sigma(k)}\right)$;

All agents derive an utility V from the completion of their job (net of time costs). This utility is assumed to be common knowledge. We assume that $V \leq E_{(t_2; \dots; t_n)} C\left(\sum_{k=2}^n t_k\right)$ to ensure that expected waiting costs never exceed V .

Given a realization of types $(t_1; \dots; t_n)$ and a schedule σ , the utility of an agent i with $\sigma(i) = j$ is given by $V_i - C\left(\sum_{k=1}^j t_{\sigma(k)}\right)$ since her waiting time is given the sum of processing times for all predecessors and her own processing time. Hence we have a model with inter-dependent valuations: an agent's utility is directly influenced by private information available only to other agents (the processing time of her predecessors in the queue).

Denote by $\Delta(S_n)$ the set of all probability distributions on S_n : In a direct revelation mechanism $(\hat{A}; p)$ each agent i reports a processing time $b_i \in [\underline{t}; \bar{t}]$. Given reports $b = (b_1; \dots; b_n)$, the designer implements a schedule $\sigma \in S_n$ according to the scheduling rule $\hat{A}: [\underline{t}; \bar{t}]^n \rightarrow \Delta(S_n)$; and receives payments $p = (p_1; \dots; p_n)$, where $p_i(b_1; \dots; b_n)$ denotes the payment of agent i .

A mechanism $(\hat{A}; p)$ truthfully implements the rule \hat{A} if truth-telling is a Bayes-Nash equilibrium in the game induced by $(\hat{A}; p)$ and the agents' utility functions. Such a mechanism is called incentive compatible (IC).

The interim utility of agent i given her type t_i , her announcement b_i and given truth-telling of the other agents is given by:

$$\begin{aligned}
 U_i(t_i; b_i) &= E_{t_{-i}} u_i(t_i; b_i; t_{-i}) \\
 &= V_i - E_{t_{-i}} \left[\sum_{\sigma \in S_n} \hat{A}(\sigma; b_i; t_{-i}) \left(C\left(\sum_{k=1}^{\sigma(i)} t_{\sigma(k)}\right) - p_i(b_i; t_{-i}) \right) \right]
 \end{aligned}$$

We use the following notation:

$$U_i(t_i) := U_i(t_i; t_i):$$

The overall performance of a mechanism $(\bar{A}; p)$ is measured by the total expected cost due to waiting, i.e. by

$$P(\bar{A}) = E_t \sum_{i=1}^n C_i(t_i; t_i) \quad \bar{A}(t_i; t_i) \quad C_i(t_i; t_i) \quad t_i^{(k)}$$

We call an incentive compatible mechanism optimal if it minimizes $P(\bar{A})$ in the class of all incentive compatible mechanisms. Note that, for a given and fixed vector of processing times, total waiting cost is minimized by the well known shortest processing time first (SEPT) schedule where agents with a shorter processing time are served before those with longer processing time. A scheduling rule \bar{A}^* is called ex-post efficient if it yields the SEPT schedule for any realization of the random vector of processing times. Obviously, an incentive compatible mechanism that implements an ex-post efficient allocation rule is optimal.

3 Auction-based queue disciplines

In this section we analyze the equilibrium schedules arising in the following auction: agents simultaneously submit sealed bids; the highest bidder is served first, the second highest bidder is served second, and so forth; finally, each bidder has to pay his bid. This multi-object auction for n slots constitutes a simple and natural way of implementing a schedule based on the agents' information about processing times. Obviously, bidders do not making losses by participating in this auction since they can assure themselves a positive payoff by bidding zero and being queued last. The next result shows that equilibrium behavior crucially depends on the structure of the cost function C : We use the following notation: $\sum_{j=2}^n t_j = 0$:

Theorem 1 Define

$$\begin{aligned} \bar{C}_1(x) &:= C_1(x) = C(x) + C(x) \quad \bar{A} \quad \bar{A} \quad \# \\ \bar{C}_i(x) &:= E_{(t_2, \dots, t_n)} C(x + t_j) + C(x + t_j) \quad \bar{A} \quad \bar{A} \quad \# \\ \underline{C}_i(x) &:= E_{(t_2, \dots, t_n)} C(x + t_j) + C(x + t_j) \quad \bar{A} \quad \bar{A} \quad \# \end{aligned}$$

The unique symmetric equilibrium of the slot auction is given by the strictly decreasing function

1.

$$b_{cave}(t) = (n_i - 1) \prod_{l=1}^{n_i-2} \frac{1}{l_i - 1} \int_t^{\infty} (1 - F(x))^{n_i - l_i - 1} F^{l_i - 1}(x) \bar{C}_l(x) f(x) dx \quad (1)$$

if C is strictly concave, and by the strictly increasing function

2.

$$b_{vex}(t) = (n_i - 1) \prod_{l=1}^{n_i-2} \frac{1}{l_i - 1} \int_t^{\infty} F(x)^{n_i - l_i - 1} (1 - F(x))^{l_i - 1} \underline{C}_l(x) f(x) dx$$

if C is strictly convex.

