

Calendar mechanisms*

Toomas Hinnosaar[†]

November 12, 2011

Abstract

I study a repeated mechanism design problem where a revenue-maximizing monopolist sells a fixed number of service slots to randomly arriving buyers with private values and increasing exit rates. In addition to characterizing the fully optimal mechanism, I study the optimal mechanisms in two restricted classes. First, the pure calendar mechanism, where the seller allocates future service dates instead of general promises. The unique optimal pure calendar mechanism is characterized in terms of the opportunity costs of allocating additional service slots. Second, I analyze the waiting list mechanism, where promises of delayed service can depend on future arrivals, but the seller cannot discriminate among buyers who are offered the same position in the waiting list. Both the waiting list and the fully optimal mechanism are implemented by non-standard auctions with a scoring rule where the distance between buyers' bids affects the allocation. A novel property of these auctions is that for buyers it is better to win by a close margin and it is worse to lose by a close margin. Finally, I model partial commitment power as a penalty that the seller has to pay when forfeiting a promise. All the results are given for general partial commitment and therefore include full commitment and no commitment as special cases.

*I am grateful to my advisor Jeff Ely for continued guidance and support, Alessandro Pavan and Asher Wolinsky for many helpful discussions, and Eddie Dekel, Dan Garrett, Marit Hinnosaar, as well as seminar participants at Northwestern University for their feedback. The financial support from the Center for Economic Theory of Northwestern University is acknowledged.

[†]E-mail: t-hinnosaar@northwestern.edu. Address: Department of Economics, Northwestern University, Evanston IL, 60208, USA.

1 Introduction

I study a mechanism design problem where a monopolist sells services to randomly arriving privately informed forward-looking customers. Typical applications include Internet ad-auctions, shuttle or concert ticket sales, restaurant reservations or the assignment of service slots by service providers ranging from hair salons to hospitals. The common features in all these examples are that 1) goods are perishable—once the corresponding date is passed, the good disappears; 2) the situation is repeated; and 3) demand is random, with customers willing to wait for some time but not for very long. In addition, in real life buyers generally interact with the seller only once: for example buy a ticket for the closest available date. In most cases continuous involvement would be too costly. Furthermore, if promises are for the too distant future, buyers may find alternative ways to solve their problems. Since the seller makes promises about future service, the promises are credible only if the seller has sufficient commitment power.

In the model there is a monopolistic seller who has a fixed number of identical service slots every period. Buyers arrive randomly in time and have independent private values. Buyers discount the value of service, but not necessarily at the same rate as the seller does. The seller's goal is to maximize the expected discounted revenue¹ by choosing the optimal mechanism from a potentially restricted class of mechanisms. In addition to the fully optimal mechanism, I study mechanisms where the seller has only a limited set of allocation rules available.

First, the pure calendar mechanism where the seller can give only promises of future service that are fulfilled with certainty. In other words, the seller has a calendar according to which the service is supplied. In each period, the seller assigns available dates from the calendar to new buyers. Selling tickets and making appointments are examples of this kind of mechanism. Second, with the waiting list mechanism the seller can give promises that are conditional on future arrivals but cannot personalize the promises to buyers' types. This restriction operates as if the seller keeps a list of buyers who were promised delayed service and serves them according to some rule. Customers in different positions in the list may be served differently, but only the position in the list—and not the buyers type—will determine the allocation she gets. In some sense this is a natural relaxation of the pure calendar mechanism assumption—the seller may want to schedule appointments but with an understanding that when a very valuable buyer arrives in the future, some of the appointments will be canceled. Finally, in the fully optimal mechanism, the seller is free to choose any kind of promises about the future service. In contrast to the waiting list the promises will be personalized to each buyer according to their valuations.

To study commitment power we assume that the seller is able to commit only to promises of delayed services that are not too tempting to break in the future. This is modeled using a partial commitment constant that can be interpreted as a cost or fine—if the seller does not serve a customer who was promised service, she must pay a penalty. All the results presented in this paper are for general partial commitment constant; therefore full commitment and no

¹As it is often the case in mechanism design literature, maximizing efficiency is not fundamentally different. In particular, all the results presented in this paper can be interpreted as efficient mechanisms when virtual values in the expressions are replaced by true values.

commitment come as special cases.

This paper takes a different approach from the related literature by assuming that each buyer interacts with the seller only once. This assumption is motivated by applications, where buyers do not interact with the seller for a very long time. This is presumably because it would be too costly for them to participate in the mechanism more than necessary or that they prefer to have the information about the future service as soon as possible (for example to purchase complementary services, such as to reserve hotel rooms at the same time as buying plane tickets). Note that if the class of mechanisms that the seller can use is unrestricted, the one-time-interaction assumption is without loss. For any multiple-interactions mechanisms, there exist a corresponding one-time-interaction mechanism, where the seller gives each buyer a complete history-dependent allocation rule as a promise. The restrictions to the mechanism designer, including pure calendar mechanism and partial commitment, are defined relative to the one-time-interaction assumption.

There are two special cases of this model, where the fully optimal mechanism is relatively simple to characterize. First, if the number of objects is large enough to serve all potential customers at their arrival, the optimal mechanism would be a constant posted-price mechanism. This is a version of Stokey's (1979) no pure price-discrimination result. Here is the reason. From the buyers' perspective, being offered delayed service is equivalent to getting fractional allocation. Without the need for rationing and without taking the continuation values into account, the optimal solution would use only corner solutions² (Myerson, 1981). Offering delayed service can only decrease the continuation values, so the dynamic feasibility considerations would not change the conclusion. This result does not hold when at least in some periods there are fewer objects than new buyers, so that rationing is needed.

Second, when the buyers discount the future service at a constant rate and the seller has the option to give promises of arbitrary length, it is optimal to give the current service to the buyers who have the highest values³ among all buyers who have arrived but not yet received service. This result follows from Said (2008). The reason is that two buyers who have arrived in different periods look at the future the same way, which makes it easy to compare buyers—only their valuations matter. The result fails, however, if buyers do not discount the future at a constant rate or, equivalently, if they exit from the model at an increasing rate. In these cases the seller needs to treat buyers from different periods differently. The buyers' types would be multi-dimensional, characterized by both the valuation and the exit process.

This paper concentrates on the remaining case, where rationing is needed and buyers' continuation values are dependent on their arrival time. There are many applications that fall in this category. First, there could be practical or legal reasons why buyers interact with the seller once and the seller cannot give promises about service in the too distant future. Many types of selling tickets and making appointments fall into this situation. Second, it may be because of buyers' preferences—it is possible that buyers accept service if it is offered in the near future,

²In particular give a buyer the instant service if her virtual surplus is positive and no service otherwise.

³Highest positive virtual surplus if the goal is to maximize revenue.

but they cannot wait for a long time. Getting a haircut, appointments for emergency services, and transportation seem to satisfy this description. Finally, it could be because of competition that we do not model explicitly. It is possible that if buyers expect to get the service soon they would wait, but if the delay is long they would look for alternatives.

We will focus attention on the extreme case where the characterizations given above fail. In particular, we assume that buyers interact with the seller only at their arrival, and the seller cannot give promises for more than one period ahead. In this case two buyers who have arrived in two consecutive periods are different, since the buyer who has already waited one period must be served now or will not be served, whereas a new buyer can be offered either current or delayed service. The heterogeneity of customers is now two-dimensional—not only their values but also their exit rates differ. As noted, there are at least three mathematically equivalent interpretations of this assumption. First, the seller simply cannot contract service more than one period ahead. Second, buyers have time preferences such that they get a positive value from current and next-period consumption, but they do not get any value from the service after that. Third, buyers exit the model two periods after their arrival. As we would expect, the optimal mechanism in this framework offers delayed service with strictly positive probability and treats buyers from different periods differently.

We show that the fully optimal mechanism is implemented by a non-standard auction with scoring rule, which favors buyers who are waiting and in which the allocation that new buyers receive depends not only on the order of their values, but also on the differences in their values. In particular we show some types of buyers strictly prefer to win in the auction by a small margin, because this is the only way they receive instant service. Similarly, some types of buyers prefer to lose by a large margin rather than a small margin, since in this case they would be refused service rather than be offered delayed service.