Proof. Assume first that C is concave and all agents other than agent 1 bid according to a strictly decreasing function b. The expected utility of agent 1 with processing time t_1 who bids as if she were of type b_1 is given by

$$U^{i,t_1; b_1^c} = V_i \prod_{l=1}^{n_i-1} \frac{1}{l_i - 1} \int_{t_1}^{\infty} (1 - F(x))^{n_i - l_i - 1} F^{l_i - 1}(x) f(x) dx \cdot E_{(t_2, \dots, t_l)} \left[C \left(t_1 + \sum_{j=1}^{l-1} t_j \cdot b_j; j = 2, \dots, l \right) \right]$$

Differentiating with respect to b_1 and rearranging the terms yields

$$\frac{\partial U^{i,t_1; b_1^c}}{\partial b_1} = i (n_i - 1) \prod_{l=1}^{n_i-1} \frac{1}{l_i - 1} \int_{t_1}^{\infty} (1 - F(x))^{n_i - l_i - 1} F^{l_i - 1}(x) f(x) dx \cdot E_{(t_2, \dots, t_l)} \left[C \left(t_1 + \sum_{j=1}^{l-1} t_j \cdot b_j; j = 2, \dots, l \right) \right] \frac{db^{i,b_1^c}}{db_1}$$

The first order condition $\frac{\partial U(t_1; b_1)}{\partial b_1} = 0$ gives

$$i \frac{db^{i,b_1^c}}{db_1} \Big|_{b_1=t_1} = (n_i - 1) \prod_{l=1}^{n_i-1} \frac{1}{l_i - 1} \int_{t_1}^{\infty} (1 - F(x))^{n_i - l_i - 1} F^{l_i - 1}(x) f(x) dx \cdot E_{(t_2, \dots, t_l)} \left[C \left(t_1 + \sum_{j=1}^{l-1} t_j \cdot t_1; j = 2, \dots, l \right) \right]$$

This necessary condition is obviously fulfilled for the decreasing function b_{cave} as given by (1) since C is strictly increasing (and hence the conditional expectation is positive). The sufficient condition $\frac{\partial^2 U(t_1; b_1)}{\partial b_1^2} > 0$ for all $t_1; b_1$ is satisfied, since

$$\frac{\partial^2 U^{i,t_1; b_1^c}}{\partial b_1^2} \Big|_{b_1=t_1} = i (n_i - 1) \prod_{l=1}^{n_i-1} \frac{1}{l_i - 1} \int_{t_1}^{\infty} (1 - F(x))^{n_i - l_i - 1} F^{l_i - 1}(x) f(x) dx \cdot E_{(t_2, \dots, t_l)} \left[C \left(t_1 + \sum_{j=1}^{l-1} t_j \cdot t_1; j = 2, \dots, l \right) \right]$$

and since the derivative $C' = \frac{d}{dt}C(t)$ is strictly decreasing if C is strictly concave. If C is convex the proof is similar and therefore omitted here. ■

The above result shows that the efficient SEPT schedule is implemented by the auction if the cost function C is concave. In the case of a convex cost function the equilibrium bidding function is strictly increasing in processing time. Hence, agents are queued in the reverse order meaning that those with longer processing time are served first (this is the well known LEPT policy). In this case the auction yields the worst possible schedule since it maximizes $P(\bar{A})$.

The main reason behind these contrasting results is as follows: Agent's i cost from being delayed for a period of time t , $C(t + t_i) - C(t_i)$; is increasing in t_i if C is convex and decreasing if C is concave. Hence, if C is convex, it is more costly for an agent with a longer processing time to queue for some time t (before her own job is processed) than for an agent with a shorter processing time. The need to avoid the higher cost is expressed by a higher bid in the auction, yielding the increasing bid function for the case of convex cost functions. The opposite occurs for a concave cost function.

The "anti-efficient" auction outcome for convex cost functions suggests to look for mechanisms that perform better in this case. We study such mechanisms in the next section.

4 Slot auctions with bid caps

Assume that bidders are not allowed to submit bids that are above a predetermined common-knowledge maximum bid \bar{b} . If $m > 1$ agents submit bids equal to \bar{b} , then each of these bidders is given any of the first m slots in the queue with probability $\frac{1}{m}$.