Characterizing the fully optimal mechanism is one contribution of the paper. In addition to that I derive the optimal in the two classes of restricted mechanisms described above. The optimal pure calendar mechanism is characterized by a unique vector of opportunity costs of the corresponding service slots. It can be implemented by a simple two-stage mechanism. The qualitative properties of the optimal waiting list mechanism are similar to the fully optimal mechanism, but the delayed service contracts cannot be personalized to the buyers receiving the delayed service. The optimal contracts are therefore designed for the “average” buyer who will be offered delayed service. The optimal waiting list will have non-trivial dynamics even if the arrival process is constant. Since the value of an average buyer to whom the seller expects to assign each contract depends on previous promises, it is possible that the optimal contracts change every period.

Finally, I characterize all three classes of mechanisms under the general partial commitment assumption, which shows how the optimal mechanism changes with commitment power. In particular there is a sufficient and, in most cases, necessary level of commitment that allows the seller to use the optimal mechanism in all three classes. If the commitment power is relatively low, the optimal calendar mechanism is static, whereas the revenue from the optimal waiting list

mechanism and the unrestricted mechanism are strictly increasing in commitment power.

As with most of the mechanism design literature, this modeling approach is similar to the seminal papers by Myerson (1981) and Riley and Samuelson (1981). In particular, Myerson's optimal mechanism is the optimal static mechanism in the model analyzed in this paper. Static mechanism belongs to all three classes of mechanisms that we consider, and it is therefore a benchmark against which to compare all the dynamic mechanisms. We show, however, that it will not generally be optimal.

This paper belongs to a growing body of literature on dynamic mechanism design. A related branch of literature studies the sale of durable goods, where the seller has a fixed number of durable goods and a deadline by which the goods must be sold. Gershkov and Moldovanu (2009) characterize the optimal online mechanism where buyers must be assigned an object at their arrival, and Board and Skrzypacz (2010) characterize the optimal mechanism when the buyers are patient. Pai and Vohra (2008) consider a model where buyers' unobservable types include their arrival and departure times and give partial characterization of the optimal mechanism. In these papers the most important trade-offs are between extraction of rents and the option value—giving away an object before the deadline means that the object cannot be assigned to potentially higher-valuing buyers who arrive later. The option value decreases in time, however, since it becomes less likely that such buyers will arrive. In contrast to these papers, our model is designed to analyze the situation where the seller has a new set of goods every period, and one set of goods perishes every period. This situation creates a different set of trade-offs.

Bergemann and Välimäki (2010) and Pavan, Segal, and Toikka (2011) give results for general classes of dynamic mechanisms, where both adverse selection and moral hazard issues could be studied. Pavan, Segal, and Toikka give the envelope formula for a large class of dynamic mechanism design questions and show how to compute transfers. One application of their general result is the sale of a durable good to buyers with changing types. Similarly to our results they show that the optimal mechanism is implementable with a non-standard auction with scoring rule, where the differences between bids affect the allocation. However, in their model, winning by a large margin is good and losing by a large margin bad, which is the opposite result from the implications of our model. Bergemann and Välimäki's main result is the characterization of transfers that implement any efficient mechanism.

All the papers noted above study the fully optimal mechanism. The pure calendar mechanism and waiting list mechanism are new to this paper. The particular way we model partial commitment is also new, but analogous to renegotiation-proofness models—for example Hart and Tirole (1988). In this literature the mechanism-designer cannot offer a contract if it is known that it will be mutually beneficial to renegotiate this contract in the future. In our model buyers interact with the seller only once, which means that renegotiation is impossible, but the approach we take is the same: the seller with partial commitment power can give only those promises that she will not find it optimal to break in the future. For simplicity, we assume that the cost of commitment is exogenous⁴.

⁴There would be several interesting ways to endogenize it, including reputation or direct compensation to

Section 2 shows that most of the trade-offs of the problem can be studied in a simple illustrative example, with only one object and two new buyers every period. In particular, Section 2.1 characterizes the optimal pure calendar mechanism, Section 2.2 the optimal waiting list mechanism, and Section 2.3 the fully optimal mechanism. Section 2.4 introduces the partial commitment assumption, and shows how all the previous results generalize for partial commitment. Section 3 generalizes the results for an arbitrary number of service slots and random arrivals. Section 4 concludes and discusses the limitations of the model. Proofs are in Appendix A.

2 Illustrative example

In this section we discuss an example of the model with just one object and two new buyers each period. The results in this section are special cases of the general results in Section 3, and therefore formal statements and full proofs are omitted.

A monopolistic seller has one service slot each period $t \in \{0, 1, 2, \dots\}$. It could be a time slot for providing service, a seat, or a perishable good produced on the same day. The seller maximizes expected discounted revenue, with discount factor $\delta \in (0, 1)$.

Each period two new buyers arrive. Each buyer i has unit demand for the service, with independent private value $v_i \in [0, 1]$. Value v_i is a random variable with cumulative density function F and probability density function $f(v_i) > 0$ for all $v_i \in [0, 1]$. Function $w(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$ denotes the standard static virtual surplus function, which is assumed to be strictly increasing⁵. To simplify the notation, we rename the buyers so that buyer 2 is the one with higher value and buyer 1 with lower value. Denote the CDF⁶ of buyer i by F_i and corresponding PDF⁷ by f_i . Note that we are not following this convention in the figures, where we allow v_1 and v_2 to be arbitrary.

Buyer i with type v_i who is served s_i periods after her arrival and pays p_i gets payoff $\beta^{s_i} v_i - p_i$. In particular, if she is served at the same period she arrives, she gets payoff $v_i - p_i$. Constant $\beta \in (0, 1)$ is the same for all buyers and is referred to as buyers' discount factor. We allow buyers' discount factor β to be different from the seller's discount factor δ .

We concentrate on mechanisms where a buyer who is not served one period after arrival is never served. As argued in the introduction, this may be the case for several reasons—restriction on the mechanism, preferences, or exit process. It captures the situation where in the short run the dates are imperfect substitutes for the buyers, but in the long run the value decreases faster than a constant rate.

Each buyer interacts with the seller only at arrival. This means that each period the two new buyers announce simultaneously their types, and the seller responds by giving both of them either instant service or promises about future service and requests transfers. A promise of future service could be a refusal of service or delayed service conditional on next period arrivals.

buyers.

⁵Monotone hazard rate condition is a sufficient condition.

⁶ $F_2(v_2) = F(v_2)^2$, $F_1(v_1) = F(v_1)[2 - F(v_1)]$.

⁷ $f_2(v_2) = 2F(v_2)f(v_2)$, $f_1(v_1) = 2[1 - F(v_1)]f(v_1)$.

A promise of delayed service is a complete description of the allocation that the buyer gets conditional on the future arrivals. In the illustrative example it is characterized by a set $\hat{D} \subset [0, 1]^2$, such that if buyer i receives promise \hat{D} she will be served in the next period if and only if the new buyers who arrive in the next period announce values $v = (v_1, v_2) \in \hat{D}$. This means that the expected payoff from promise \hat{D} to buyer i with value v_i is $Pr(\hat{D})\beta v_i + Pr(\hat{D}^c)0 - p_i = \beta \hat{d}v_i - p_i$, where $\hat{d} = Pr(\hat{D})$. The set \hat{D} can be interpreted as a contract or a ticket that the seller can give to a buyer one period before the service in exchange for money.

A state captures all relevant details from the history up to a particular period. In particular, it characterizes the promises of delayed service given to buyers in the previous period. In the illustrative example a state is characterized by set $D \subset [0, 1]^2$, which includes all the combinations of types v for which the instant service is unavailable. It is clear that a state is equivalent to a promise given to a buyer in the previous period—the current object is unavailable because it was promised to a buyer in the previous period.⁸ Unless explicitly stated otherwise, we denote the state by D and a contract by \hat{D} .

The goal of the mechanism design problem is to find mechanisms that are Bayesian incentive-compatible and ex-ante individually rational. As it turns out, all mechanisms that we will discuss below are implementable in dominant-strategy incentive-compatible and ex-post individually rational transfers.