Theorem 2 1) Let C be concave and assume that the bid cap \bar{b} satisfies

$$E_{(t_2, \dots, t_n)} \left[C\left(\bar{t} + \sum_{j=2}^n t_j\right) - C\left(\bar{t} + \sum_{j=2}^k t_j\right) \right] < \bar{b} < b(\bar{t}) \quad (2)$$

The schedule auction with bid cap \bar{b} has a unique symmetric equilibrium characterized by a type $e = e(\bar{b})$, strictly decreasing in \bar{b} , such that

$$e_{cave}(\bar{b}; t) = \begin{cases} b_{cave}(t) & \text{if } t > e \\ \bar{b} & \text{if } t \leq e \end{cases} \quad (3)$$

2) Let C be convex and assume that the bid cap \bar{b} satisfies

$$E_{(t_2, \dots, t_n)} \left[C\left(\bar{t} + \sum_{j=2}^n t_j\right) - C\left(\bar{t} + \sum_{j=2}^k t_j\right) \right] < \bar{b} < b\left(\bar{t} + \sum_{j=2}^n t_j\right) \quad (4)$$

The schedule auction with bid cap b has a unique symmetric equilibrium characterized by a type $e = e(b)$, strictly increasing in b , such that

$$b_{\text{vex}}(t) = \begin{cases} b_{\text{vex}}(t) & \text{if } t \leq e \\ b & \text{if } t > e \end{cases} \quad (5)$$

Proof. See Appendix. ■

By choosing an appropriate bid cap the auctioneer is able to determine the types of agents that pool in equilibrium. Whereas the auctioneer can only decrease efficiency by introducing a bid cap $b < b(t)$ in the case of a concave cost function (note that performance is monotonically decreasing in b), an effective bid cap $b < b(t)$ necessarily increases performance in the case of a convex cost function since, in that case, the auction performs worst among all possible scheduling mechanisms. Moreover, overall performance is increasing in b if C is convex. In particular, by setting a bid cap equal to the lower bound in (4) all agents bid b and hence are given a certain position in the queue with probability $\frac{1}{n}$: In this case all agents pay a fixed fee, and the schedule is determined by an equal chance lottery among all agents.

4.1 Revenue considerations

We now turn to an analysis of bid caps on the designer's revenue. We first show that, if the cost function is concave, the designer cannot profit by setting an effective bid cap. Thus, the ensuing efficiency loss in this case has necessarily a negative effect on her revenue. The result is not trivial since some bidders bid in the constrained auction more than in the unconstrained one. Hence, an auctioneer interested in revenue needs to balance this positive effect of bid caps with the negative effect caused by the fact that some other bidders cannot bid more than b .

Theorem 3 Assume that C is concave. Then the designer's revenue in an auction with an effective bid cap is lower than her revenue in the auction without bid cap.

In the case of a convex cost function it is interesting to note that the introduction of a bid cap causes some low-type bidders to bid b even if they submitted lower bids in the unrestricted auction (because this significantly increases their probability of getting a better slot). The next result shows that if the cost function is not "too convex", then it is optimal for a revenue-maximizing designer to set the lowest relevant effective bid cap. Thus, increasing the efficiency of the auction is also beneficial in terms of revenue, and the optimum is achieved with a lottery that does not condition on the agents' private information


Theorem 4 Assume that C is convex. For any distribution function F and for any number of jobs n , there exists a constant $K_n > 0$ such that if $C'' < K_n$ the revenue maximizing bid cap is given by

$$b_M = E_{(t_2, \dots, t_n)} \left[C \left(t + \frac{1}{n-1} \sum_{j=2}^n t_j \right) \right]$$

Proof. see Appendix ■

In contrast to the finding above, there exist environments with strongly convex cost functions where the unrestricted auction yields more revenue than any auction with an effective bid cap. The intuition is that more convexity leads to more aggressive bidding by agents with high processing costs. The revenue loss resulting from limiting these bidders cannot be anymore compensated by the low-type bidders' willingness to increase their bids up to the bid cap.

Example 5 Assume $n = 2$, $C(t) = e^t$ and $t_i \gg U[0; 1]$: The following graph shows the dependence of a bidder's expected bid on the marginal type θ . Since the marginal bid increases in θ (which, in turn, increases in the bid cap θ) it is optimal not to restrict the auction with an effective bid cap.



Note that in many standard auctions the auctioneer can increase her revenue by rationing the auctioned good through the imposition of a reserve price. Apart from the fact that this has a negative effect on overall efficiency, and hence is not desirable if efficiency is the auctioneer's concern, it might be difficult to commit to an exclusion of bidders ex-post. In contrast, we analyzed here the effect of an instrument (bid caps) that ensure that all agents are served.

5 Optimal Mechanisms

In the previous section we showed that, for the case of convex cost functions, it may be advantageous (both for revenue and for efficiency purposes) not to use any private information

about processing times, i.e., both the revenue maximizing and the performance maximizing auction may sometime coincide with a lottery where slots in the queue are allocated at random. But, it is conceivable that scheduling mechanisms other than the above described auctions do perform better. In this section we show that mechanism that do not depend on the available private information do minimize expected waiting costs in an important class of incentive compatible mechanisms.

Definition 6 A direct revelation mechanism $(\bar{A}; p)$ is called ex-post incentive compatible if truthful announcement of processing times constitutes an ex-post equilibrium, i.e. for any realization of types, it is optimal for any agent i to announce his true type given that all other agents also announce their processing times truthfully.

Ex-post equilibria do not depend on the underlying distribution functions (In particular, they do not depend on the assumption of independency of agents' processing times).

A mechanism $(\bar{A}; p)$ that minimizes $P(\bar{A})$ in the class of all ex-post incentive compatible mechanisms is called an optimal ex-post incentive compatible mechanism.