As an illustration of possible dynamics, consider the following example. In period 0 the seller has not made any promises, so the initial state $D_0 = \emptyset$. Two new buyers arrive and announce their types. The seller serves one of them instantly (at period 0) and promises the other that she will be served in the next period if and only if all new buyers announce values below \underline{v} . This is described by a contract $\hat{D}_0 = [0, \underline{v}]^2$, which is the state in period 1, $D_1 = \hat{D}_0 = [0, \underline{v}]^2$. Now, in period 1 two new buyers arrive and announce their values $v = (v_1, v_2)$. If $v \in D_1$, the seller has to serve the buyer from the previous period and therefore cannot serve any of the new buyers instantly, whereas if $v \notin D_1$, the instant service is available. When the seller does not offer delayed service to either of the new buyers, we return to the initial state.

One possible mechanism in this framework is the repetition of the static mechanism, where the good is always sold to one of the buyers arriving in the particular period. In other words, the seller never gives promises of delayed service, so the state is $D = \emptyset$ in all periods. By Myerson (1981) the optimal static mechanism assigns instant service to the buyer with the higher value if his virtual valuation is positive, and refuses service to both buyers otherwise. This allocation rule is illustrated by Figure 2.1.

The optimal static mechanism can be implemented by the second price auction with reserve price p^* such that $w(p^*) = p^* - \frac{1-F(p^*)}{f(p^*)} = 0$. This ensures expected revenue $\pi^s = \int_{p^*}^1 w(v_2)dF_2(v_2)$. This mechanism will typically not be optimal in a dynamic setting.

⁸There is a minor abuse of notation here, since it is possible that two of the buyers from the previous period received promises of delayed service, say in sets D_1 and D_2 respectively. At the optimum this happens with zero probability, but is feasible when $D_1 \cap D_2 = \emptyset$ and $D_1 \cup D_2 \subset [0, 1]$. In this case we can simply redefine the state as $D = D_1 \cup D_2$.




Figure 2.1: Optimal static allocation. Shaded area shows when buyer 1 is served instantly.

2.1 Pure calendar mechanism

A pure calendar mechanism is a mechanism where the promise of delayed service is unconditional on next period arrivals. It is like marking dates on a calendar and then providing services according to this calendar, or selling tickets that guarantee the service. This means that a promise of delayed service needs to be fulfilled with certainty, and the seller cannot serve anyone instantly when a buyer from the previous period was promised delayed service.

In the notation introduced above, the promise of delayed service can be only $D = \emptyset$ or $D = [0, 1]^2$, which correspond to probabilities 0 and 1 respectively. There are only two possible states: the instant service is either available, or it is already promised to someone and is therefore unavailable. To shorten the notation, we denote the first state by 0 and the second by 1, both of which correspond to the promises given to buyers in the previous periods.

The pure calendar mechanism assumption also means that the seller cannot personalize the promise of delayed service depending on the type of buyer to whom it is promised. In the next subsections we relax both assumptions. First, the waiting list mechanism relaxes the assumption that the promises have to be unconditional on the next period arrivals; later, the fully optimal mechanism also allows personalized promises.

We will now characterize the optimal allocation rule. There are two cases to consider. First is the interior case, where the seller sometimes promises delayed service. Second is the corner solution, where delay is never promised. The latter case is therefore a static mechanism design problem that we already know how to solve. We start with the interior solution and will later discuss the corner solution.

Suppose that the parameters are such that the seller promises delayed service for some types of buyers. There are two states—the instant service is either available, or it was promised to a

buyer who arrived in the previous period and is therefore unavailable. Denote the revenue from the state when the instant service is available by $\pi(0)$ and when it is unavailable by $\pi(1)$. Clearly, the expected revenue from the former is higher than from the latter. Denote this difference by $\Delta = \pi(0) - \pi(1) > 0$.

The difference in revenues in the two states is an opportunity cost of a promise one period ahead. In particular, if the seller promises delayed service to a current buyer, the discounted loss is $\frac{\delta\Delta}{1-\delta}$. The denominator $1 - \delta$ is needed to compare the difference in continuation values, Δ , to the difference in flow values, which is the revenue extracted from current buyers.

Consider the state where instant service is unavailable. Then the only object that can be allocated is a promise of next-period service. Buyers' values of the object are βv_1 and βv_2 , and allocating this object costs $\frac{\delta\Delta}{1-\delta}$ to the seller. The optimal mechanism therefore assigns the object to the buyer with the higher value, if his discounting-adjusted virtual surplus $\beta w(v_2)$ is greater than the opportunity cost $\frac{\delta\Delta}{1-\delta}$. It is illustrated by Figure 2.2(a). This allocation rule can be implemented by a standard second price auction for the delayed service with reserve price \underline{v} , such that $\beta w(\underline{v}) = \frac{\delta\Delta}{1-\delta}$.

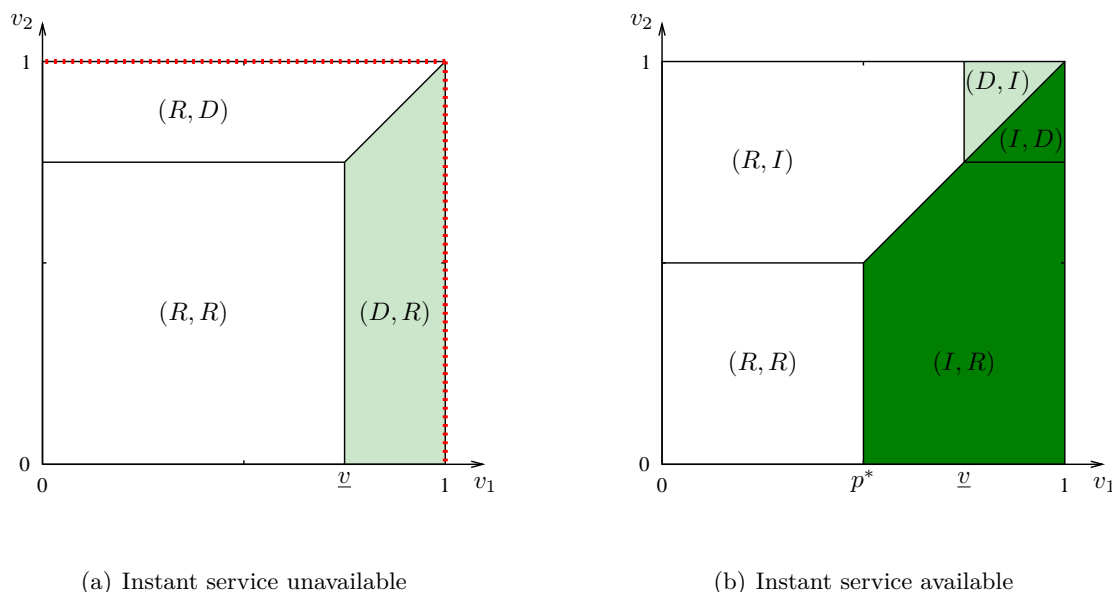


Figure 2.2: Optimal pure calendar mechanism. Intensity of shading color describes the discounted expected quantity for buyer 1, the area below dotted line is the set where previous-period customer is served.

Consider now the state where instant service is available, so that the seller sells two goods to two buyers. The instant service is more valuable to both buyers and costless to allocate, whereas allocating delayed service changes the state exactly as in the case studied above and has the same opportunity cost $\frac{\delta\Delta}{1-\delta}$ as above. This means that the instant service is allocated often, and the delayed service is promised only if both buyers have sufficiently high values.

This is a demand-smoothing argument. When in some period both buyers have high values,

the demand is high. Since it is unlikely that in the next period the values would be as high, it is optimal to extract more revenue from the current high-valuing buyers and lose some revenue from tomorrow's buyers (whose values are uncertain).

In particular, the optimal mechanism allocates the instant service to the buyer with higher value when his virtual surplus is positive. It promises delayed service to the buyer with lower type, when his discounting-adjusted virtual surplus exceeds the opportunity cost. This allocation rule is illustrated by Figure 2.2(b).

This allocation rule is implemented by a simple two-stage auction mechanism. At the first stage, the instant service is allocated at a standard second price auction with reserve price p^* . In the second stage, the loser is offered the next period service at a fixed price $\beta\underline{v}$. When the loser accepts the offer, the winner of the first-round auction gets a discount $\beta(v_l - \underline{v}) > 0$, where v_l is the loser's bid in the auction.