The following result shows that, for the case of two agents, the best allocation that can be implemented in ex-post equilibrium is given by a lottery giving the first position in the queue to both agents with probability $\frac{1}{2}$. In fact, any mechanism that does not depend on reports yields that same allocation. More generally, if the types $t_1; t_2$ are distributed on $[\underline{t}_1; \bar{t}_1]$ and $[\underline{t}_2; \bar{t}_2]$ according to different distribution functions F_1 and F_2 (with densities f_1 and f_2) the best allocation in ex-post equilibrium is given by the ex-ante efficient allocation, i.e., by scheduling first the agent with the shortest expected processing time.

Theorem 7 Assume that C is convex. The optimal ex-post incentive compatible mechanism is to always choose the SEPT schedule based on the ex-ante expected processing costs $\int_{\underline{t}_i}^{\bar{t}_i} C(t_i) f_i(t_i) dt_i$. If the ex-ante expected processing costs are the same for both agents, then any random schedule is optimal.

Proof. Appendix ■

The situation is more complicated if the mechanism designer is not only interested in minimizing expected total waiting costs for the agents, but additionally prefers that this cost stays below a given threshold as often as possible. Such a constraint arises for example if there is an overall budget constraint with costly overdraws.

The following stylized example (which can be approximated in our model) shows that these two goals might not be attainable simultaneously.

Example 8 Assume two agents and the following distribution of processing time:

$$t_1 = \begin{cases} 1 & \text{with probability } \frac{1}{5} \\ 3 & \text{with probability } \frac{4}{5} \end{cases}; \quad t_2 = \begin{cases} 2 & \text{with probability } \frac{1}{5} \\ 5 & \text{with probability } \frac{4}{5} \end{cases}$$

The (convex) cost function is given by $C(t) = t^2$. Note that t_2 stochastically dominates t_1 hence, according to Theorem 7, it is optimal to serve agent 1 first since this minimizes expected processing time. On the other hand, the probability that overall processing time is below 30 is higher if agent 2 is queued first, i.e.

$$\Pr\{C(t_1) + C(t_1 + t_2) \leq 30\} = \frac{4}{25} < \frac{1}{5} = \Pr\{C(t_2) + C(t_1 + t_2) \leq 30\}$$

The following result offers a sufficient condition guaranteeing that serving the agent with lower expected processing time first maximizes the probability that total costs are below a given (arbitrary) threshold among all announcement-independent mechanisms. We use the following definition:

Definition 9 The random variable t_2 is larger than t_1 in the sense of likelihood ratio, $t_2 \succ_{LR} t_1$, if $\frac{f_2(t)}{f_1(t)}$ is (weakly) increasing in t :

Note that the likelihood ratio order implies standard stochastic dominance. In particular, $t_2 \succ_{LR} t_1$ implies that $E(t_2) \geq E(t_1)$:

The following is a well-known property of the likelihood ratio ordering:

Proposition 10 Assume that $t_2 \succ_{LR} t_1$; and let $h(x; y)$ be a real-valued function satisfying $h(t_2; t_1) \geq h(t_1; t_2)$ for all $t_2 \geq t_1$. Then the random variable $h(t_2; t_1)$ stochastically dominates $h(t_1; t_2)$:

Proof. See Ross (1983), Proposition 8.4.2 on page 268.⁵ ■

It is immediate that the above condition is satisfied for the function h that measures total waiting costs, i.e., $h(t_1; t_2) = C(t_1) + C(t_1 + t_2)$: Hence, if $t_2 \succ_{LR} t_1$; it follows that the probability of overall costs being below an arbitrary threshold is higher if agent 1 is served before agent 2 than the other way round. In other words, if in a deterministic scheduling problem with known processing times it is optimal to interchange two jobs, then it is also stochastically optimal to interchange two jobs with random processing times that are likelihood ratio ordered. In conjunction with Theorem 7, these observations yield the following result:

Theorem 11 Assume that $t_2 \succ_{LR} t_1$ and that C is convex. Given an arbitrary constant $K \geq 0$; the optimal ex-post incentive compatible mechanism maximizes the probability that overall costs are below K in the class of all mechanisms that do not depend on announcements.

⁵For an extensive analysis of the uses of stochastic orders in queueing see Chang and Yao (1993) and Shanthikumar and Yao (1991).

6 Conclusion

We have combined a simple queueing problem with an incentive problem arising if impatient agents are privately informed about the processing time needed to complete their respective jobs. We have shown that auction-based queue disciplines are efficient if the delay cost function is concave, and anti-efficient if it is convex. For the later case, the auction's performance can be enhanced by imposing bid caps. Finally, for the case of convex cost functions, we have shown that the best performance is attained by mechanisms that do not attempt to condition on the private information. These mechanisms take into account ex-ante available information.

The model analyzed here is very simple, and the vast queueing literature has considered much more complicated models with random arrivals, multiple servers, preemptive service, multi-stage service, etc...In principle, the performance of various pricing mechanisms similar to those studied here can be applied to such models. We think that the combination of queueing and incentive models constitutes a fruitful avenue and that such studies will have many significant real-life applications.