Offering the delayed service to the loser of the second price auction creates surplus for the loser. This means that the winner must be compensated, since losing is not as bad as it would be without the second stage. If types of buyers are $v_2 \geq w_1 \geq \underline{v}$, the high type 2 would prefer winning to losing in the first round if compensation p is high enough to satisfy $v_2 - v_1 + p \geq \beta(v_2 - \underline{v})$. In particular, in the critical case when $v_1 = v_2$, we get the compensation given above, whereas if $v_2 > v_1$, this compensation is sufficient and even gives some extra surplus to the buyer.

Note that the threshold \underline{v} that we used to determine whether or not to allocate the delayed service is the same in both states. The opportunity cost of a promise comes from being in a less restrictive state in the next period. This cost is independent of the type who receives the delayed service.

The discussion above took the opportunity cost $\frac{\delta\Delta}{1-\delta}$ as given, but the difference Δ between the two states is endogenous and yet to be determined. The difference in revenue in the two states can be decomposed to two parts. First, if the instant service is available, it is assigned to the higher-valuing buyer whenever her virtual valuation is positive. This is exactly the same as static optimum and thus gives the static revenue π^s .

The second difference is dynamic and comes from the promises of delayed service. When the instant service is unavailable, the delayed service is promised to the higher buyer whenever his discount-adjusted virtual surplus $\beta w(v_2)$ is higher than the opportunity cost, which is equal to $\beta w(\underline{v})$. In this case the seller gets the difference. This gives expected dynamic flow revenue $\pi_2(\underline{v}) = \int_{\underline{v}}^1 [w(v_2) - w(\underline{v})] dF_2(v_2)$. In the state where the instant service is available, the delayed service is offered under the same condition, but only to the lower-valued buyer. This ensures expected dynamic flow revenue $\pi_1(\underline{v}) = \int_{\underline{v}}^1 [w(v_1) - w(\underline{v})] dF_1(v_1)$, which is smaller than $\pi_2(\underline{v})$.

Using these observations, we can express the opportunity cost $\frac{\delta\Delta}{1-\delta}$ in two ways as a function of \underline{v} to get equation (2.1) that uniquely determines \underline{v} and therefore Δ , assuming that the interior solution exists.

$$\beta w(\underline{v}) = \frac{\delta\Delta}{1-\delta} = \delta (\pi^s + \beta[\pi_1(\underline{v}) - \pi_2(\underline{v})]). \quad (2.1)$$

We can now address the corner solution. To see why equation (2.1) may not have a solution in $\underline{v} \in [0, 1]$, consider the case when the buyers' discount factor β is very low compared to the seller's discount factor δ . With sufficiently small β , the value of delayed service is negligible to the buyers and therefore it is never optimal to offer delayed service. In this case the optimal calendar mechanism is static. We already know that the optimal static mechanism is the second price auction with reserve price p^* , which guarantees revenue π^s .

Suppose now that β is slightly higher, so that the seller is exactly indifferent between offering delayed service to a very small set of potential buyers and using the optimal static mechanism. The incentive to offer delayed service is clearly highest for highest types. In particular, the gain from offering delayed service to a buyer with type $v_i = 1$ would be β , but the cost would be lost revenue π^s next period, with discounted value $\delta\pi^s$. This gives $\beta \geq \delta\pi^s$ as a necessary and sufficient condition for optimal delay. If $\beta/\delta > \pi^s$, it is optimal to offer delayed service with positive probability, whereas when $\beta/\delta < \pi^s$ the optimal pure calendar mechanism is always static.

2.2 Waiting list

Another commonly used method to allocate future service is a waiting list. In this paper we model the waiting list as a restriction to promises, where the seller can choose to delay service conditionally to some buyers but cannot discriminate between two types of buyers who are offered delayed service. In particular, the waiting list assumption means that at any state D , possible promises are such that each buyer gets either instant service, no service, or a promise of delayed service in the next period, when the new arrivals are in some fixed set $\hat{D} \subset [0, 1]^2$. The waiting list assumption is that this set is the same for all buyers who are offered delayed service. Note that we allow \hat{D} be different for all states D or even for all periods.

There are potentially many ways for how a waiting list could be organized. The interpretation of our assumption is that people who wait are put on a list, and once the person is on this list in a certain position, the type is forgotten and the process becomes the same for all types buyers who wait in the same position. In the case when the number of objects is one, the optimal waiting list will include at most one individual, which is a very simple list. In the general model we will analyze in Section 3 the optimal waiting list includes up to m individuals.

We have already studied two cases of waiting lists. At one extreme, when the promise for delayed service is always the empty set, we get the static mechanism. At the other extreme, where the promise of delayed service is unconditional on arrivals, we get the pure calendar mechanism. As we will show below, the pure calendar mechanism can be the optimal waiting list mechanism with certain parameters of the model, whereas static mechanisms will never be optimal. In particular, we show that the optimal contract \hat{D} always includes the set⁹ $[0, p^*]^2$ as a strict subset.

We will show that the optimal waiting list mechanism has some interesting properties. Under

⁹Remember that p^* is defined so that the virtual surplus $w(p^*) = p^* - \frac{1-F(p^*)}{f(p^*)} = 0$.

some combinations of parameters¹⁰, the optimal waiting list mechanism is implemented by non-standard auctions with scoring rule where the services that buyers are offered will depend on how similar their bids are to their opponent's bids. In particular, it is sometimes better to win by a little than by a lot, and it is sometimes better to lose by a lot than by a little. After we have characterized the optimal mechanism we will make these statements precise, but the general intuition is that distortions from serving a buyer from a past period depends on the values of both new buyers. Therefore whether or not the seller uses instant service to serve new buyers or a buyer from the previous period will depend on the values of both new buyers.

Another interesting observation will be the dynamics of waiting list, which is non-trivial. It is possible that the optimal contracts are different each period. The distribution of buyers who receive the promise of delayed service will depend on the contract given to a buyer in the previous period. It is possible that in each period this contract is different from the previous contract and therefore the new optimal contract is also going to be different. As long as the new arrivals are such that one of them receives delayed service, the mechanism could therefore be in a new state every period.

The characterization of the optimal waiting list proceeds in five steps. Initially we assume that the expected discounted revenue at each state D , denoted by $\pi(D)$, is well-defined. We start by characterizing the optimal allocation rule for each state D when the contract offered at this state, denoted by \hat{D} , is fixed. This gives us the allocation rule for each $v = (v_1, v_2)$ at each state D for a given \hat{D} . The second step is to characterize the optimal contract \hat{D} when the probability of service, $\hat{d} = Pr(\hat{D})$, is known. The third step is to find the optimal probability \hat{d} . For this it is convenient to optimize with respect to another variable \bar{b} , which we define later and that describes the maximum distortion from giving a promise. The fourth step is to use the first order condition with respect to \bar{b} and the characterization of optimal allocation to fully characterize the optimal promise and allocation rule or to observe that we get the corner solution. The fifth and final step is to show that the functions $\pi(D)$ exists, which will be verified using the contraction mapping theorem. Once we have characterized the optimal waiting list mechanism, we discuss its implementation and the properties mentioned above.

Before continuing with the characterization we can make a simple observation. Note that the seller never wants to offer service to buyers with negative surplus. This means that a contract \hat{D} such that $Pr(\hat{D}) < F_2(p^*)$ can never be optimal. Instead of this contract, the seller could offer contract $\hat{D}' = [0, p^*]^2$, which ensures strictly higher revenue from all current buyers who are promised this contract and is not restrictive in the next period. Therefore in the following analysis we assume that $\hat{d} \geq F_2(p^*)$ and $[0, p^*]^2 \subset \hat{D}$ for optimal contract \hat{D} .

Step 1: allocation rule. Suppose that the expected revenue function $\pi(D)$ is a known function and denote $\pi(d) = \max_{D: Pr(D) \geq d} \pi(D)$. Let the state be D and suppose that the optimal contract of delayed service in this state is \hat{D} , with corresponding probability $\hat{d} = Pr(\hat{D})$. Suppose that the two buyers who arrived announced values $v = (v_1, v_2)$, where we reorder the values so that

¹⁰For example when buyers' discount factor β equals the seller's discount factor δ .

$v_1 \leq v_2$.