7 Appendix

We use the following abbreviation: $dF(t_{l+1}; \dots; t_n) := \int_{t_{l+1}}^{\infty} \dots \int_{t_n}^{\infty} f(t_{l+1}) \dots f(t_n) dt_{l+1} \dots dt_n$:

Proof of Theorem 2: The proof is performed for the case of a convex cost function. The case of a concave cost functions follows along the same line.

Given that other bidders bid according to a strictly increasing bidding function b if their type is smaller θ and bid θ otherwise, the interim expected utility of a bidder with type t_1 who bids according to b as if he were of type $\theta < \theta$ is given by:

$$U(t_1; b) = V_i \prod_{l=1}^{n_i-1} \frac{\mu_{n_i-1}}{\mu_{l-1}} \frac{F^{l-1}(\theta)}{F^{l-1}(t_1)} \frac{z_{\bar{t}}}{|z_{\bar{t}} - \theta|} \frac{z_{\bar{t}}}{|z_{\bar{t}} - t_1|} \frac{\tilde{A}}{C} \prod_{j=l+1}^n \int_{t_j}^{\infty} dF(t_{l+1}; \dots; t_n) \frac{1}{b(\theta)}$$

If a bidder with type t_1 bids θ , his interim expected utility is given by (we define $\int_s^s f(s) = 0$):

$$U(t_1; \theta) = V_i \prod_{l=1}^{n_i-1} \frac{\mu_{n_i-1}}{\mu_{l-1}} \frac{F^{l-1}(\theta)}{F^{l-1}(t_1)} \frac{z_{\bar{t}}}{|z_{\bar{t}} - \theta|} \frac{z_{\bar{t}}}{|z_{\bar{t}} - t_1|} \frac{1}{|z_{\bar{t}} - \theta|^{n_i-1}} \prod_{k=0}^{n_i-1} \frac{\tilde{A}}{C} \prod_{j=l+1}^n \int_{t_j}^{\infty} dF(t_{l+1}; \dots; t_n) \frac{1}{\theta}$$

Taking the derivative of $U^i(t_1; b^i)$ with respect to b^i yields

$$\begin{aligned} \frac{\partial U^i(t_1; b^i)}{\partial b^i} &= \sum_{l=2}^n \mu_{l-1}^{n_i} (l-1) F^{l-2}(b^i) \frac{z_{\bar{t}}}{|z_{\bar{t}}|} \frac{\tilde{A}}{C_{t_1+b^i}} \sum_{j=l+1}^n t_j dF(t_{l+1}; \dots; t_n) \\ &\quad + \sum_{l=1}^n \mu_{l-1}^{n_i} (n_i-l) F^{l-1}(b^i) \frac{z_{\bar{t}}}{|z_{\bar{t}}|} \frac{\tilde{A}}{C_{t_1+b^i}} \sum_{j=l+2}^n t_j dF(t_{l+2}; \dots; t_n) \\ &\quad + i \frac{d}{db^i} b^i \\ &= (n_i-1) \sum_{l=1}^n \mu_{l-1}^{n_i} F^{l-1}(b^i) \frac{z_{\bar{t}}}{|z_{\bar{t}}|} \frac{\tilde{A}}{C_{t_1+b^i}} \sum_{j=l+2}^n t_j dF(t_{l+2}; \dots; t_n) + i \frac{d}{db^i} b^i \end{aligned}$$

The first order condition $\frac{\partial U^i(t_1; b^i)}{\partial b^i} \Big|_{b^i=t_1} = 0$ and the condition $b^i(t) = 0$ determines the equilibrium bidding function⁶ $b^i(t_1)$ for $t_1 < \bar{e}$. This function coincides with $b_{\text{vex}}^i(t_1)$ on that interval. Given \bar{e} , the marginal type \bar{e} is determined by the condition $U^i(\bar{e}, b^i(\bar{e})) = U^i(\bar{e}, \bar{e})$; i.e. by

$$\bar{e} = b^i(\bar{e}) + \sum_{l=1}^n \mu_{l-1}^{n_i} F^{l-1}(\bar{e}) \frac{z_{\bar{t}}}{|z_{\bar{t}}|} \frac{\tilde{A}}{C_{\bar{e}+}} \sum_{j=l+1}^n t_j + \frac{1}{n_i-l+1} \sum_{k=0}^l \frac{\tilde{A}}{C_{\bar{e}+}} \sum_{j=l+1}^n t_j dF(t_{l+1}; \dots; t_n); \quad (6)$$

For $\bar{e} = b^i(\bar{t})$ we get $\bar{e} = \bar{t}$ whereas for

$$\bar{e}_M = \frac{z_{\bar{t}}}{|z_{\bar{t}}|} \frac{\tilde{A}}{C_{\bar{t}+}} \sum_{j=2}^n t_j + \frac{1}{n_i-l+1} \sum_{k=0}^l \frac{\tilde{A}}{C_{\bar{t}+}} \sum_{j=2}^n t_j dF(t_{l+1}; \dots; t_n) \quad (7)$$

we have $\bar{e} = \bar{t}$: The derivative with respect to \bar{e} of the r.h.s. of (6) is strictly positive. This