When buyers announced types $v \in D$, the seller has to serve the waiting buyer. This means that the only remaining question is when to allocate the delayed service with contract \hat{D} and when not to do so. The trade-off is analogous to the state where the instant service was unavailable in the pure calendar mechanism. The seller has one object, namely contract \hat{D} , to allocate. The value of the contract \hat{D} to buyers is $\beta \hat{d} v_1$ and $\beta \hat{d} v_2$. The cost of allocating the contract is the opportunity cost that comes from being in a more restrictive state in the next period. In particular, the discounted opportunity cost is $\frac{\delta}{1-\delta}[\pi(\emptyset) - \pi(\hat{D})]$. The optimal allocation rule gives the contract \hat{D} to the buyer with higher valuation if and only if his discount-adjusted virtual valuation is higher than the opportunity cost. Equivalently, there is a threshold \underline{v} so that the higher type receives promise of delayed service if and only if his type is higher than the threshold, $v_2 > \underline{v}$. The threshold \underline{v} must be such that it equalizes the discount-adjusted virtual surplus to the opportunity cost, which gives equation (2.2).

$$\beta \hat{d} w(\underline{v}) = \frac{\delta}{1-\delta}[\pi(\emptyset) - \pi(\hat{D})]. \quad (2.2)$$

When the new buyers have types $v \notin D$, the seller has both instant service and delayed service to allocate to them. The trade-offs are analogous to the state when the instant service is available in the pure calendar mechanism case. Again, instant service is more valuable for both buyers and costless to produce, whereas promising the contract \hat{D} has an opportunity cost $\frac{\delta}{1-\delta}[\pi(\emptyset) - \pi(\hat{D})]$. Therefore the buyer with higher value receives the instant service if and only if her virtual surplus is positive, and the buyer with lower value gets the delayed service if and only if her value is above \underline{v} , which satisfies (2.2).

The new question is compared to the pure calendar mechanism case is the choice of optimal contract \hat{D} . Since buyers' values depend only on probability of getting the service, we can first compute the optimal probability $\hat{d} = Pr(\hat{D})$. This is as if the seller describes only the probability of service \hat{d} instead of the full contract \hat{D} and there is agreement that the promise is going to be fulfilled in the next period. Then in the next period the seller has \hat{d} as the constraint and chooses optimal set \hat{D} according to which serve the waiting customer.

Step 2: fulfilling promises optimally. Suppose that contract \hat{D} was given optimally to previous a period customer. Remember that this implies the probability of service $\hat{d} \geq F_2(p^*)$. From the allocation rule above we can calculate the expected revenue for each v in both of the two cases discussed. Denote the expected revenue from v such that the instant service is unavailable by $\underline{\pi}(v)$ and the expected revenue when the instant service is available by $\bar{\pi}(v)$. The difference between the two, $\bar{\pi}(v) - \underline{\pi}(v)$ can be interpreted as the opportunity cost of adding v to the contract \hat{D} . The expected revenue from probability \hat{d} is therefore

$$\pi(\hat{d}) = \max_{\hat{D}: Pr(\hat{D}) \geq \hat{d}} \left\{ \hat{d} E_v \left[\underline{\pi}(v) | v \in \hat{D} \right] + (1 - \hat{d}) E_v \left[\bar{\pi}(v) | v \in \hat{D}^c \right] \right\}$$

$$= E_v [\bar{\pi}(v)] - \hat{d} \min_{\hat{D}: Pr(\hat{D}) \geq \hat{d}} E_v [\bar{\pi}(v) - \underline{\pi}(v) | v \in \hat{D}].$$

It is clear that in optimum \hat{D} must include points where the distortion from including this profile of promises to \hat{D} , computed as $\bar{\pi}(v) - \underline{\pi}(v)$, is the smallest. It is easy to see that when $v_2 \geq p^*$, distortion $\bar{\pi}(v) - \underline{\pi}(v)$ is strictly increasing in v_2 , and therefore the constraint must be binding.

Then the contract \hat{D} is characterized by the upper bound for the distortion that we denote by \bar{b} . That is, the optimal way to fulfill promised probability \hat{d} is to serve the buyer from the previous period if and only if the distortion from doing so is less than \bar{b} . Therefore we can write¹¹ $\hat{D} = \{v : \bar{\pi}(v) - \underline{\pi}(v) \leq \bar{b}\}$ for \bar{b} such that $Pr(\hat{D}) = \hat{d}$.

Step 3: optimal probabilities of delayed service. The third step is to find the optimal probability \hat{d} that will be promised to the current buyers. We have already discussed how the promise will be fulfilled—in the set $\hat{D} = \{v : \bar{\pi}(v) - \underline{\pi}(v) \leq \bar{b}\}$ for \bar{b} such that $Pr(\hat{D}) = \hat{d}$.

Before optimizing with respect to the probability of delay \hat{d} , we have to find how it affects the expected revenue. An increase in probability \hat{d} has two opposite effects on revenue. First, the seller extracts a higher proportion of discounting-adjusted virtual surplus $\beta w(v_i)$ from each buyer who receives the promise. Second, the contract will be more restrictive and therefore it decreases the continuation revenue whenever delayed service is promised. To study both effects we have to know the probability of assigning the contract to someone and the distribution of buyers' types who receive the promise. We already showed that inside set D the delayed service is promised to the high type and outside the set D it is promised to the low type. In both cases the condition to assign the delayed service is that the corresponding type must be higher than threshold \underline{v} . Let us define probability density function¹² $g_v(\cdot|D)$ as

$$g_v(v_i|D) = \int_0^{v_i} \mathbf{1}[(v_1, v_i) \in D] f_1(v_1) f_2(v_i) dv_1 + \int_{v_i}^1 \mathbf{1}[(v_i, v_2) \notin D] f_1(v_i) f_2(v_2) dv_2.$$

Let the corresponding cumulative density function be $G_v(v_i|D) = \int_0^{v_i} g_v(v'_i|D) dv'_i$.

The distribution G_v that we defined here plays a crucial role in the analysis. Remember that the seller cannot discriminate between buyers who receive the delayed service contract. The distribution G_v characterizes the distribution of the type of buyer who receives the delayed service contract. It is a mixture of f_1 and f_2 , the distributions of high and low type, where the precise mixture depends on the current state D . In particular, if the seller did not give any promises in the previous period, $D = \emptyset$, then it is the same as the distribution of low type, whereas if the seller promised delayed service with certainty (as in pure calendar mechanism), it

¹¹Note that this analysis took the contract offered in state \hat{D} as fixed. In optimum it depends on the state \hat{D} , but as a necessary condition for optimality, the analysis is still valid. Suppose that at the state \hat{D} with the optimal contract \hat{D} the set \hat{D} does not have the form described above. Then we could increase the revenue slightly while keeping the probability of service fixed. Because of the envelope theorem the effect through the change in optimal contract \hat{D} is negligible.

¹²The function g_v is a PDF, since by construction $g_v(v_i|D) \geq 0$ for each v_i and $\int_0^1 g_v(v_i|D) dv_i = Pr(D) + Pr(D^c) = 1$.

is the distribution of high type.

From the allocation rule above we know that the seller assigns the delayed service contract to a corresponding buyer if and only if her type $v_i > \underline{v}$. This happens with probability $1 - G_v(\underline{v}|D)$. The expected revenue extracted from the buyer who receives the promise of delayed service is now $\int_{\underline{v}}^1 \beta \hat{d} w(v_i) dG_v(v_i|D)$.

We can now characterize expected revenue function for a given contract \hat{D} and in particular, probability \hat{d} . It includes two parts. First, the static component of the revenue that is independent of \hat{D} is simply

$$(1 - \delta)(1 - d)E_v[\max\{0, w(v_2)\}|v \notin D] + \delta\pi(\emptyset).$$

Second, the dynamic part, where the seller extracts $\beta \hat{d} w(v_i)$ from each buyer who receives delayed service, but loses $\delta[\pi(\emptyset) - \pi(\hat{D})]$ in continuation revenue. This gives

$$(1 - \delta)\beta \hat{d} \int_{\underline{v}}^1 w(v_i) dG_v(v_i|D) - \delta[1 - G_v(\underline{v}|D)][\pi(\emptyset) - \pi(\hat{D})].$$

The last step before optimization is a change of variables. Instead of working with probability \hat{d} , it is more convenient to work with maximum distortion \bar{b} that characterizes the contract \hat{D} and therefore probability \hat{d} . We will first argue why we can do it and then show how to compute \hat{d} and $[\pi(\emptyset) - \pi(\hat{D})]$ more conveniently using maximum distortion \bar{b} .