⁶It can easily be checked that the second order condition is also satisfied for $t < \bar{e}$.

shows that θ is uniquely defined and that a higher bid cap θ yields a higher marginal type θ :

$$\begin{aligned}
 \frac{d\theta}{d\theta} &= \frac{d b_i(\theta)}{d\theta} + \sum_{l=1}^{n_i-1} \frac{\mu_{n_i-1}^{l+1}}{l!} F^{l+1}(\theta) f_i(\theta) \\
 &\quad \frac{z_i}{z_i} \frac{z_i}{z_i} \frac{\tilde{A}}{\tilde{A}} \sum_{j=l+1}^{n_i-1} C_{\theta} \theta^{j-1} \frac{1}{n_i-l+1} \sum_{k=0}^{n_i-l} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+1}; \dots; t_n)}{d\theta} \\
 &\quad \sum_{l=1}^{n_i-1} \frac{\mu_{n_i-1}^{l+1}}{l!} F^{l+1}(\theta) f_i(\theta) \frac{z_i}{z_i} \frac{\tilde{A}}{\tilde{A}} \sum_{j=l+2}^{n_i-1} C_{2\theta} \theta^{j-1} \frac{dF(t_{l+2}; \dots; t_n)}{d\theta} \\
 &\quad \sum_{k=0}^{n_i-1} \frac{\tilde{A}}{\tilde{A}} \frac{n_i-l-k}{n_i-l+1} C_{2\theta} \theta^{k-1} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+2}; \dots; t_n)}{d\theta} + \frac{k+1}{n_i-l+1} C_{\theta} \theta^{k-1} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+2}; \dots; t_n)}{d\theta} \\
 &\quad + \sum_{l=1}^{n_i-1} \frac{\mu_{n_i-1}^{l+1}}{l!} F^{l+1}(\theta) f_i(\theta) \frac{z_i}{z_i} \frac{\tilde{A}}{\tilde{A}} \sum_{j=l+1}^{n_i-1} C_{\theta} \theta^{j-1} \frac{1}{n_i-l+1} \sum_{k=0}^{n_i-l} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+1}; \dots; t_n)}{d\theta} \\
 &= \frac{d b_i(\theta)}{d\theta} + (n_i-1) \sum_{l=1}^{n_i-1} \frac{\mu_{n_i-1}^{l+1}}{l!} F^{l+1}(\theta) f_i(\theta) \\
 &\quad \frac{z_i}{z_i} \frac{z_i}{z_i} \frac{\tilde{A}}{\tilde{A}} \sum_{j=l+2}^{n_i-1} C_{\theta} \theta^{j-1} \frac{1}{n_i-l+1} \sum_{k=0}^{n_i-l} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+2}; \dots; t_n)}{d\theta} \\
 &\quad + \frac{1}{n_i-l} \sum_{k=0}^{n_i-1} \frac{\tilde{A}}{\tilde{A}} \frac{n_i-l-k}{n_i-l+1} C_{2\theta} \theta^{k-1} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+2}; \dots; t_n)}{d\theta} + \frac{k+1}{n_i-l+1} C_{\theta} \theta^{k-1} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+2}; \dots; t_n)}{d\theta} \\
 &\quad + (n_i-1) F^{n_i-2}(\theta) f_i(\theta) \frac{1}{2} C_{2\theta} \theta^{n_i-2} \frac{dF(t_{l+1}; \dots; t_n)}{d\theta} + \sum_{l=1}^{n_i-1} \frac{\mu_{n_i-1}^{l+1}}{l!} F^{l+1}(\theta) f_i(\theta) \\
 &\quad \frac{z_i}{z_i} \frac{z_i}{z_i} \frac{\tilde{A}}{\tilde{A}} \sum_{j=l+1}^{n_i-1} C_{\theta} \theta^{j-1} \frac{1}{n_i-l+1} \sum_{k=0}^{n_i-l} \frac{\tilde{A}}{\tilde{A}} \theta^k \frac{dF(t_{l+1}; \dots; t_n)}{d\theta}
 \end{aligned}$$

$$\begin{aligned}
&= (n_i - 1) \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) \\
&\quad + (n_i - 1) F^{n_i-2}(e) \frac{1}{2} C(e) \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1}) \\
&\quad + \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1})
\end{aligned}$$

Obviously, it is not optimal to deviate to a bid $b < e$ if $t_1 > e$ or to a bid $b > e$ if $t_1 < e$. Q.E.D.