Remember that for each $\hat{d} \geq F_2(p^*)$, we have $\hat{D} = \{v : \bar{\pi}(v) - \underline{\pi}(v) \leq \bar{b}\}$ for \bar{b} such that $Pr(\hat{D}) = \hat{d}$. The distortion function $\bar{\pi}(v) - \underline{\pi}(v)$ is strictly monotone in v_2 and the distribution is continuous, so \bar{b} must be strictly increasing with \hat{d} . This means that there is one-to-one relationship between $\hat{d} \in [F_2(p^*), 1]$ and \bar{b} which will be¹³ in $[0, 1 - \delta]$. The value of \bar{b} that corresponds to \hat{d} can be interpreted as the maximum distortion that the seller must incur in the next period to fulfill the promise \hat{d} . Instead of choosing optimal \hat{d} , it is convenient to optimize with respect to \bar{b} .

In particular, let $G_b(\bar{b}) = Pr(\bar{\pi}(v) - \underline{\pi}(v) \leq \bar{b})$. Then $\hat{d} = G_b(\bar{b})$. Moreover, we also know how to compute $\pi(\emptyset) - \pi(\hat{D})$. Compared to non-constrained optimization problem at state \emptyset , state \hat{D} forces the seller to serve the previous buyer at set \hat{D} and therefore lose $\bar{\pi}(v) - \underline{\pi}(v)$ at all these realizations. A convenient way to compute this value is to integrate over G_b . We get that $\pi(\emptyset) - \pi(\hat{D}) = \int_0^{\bar{b}} b dG_b(b)$.

We can therefore rewrite the contract design problem in terms of \bar{b} to obtain the following maximization problem, where $E_{G_v}[w(v_i)|v_i \geq \underline{v}] = \int_{\underline{v}}^1 w(v_i) \frac{dG_v(v_i|D)}{1 - G_v(\underline{v}|D)}$ is the expected surplus of a buyer who receives the delayed service contract.

$$\max_{\bar{b} \in [0, 1 - \delta]} \left\{ (1 - \delta)\beta G_b(\bar{b}) E_{G_v}[w(v_i)|v_i \geq \underline{v}] - \delta \int_0^{\bar{b}} b dG_b(b) \right\},$$

¹³Easy to verify: distortion is minimal if both virtual surpluses are close to 0, which gives lower bound 0; and distortion is maximal if both virtual surpluses are 1, which gives $1 - \delta$.

which gives a first order condition¹⁴

$$(1 - \delta)\beta E_{G_v}[w(v_i)|v_i \geq \underline{v}] - \delta\bar{b} = 0 \quad (2.3)$$

Equation (2.3) balances the trade-off from extracting slightly more revenue from the average type who receives the contract \hat{D} and the loss from added distortion in the next period. If the first term is large enough to be higher than \bar{b} even at its upper bound $1 - \delta$, we get the corner solution where $\bar{b} = 1 - \delta$ or equivalently $\hat{d} = 1$. Since the first term is non-negative and the lower bound of \bar{b} is 0, we can never have the lower corner solution.

Step 4: finding \underline{v} and \bar{b} . The fourth step of the analysis is to combine the results above to characterize the interior and corner solutions of the optimal mechanism design problem. We have now characterized the optimal mechanism at state D with two variables, the maximum distortion \bar{b} from contract \hat{D} and threshold \underline{v} that characterizes the opportunity cost of promising the contract \hat{D} . The two variables are related by equations (2.2) and (2.3), which can be rewritten as (2.2') and (2.3') respectively.

$$w(\underline{v}) = \frac{\delta}{1 - \delta} \frac{\int_0^{\bar{b}} b dG_b(b)}{\beta G_b(\bar{b})} = E_{G_b} \left[\frac{\delta b}{\beta(1 - \delta)} \mid \frac{\delta b}{\beta(1 - \delta)} \leq \frac{\delta \bar{b}}{\beta(1 - \delta)} \right], \quad (2.2')$$


$$\frac{\delta \bar{b}}{\beta(1 - \delta)} = E_{G_v} [w(v_i) | w(v_i) \geq w(\underline{v})]. \quad (2.3')$$

The first condition (2.2') characterizes the trade-off between promising a delayed service to marginal type \underline{v} and some revenue next period. In particular, the opportunity cost of \hat{D} is the average distortion it creates next period. The distortion has conditional distribution function $\frac{g_b(b)}{\beta G_b(\bar{b})}$, which gives us the condition. The second condition (2.3') characterizes the trade-off between increasing the set \hat{D} slightly and therefore extracting slightly more revenue from average type who receives the promise \hat{D} and the additional distortion this increase creates.

Both equations (2.2') and (2.3') describe strictly increasing continuous relationships between $\frac{\delta \bar{b}}{\beta(1 - \delta)}$ and $w(\underline{v})$, so the remaining question is when these two equations can and cannot be simultaneously satisfied. Note that by assumptions $w(\underline{v}) \in [0, 1]$ and $\bar{b} \in [0, 1 - \delta]$. If the solution to these equations is such that $\bar{b} > 1 - \delta$, then it means that the revenue is strictly increasing for each \bar{b} , so that we get the corner solution where $\bar{b} = 1 - \delta$ or equivalently $\hat{d} = 1$. Depending on the ordering of β and δ we get one of two cases illustrated by Figure 2.3.

Let's first consider the case when the seller's discount factor δ is not smaller than the buyers' discount factor β . Intuitively, when β is close to 0, we should find that it is optimal to choose very low \hat{d} . This case is illustrated by Figure 2.3(a). In this case there is always an interior solution (\bar{b}, \underline{v}) . Indeed, when β is very low, the solution has to be such that \bar{b} is very small, which

¹⁴Note that we can ignore the effect through \underline{v} by using the envelope theorem.



(a) Case $\delta \geq \beta$: only interior solutions.

(b) Case $\delta < \beta$: corner solution possibility.

Figure 2.3: Plots illustrating the characterization of the optimal waiting list mechanism.

corresponds to \hat{d} close to $F_2(p^*)$ and therefore \hat{D} close to $[0, p^*]^2$.

Note that whenever $\beta > 0$ the optimal $\bar{b} > 0$, and therefore \hat{D} includes strictly more values than just $[0, p^*]^2$. To see why this is the case suppose that the proposed contract is $[0, p^*]$. By increasing the set marginally and therefore increasing the probability marginally, the seller extracts additional revenue from all types of buyers who receive the promise of delayed service, which is strictly positive gain. The loss from this change is that the seller loses an option to offer instant service next period to the types of buyers who were added to the set. But these buyers have virtual valuation close to 0, which means that the lost revenue is close to 0.

Consider now the case when $\delta < \beta$, which is illustrated by Figure 2.3(b). Now there are two possibilities. Either there is an interior solution (\bar{b}, \underline{v}) satisfying both equations, or we get the corner solution where the contract ensures service in the next period with certainty. If this is the case both at states $D = \emptyset$ and $D = [0, 1]^2$ we get the calendar mechanism studied in Section 2.1.

The condition for getting the corner solution is relatively simple. Let \underline{v}^{cm} be the threshold used to offer delayed service in the calendar mechanism case. The average discount-adjusted virtual surplus of a buyer who receives a promise of delayed service with threshold \underline{v}^{cm} is $\beta E_{G_v}[w(v_i) | v_i \geq \underline{v}^{cm}]$. If this value is higher than δ , then even at the highest possible distortion \bar{b} , the revenue would be increasing in \bar{b} , so we would be in the corner solution. Note that by definition, $\beta w(\underline{v}^{cm}) = \frac{\delta}{1-\delta}[\pi(\emptyset) - \pi([0, 1]^2)] = \frac{\delta}{1-\delta} E_{G_b}[b]$, which gives the formal condition for the corner solution in (2.4).

$$\delta \leq \beta E_{G_v} \left[w(v) \mid w(v) \geq \frac{\delta E_{G_b}[b]}{\beta(1-\delta)} \right] = \beta E_{G_v} [w(v) | v \geq \underline{v}^{cm}] \quad (2.4)$$

Step 5: the existence of $\pi(D)$ We show that the function $\pi(d)$ exists. Since in the characterization above, only the continuation values in terms of probabilities were needed, this is sufficient to characterize $\pi(D)$ for arbitrary state.

For a given function $\pi(d)$ the characterization above defines expected revenues for each possible state and therefore for each probabilistic promise d' . Let $T\pi(d')$ denote the profit for any given d' . We want to argue that there is a unique fixed point $\pi = T\pi$. We use Blackwell's sufficient conditions to argue that T is a contraction with rate δ and then apply the Contraction mapping theorem to get the existence result. First, monotonicity. If $\hat{\pi}(d) \geq \pi(d)$ for all d , then any allocation rule with continuation value function $\hat{\pi}$ ensures higher revenue than with π and therefore the optimum cannot give lower expected revenue. Second, to verify discounting, suppose $\hat{\pi}(d) = \pi(d) + a$. Now, since the dynamic part of the revenue depends only on $\pi(0) - \pi(d)$, the optimal allocation is unchanged. The only remaining dependence is through static continuation value $\delta\hat{\pi}(1)$ and therefore $T\hat{\pi}(d) = T\pi(d) + \delta a$ for all d .

Implementation We have now showed how to choose optimal probability \hat{d} of delayed service at each state D , how to design a contract \hat{D} that fulfills probabilistic promise \hat{d} , and how to allocate the promises at each state D with a given contract \hat{D} . Combining the results we can describe the shape of optimal contracts and conclude how to implement the optimal mechanism.

There are two possible types of allocation rules in terms of qualitative results. One possibility is that the contract \hat{D} is characterized by a simple threshold rule, $\hat{D} = [0, \bar{v}]^2$ as described by Figure 2.4(a). This threshold is defined such that $(1 - \delta)w(\bar{v}) = \bar{b}$. It is optimal to use this type of contract only if optimal \bar{b} and therefore probability of service is small enough. The second possibility is slightly more complex contract described in Figure 2.4(b). Next we will discuss each case separately.

Before studying the two cases, note that from step 4 of the analysis and in particular from Figure 2.3 we see that $\frac{\delta\bar{b}}{\beta(1-\delta)} \geq w(\underline{v})$. Using the definition of \bar{v} we get $\beta w(\underline{v}) \leq \delta w(\bar{v})$. Therefore \underline{v} can be higher than \bar{v} only if buyers' discount factor β is sufficiently larger than seller's discount factor δ . A sufficient condition for being in the "high β/δ " region is when $\beta \leq \delta$, whereas when β is close to 0 we are clearly in the low " β/δ " region.

In the first case, when β/δ is relatively low (Figure 2.4(a)), the optimal mechanism promises delayed service with relatively low probability and it is possible to fulfill this promise in a region characterized by a simple threshold \bar{v} , so that $\hat{D} = [0, \bar{v}]^2$. In this case the implementation of the mechanism is analogous to the pure calendar mechanism, but with a different reserve price. The mechanism is the following: In the first stage the instant service is sold at a second price auction with reserve price \bar{v} . If neither of the buyers bid above this value, the waiting buyer is served and new buyers are refused service. However, if both buyers bid above the reserve price, the loser is offered delayed service with price $\beta\hat{d}\underline{v}$. If he accepts the offer, the winner will be given discount $\beta\hat{d}[v_l - \underline{v}]$, where v_l is the loser's bid in the second price auction. The reason for this discount is the same as in pure calendar mechanism case—the second stage creates surplus for the loser of the auction, therefore extra surplus must be given to the winner to avoid bid shading.

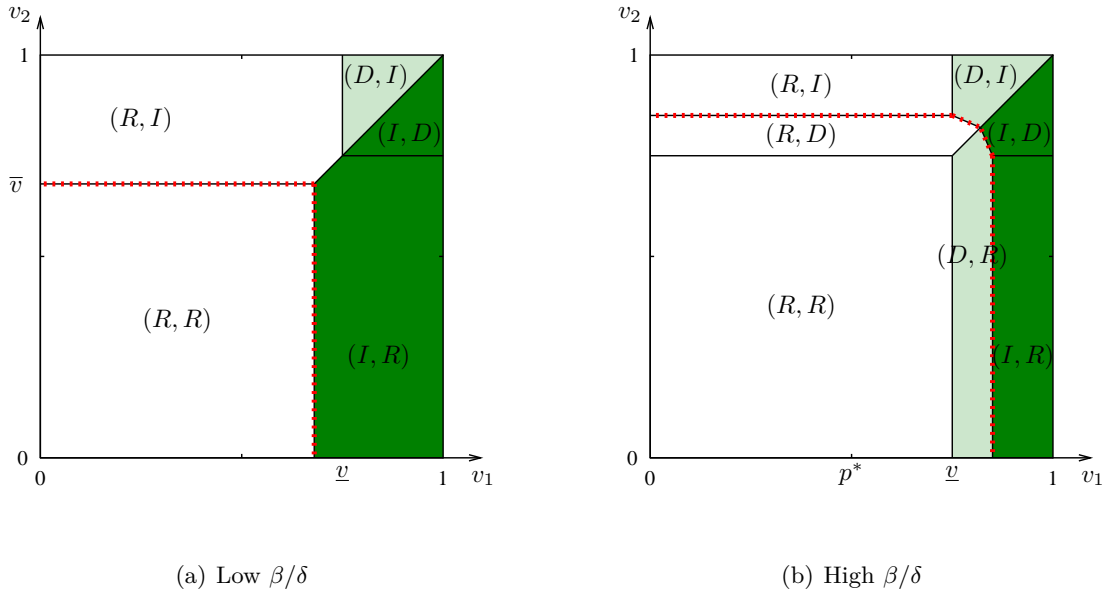


Figure 2.4: Optimal waiting list mechanism. Intensity of shading color describes the discounted expected quantity for buyer 1, the area below dotted line is the set where the previous-period customer is served.

In the second case, when β/δ is relatively high (Figure 2.4(b)), the optimal contract is more complex than a single threshold. The exact transfers can be computed from ex-post envelope condition. The optimal mechanism in this case will be implemented by a non-standard auction with scoring rule, where whether or not buyers receive instant or delayed services depends not only on their own types and order of types, but also on the similarity of their types.

In particular, there are some types of buyers who receive delayed service unless the other buyer has only slightly smaller value (in this case they get instant service). There are also types of buyers who receive delayed service unless the other buyer has only slightly higher value. Figure 2.5 focuses on these types of buyers. We will now discuss the reasons for these effects.

Some types of buyers get instant service only if they win by a close margin, receiving delayed service otherwise. An example of such type is v'' in Figure 2.5. This can be interpreted via opportunity costs. Suppose one buyer has value v'' and the opponent has negative virtual surplus. In this case it is optimal to serve the waiting buyer and offer the high type delayed service. This is a distortion to current optimum—instead of serving a buyer with positive virtual surplus instantly the seller is offering him delayed service. But ex-ante, in the previous period, this distortion was low enough to include it in the optimal contract that was given to the waiting buyer. Now, when the value of the lower type increases, then at some point there will be extra distortion from serving the waiting customer—in addition to delaying service to buyer with v'' , the seller is losing the value potentially extracted from low type. In other words, in this case delaying the service to high type comes with increased opportunity cost of not being able to serve the low type. Therefore when the opponent's value gets close enough to v'' , the waiting




Figure 2.5: Allocation given to buyer 1 under optimal waiting list mechanism when β/δ is not small. Type v' is refused service if he loses by a close margin, but gets delayed service otherwise. Type v'' gets instant service if he wins by a close margin, but gets delayed service otherwise.

buyer is not served and instead both new buyers are served, and the one with high value gets instant service. Of course, when the opponent's value rises even above v'' , the opponent will be the one receiving instant service and therefore the buyer with v'' again gets delayed service.