Proof of Theorem 4: It suffices to show that for the linear cost function $C(x) = x$ imposing a bid cap θ_M leads to a strictly higher revenue than any other bid cap. Given a bid cap θ and given the uniquely defined corresponding marginal type e , the auctioneer's expected revenue per agent is:

$$R(\theta) := \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1}) + \int_{\bar{z}_i}^{\theta} b_{\text{vex}}(x) f(x) dx$$

We show below that $\frac{d}{d\theta} R(\theta) = \frac{d}{d\theta} \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1}) + b_{\text{vex}}(e) f(e) < 0$ for $e < \bar{z}_i$: Thus, the optimal marginal type is e ; and, accordingly, the optimal bid cap is θ_M :

For $C(x) = x$ we have:

$$\begin{aligned}
\frac{d}{d\theta} R(\theta) &= \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1}) \\
&\quad + (n_i - 1) F^{n_i-2}(e) \frac{1}{2} C(e) \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1}) \\
&\quad + \int_{z_i}^{\bar{z}_i} \frac{1}{n_i} \sum_{k=0}^{n_i-1} \frac{n_i - l_i - k}{n_i - l_i - 1} C(e + t_j) dF(t_{l+1})
\end{aligned}$$

$$\begin{aligned}
&= \prod_{i=1}^{n_i-1} \int_{t_{i-1}}^{t_i} \frac{1}{n_i} e^{-\frac{t_i - t_{i-1}}{n_i}} \prod_{j=i+1}^n \int_{t_{j-1}}^{t_j} \frac{1}{n_j} e^{-\frac{t_j - t_{j-1}}{n_j}} dF(t_{i+1}; \dots; t_n) \\
&= \prod_{i=1}^{n_i-1} \int_{t_{i-1}}^{t_i} \frac{1}{n_i} e^{-\frac{t_i - t_{i-1}}{n_i}} \prod_{j=i+1}^n \int_{t_{j-1}}^{t_j} \frac{1}{n_j} e^{-\frac{t_j - t_{j-1}}{n_j}} dF(t_{i+1}; \dots; t_n) \\
&= \prod_{i=1}^{n_i-1} \int_{t_{i-1}}^{t_i} \frac{1}{n_i} e^{-\frac{t_i - t_{i-1}}{n_i}} \prod_{j=i+1}^n \int_{t_{j-1}}^{t_j} \frac{1}{n_j} e^{-\frac{t_j - t_{j-1}}{n_j}} dF(t_{i+1}; \dots; t_n) \\
&= 0.
\end{aligned}$$

Q.E.D.

Proof of Theorem 7: Since the suggested mechanisms do not depend on announcements, truth-telling is an ex-post equilibrium. To simplify notation, let $k_i^i(b_1; b_2)$ denote the probability that agent i is served first, given announcements $b_1; b_2$. Obviously we must have $k_2 = 1 - k_1$. For an allocation $(k_1; k_2)$ to be implementable in ex-post equilibrium we must have for all t_1, t_2 :

$$u_i(t_1; t_1; t_1) - u_i(t_1; t_2; t_1) \geq u_i(t_2; t_1; t_1) - u_i(t_2; t_2; t_1) \quad \text{for all } t_1, t_2$$

For $i = 1$ this yields for all t_2 :

$$\begin{aligned}
&C(t_1) - C(t_1; t_2) \geq k_1 C(t_2; t_1) + (1 - k_1) C(t_2; t_2) \\
&C(t_1) - C(t_1; t_2) \geq k_1 C(t_2; t_1) + (1 - k_1) C(t_2; t_2)
\end{aligned}$$

This is equivalent to:

$$\begin{aligned}
&C(t_2; t_1) - C(t_1) \geq C(t_2; t_2) - C(t_1; t_2) \\
&C(t_2; t_1) - C(t_1) \geq C(t_2; t_2) - C(t_1; t_2)
\end{aligned}$$

Since C is convex, the above condition is equivalent to $k_1 C(t_2; t_2) \geq k_1 C(t_1; t_2)$ for all t_2 . This last condition implies in turn that ex-post implementability requires that

$$\int_{t_2}^{\infty} k_1(t_1; t_2) f_2(t_2) dt_2 \quad \text{is increasing in } t_1 \tag{8}$$

Equivalently, by looking at $i = 2$ we get the requirement that

$$\int_{\underline{t}_1}^{\bar{t}_1} k_1(t_1; t_2) f_1(t_1) dt_1 \text{ is decreasing in } t_2: \quad (9)$$

We are now looking for the solution to the following problem:

$$\min_{(k_1; k_2)} \int_{\underline{t}_1}^{\bar{t}_1} \int_{\underline{t}_2}^{\bar{t}_2} [C(t_1) k_1(t_1; t_2) + C(t_2) k_2(t_1; t_2)] f_1(t_1) f_2(t_2) dt_1 dt_2 \quad (10)$$

subject to incentive compatibility constraints:

Since $k_2 = 1 - k_1$; and since the cost $C(t_1 + t_2)$ is incurred for sure in any allocation, the above problem becomes:

$$\min_{k_1} \int_{\underline{t}_1}^{\bar{t}_1} \int_{\underline{t}_2}^{\bar{t}_2} (C(t_1) - C(t_2)) k_1(t_1; t_2) f_1(t_1) f_2(t_2) dt_1 dt_2 \quad (11)$$

s.t: (8); (9):

We have that

$$\begin{aligned} & \int_{\underline{t}_1}^{\bar{t}_1} \int_{\underline{t}_2}^{\bar{t}_2} (C(t_1) - C(t_2)) k_1(t_1; t_2) f_1(t_1) f_2(t_2) dt_1 dt_2 \\ = & \int_{\underline{t}_1}^{\bar{t}_1} C(t_1) \int_{\underline{t}_2}^{\bar{t}_2} k_1(t_1; t_2) f_2(t_2) dt_2 f_1(t_1) dt_1 \\ & - \int_{\underline{t}_2}^{\bar{t}_2} C(t_2) \int_{\underline{t}_1}^{\bar{t}_1} k_1(t_1; t_2) f_1(t_1) dt_1 f_2(t_2) dt_2: \end{aligned}$$

For the solution k_1 of (11) define:

$$\int_{\underline{t}_1}^{\bar{t}_1} \int_{\underline{t}_2}^{\bar{t}_2} k_1(t_1; t_2) f_1(t_1) f_2(t_2) dt_1 dt_2 := K \in [0; 1]:$$

Since C is increasing and because of (8) and (9) we obtain that

$$\int_{\underline{t}_1}^{\bar{t}_1} C(t_1) \int_{\underline{t}_2}^{\bar{t}_2} k_1(t_1; t_2) f_2(t_2) dt_2 f_1(t_1) dt_1$$

is minimized if $\int_{\underline{t}_2}^{\bar{t}_2} k_1(t_1; t_2) f_2(t_2) dt_2 = K$ and that

$$\int_{\underline{t}_2}^{\bar{t}_2} C(t_2) \int_{\underline{t}_1}^{\bar{t}_1} k_1(t_1; t_2) f_1(t_1) dt_1 f_2(t_2) dt_2$$

is maximized if $\int_{\underline{t}_1}^{\bar{t}_1} k_1(t_1; t_2) f_1(t_1) dt_1 = K$. Hence, for the solution of (11), we must have:

$$\begin{aligned} & \int_{\underline{t}_1}^{\bar{t}_1} \int_{\underline{t}_2}^{\bar{t}_2} (C(t_1) - C(t_2)) k_1(t_1; t_2) f_1(t_1) f_2(t_2) dt_1 dt_2 \\ = & K \int_{\underline{t}_1}^{\bar{t}_1} C(t_1) f_1(t_1) dt_1 - \int_{\underline{t}_2}^{\bar{t}_2} C(t_2) f_2(t_2) dt_2 \quad ! \end{aligned}$$

This last expression is minimized at:

$$\begin{aligned}
K = 0 & \quad \text{if } \int_{t_1}^{\bar{t}_1} C(t_1) f_1(t_1) dt_1 \leq \int_{t_2}^{\bar{t}_2} C(t_2) f_2(t_2) dt_2 \\
K = 1 & \quad \text{if } \int_{t_1}^{\bar{t}_1} C(t_1) f_1(t_1) dt_1 \cdot \int_{t_2}^{\bar{t}_2} C(t_2) f_2(t_2) dt_2 \\
\text{any } K \in [0; 1] & \quad \text{if } \int_{t_1}^{\bar{t}_1} C(t_1) f_1(t_1) dt_1 = \int_{t_2}^{\bar{t}_2} C(t_2) f_2(t_2) dt_2
\end{aligned}$$

The mechanism suggested in the statement of the theorem obviously fulfills these requirements. Q.E.D.

8 References

Chang, C. and Yao, D. (1993): "Rearrangement, majorization and stochastic scheduling", *Mathematics of Operations Research* 18, 658-684

Dasgupta, P., and Maskin, E. (2000): "Efficient Auctions", *Quarterly Journal of Economics* 115, 341-388.

Gavious, A., Moldovanu, B., and Sela, A. (2002) "Bid Costs and Endogenous Bid Caps," *Rand Journal of Economics* 33, 709-722.

Gross, D. and Harris, C. (1998): *Fundamentals of Queuing Theory*, New York: John Wiley&Sons

Hain, R. and Mitra, M. (2002): "Simple Sequencing Problems with Interdependent Costs" mimeo, University of Bonn.

Jehiel, P., and Moldovanu, B. (2001) "Efficient Design with Interdependent Valuations", *Econometrica* 69, 1237-1259.

Mitra, M. (2001): "Mechanism Design in Queueing Problems", *Economic Theory* 17, 277-305.

Ross, S. (1983): *Stochastic Processes*, New York: John Wiley&Sons

Shanthikumar, J. and Yao, D. (1991): "Bivariate Characterizations of Some Stochastic Order Relations", *Advances in Applied Probability* 23, 642-659.

Wellman, P., Walsh, E., Wurman, P. and MacKie-Mason, J. (2001): "Auction Protocols for Decentralized Scheduling", *Games and Economic Behavior* 35, 271-303.