There are also some types of buyers who are refused service only if they lose by a close margin, but who receive delayed service otherwise. An example of such type is v' in Figure 2.5. The reason is similar. If a buyer has type v' and the opponent has even lower value, the distortion is low enough so that the seller serves the waiting buyer and offers the higher of the types delayed service. That is, when the opponent's value rises above v' , the buyer with v' will be refused service, since delayed service goes to the opponent and there are no more goods available. Now, when the opponent's value raises even higher, the total distortion increases and at some point becomes high enough so that the seller will serve not the waiting buyer but both new buyers instead. Therefore the buyer with v' will get delayed service even when being the low type.

Finally note that this mechanism has potentially non-trivial dynamics in the following sense. Suppose the initial state is such that no promises have been given, $D_0 = \emptyset$. Then the seller designs some optimal contract D_1 that is offered to some types of buyers, which will always be the lower of the two. Now, at the next period the state is D_1 , which is different from D_0 . Therefore the types of buyers who now receive a promise of delayed service come from a different distribution, which is a mixture of high and low types of buyers, characterized by a function $G_v(v_i|D_1)$. The optimal contract is now different, say D_2 . If this contract is given to a new buyer, in period 2 the state is D_2 , which is again different from D_0 and D_1 . Continuing the same way, we get a sequence of promises D_1, D_2, \dots , which are all different and optimal, given that there has been a sequence of arrivals such that one buyer always got a promise of delayed service.

2.3 Fully optimal mechanism

In this section we analyze the fully optimal mechanism, where the seller can choose any allocation rule she might like. Compared to the waiting list this means removing the last remaining restriction—at each state with any vector of types the seller is now allowed to give the optimal promise of delayed service to each buyer.

The relaxation of this assumption affects the mechanism design problem. The fully optimal mechanism is similar to the optimal waiting list with the exception that the promise of delayed service will depend on the buyer's type who receives the promise. In particular, the probability of receiving the service in the next period is an increasing function of the type. This simplifies the analysis, since instead of designing a contract that works for the average buyer who gets the contract, the seller can now design a separate contract for each type.

There are some straightforward observations that make the analysis simpler. First, the seller will promise delayed service to only one buyer¹⁵. Therefore a state is still simply D , the promise given to one of the previous buyers. Second, we can again separate the optimal assignment decision to probability of delayed service \hat{d} , and contract \hat{D} which will be optimal given this probability. Note that initial state may be sub-optimal, but continuation states are always optimal. Even if the initial state is such that there are two buyers waiting, their transfers are already sunk, so we can treat of them as a single buyer.

We proceed by characterizing the fully optimal mechanism in four steps. First we take the expected revenue functions $\pi(d)$ as given and characterize the optimal allocation rule. The optimal allocation rule will now include the optimal promise function $\hat{d}(\cdot)$, which at this point is characterized as a function of continuation values. Second, we discuss how the probabilities of delay are fulfilled optimally. Third, we characterize the optimal promise function $\hat{d}(v_i)$ explicitly and discuss its properties. Finally, we use contraction mapping theorem to show that functions $\pi(d)$ exist. Since the analysis is similar to the waiting lists, we are omitting some of the details and emphasizing the differences.

Step 1: allocation rule. Again, suppose $\pi(D)$ and $\pi(d)$ functions exist and fix a state D . As before, we denote the optimal contract of delayed service offered to a buyer in the current period by \hat{D} (which will then be the state in the next period). Since the buyers are interested in probability of service, we are again characterizing the optimal promises in terms of probability $\hat{d} = Pr(\hat{D})$ and later will describe how the seller optimally fulfills these promises.

Suppose the new buyers announce types $v \in D$. Then the seller has to serve the previous buyer and therefore can offer delayed service to only one of the buyers. This is offered to the highest of the two buyers with some probability \hat{d} . The optimal choice of \hat{d} solves the following maximization problem

$$\underline{\pi}(v) = \max_{\hat{d}} \left\{ \delta \pi(0) + (1 - \delta) \beta \hat{d} w(v_2) - \delta [\pi(0) - \pi(\hat{d})] \right\},$$

¹⁵Instead of promising D_1 and D_2 to two current buyers, promising $D_1 \cup D_2$ to the higher of them would increase flow revenue while keeping .

where the first component is the static continuation revenue that the seller gains by not giving any promises; the second term is the revenue extracted from the high type by giving a promise of service with probability \hat{d} in the next period; and the third term is the loss in continuation value from giving this promise. The revenue extracted from the high type is now his standard static virtual surplus $w(v_2)$ adjusted by discount factor β and probability \hat{d} . The first order condition from the maximization problem is

$$(1 - \delta)\beta w(v_2) + \delta\pi'(\hat{d}) \geq 0. \quad (2.5)$$

We will come back to the function \hat{d} once we know how to compute $\pi(d)$ and differentiate it. From (2.5) we see that $\hat{d}(v_2)$ is only a function of v_2 and it balances the trade-off from extracting revenue from the buyer and having higher opportunity cost because of being in a more restrictive state in the next period. In particular, $\hat{d}(v) = 0$ whenever it is promised to a buyer with negative virtual surplus and $\hat{d} \geq F_2(p^*)$ for all $v_2 > p^*$.

When buyers announced $v \notin D$, the instant service is available, so that the higher of the new buyers is served instantly if and only if he has strictly positive virtual surplus, whereas the lower new buyer is offered delayed service with probability \hat{d} , that solves the following optimization problem

$$\bar{\pi}(v) = \max_{\hat{d}} \left\{ (1 - \delta) \max\{0, w(v_2)\} + \delta\pi(0) + (1 - \delta)\beta\hat{d}w(v_1) - \delta[\pi(0) - \pi(\hat{d})] \right\}.$$

There are two differences with the previous maximization problem. First, the static revenue is increased, since some revenue is extracted from serving the high type instantly. Second, the delayed revenue comes now from the low type rather than the high type, which means lower dynamic revenue in expectation. The first order condition is the same as previously given with (2.5), but for v_1 instead of v_2 . Therefore the optimal promise is the same function $\hat{d}(v_i)$ in both cases; the difference is that it is promised to a different buyer in each of the two cases.

Graphical illustration of the combined allocation rule is by Figure 2.6.

Step 2: fulfilling promises optimally. As with the waiting list, the difference of revenues in the two cases, $\bar{\pi}(v) - \underline{\pi}(v)$, can be interpreted as the distortion from serving a buyer from the previous period rather than serving new buyers optimally. This difference is strictly increasing in v_2 for $v_2 \geq p^*$. It is again optimal to choose contract \hat{D} so that it minimizes $\bar{\pi}(v) - \underline{\pi}(v)$ for a given probability of delayed service $\hat{d} \geq F_2(p^*)$, so we get $\hat{D} = \{v \in [0, 1]^2 : \bar{\pi}(v) - \underline{\pi}(v) \leq \bar{b}\}$ such that $Pr(D) = \bar{b}$.

Step 3: optimal probabilities of delayed service. As with waiting lists, it is more convenient to optimize with respect to maximum distortion \bar{b} than the probability of service \hat{d} . Let $\bar{b}(v_i)$ denote the maximum distortion corresponding to a promised probability $\hat{d}(v_i)$.

In this notation, a slight increase in promised probability \hat{d} means that a new set of values must be added to the contract \hat{D} where the instant service will be unavailable next period. If




Figure 2.6: Optimal allocation in fully optimal mechanism. Intensity of shading color describes the expected discount factor for buyer 1, the area below dotted line is the set where delayed service is offered.

the change is marginal, all those added values have distortion that is close to the maximum distortion \bar{b} . That is, $\pi'(\hat{d}(v_i)) = -\bar{b}(v_i)$.

We can now analyze equation (2.5), the first order condition that characterizes optimal promise $\hat{d}(v_i)$, which takes the form

$$(1 - \delta)\beta w(v_i) = \delta \bar{b}(v_i).$$

Here we see that whenever $w(v_i) > 0$, the maximum distortion $\bar{b}(v_i) > 0$ and so the optimal probability $\hat{d}(v_i) > F_2(p^*)$. Moreover, maximum distortion $\bar{b}(v_i)$ and therefore probability $\hat{d}(v_i)$ are strictly increasing functions of v_i for all interior solutions. Finally, we get the upper corner solution where the probability $\hat{d}(v_i) = 1$, or equivalently maximum distortion $\bar{b}(v_i) = 1 - \delta$ whenever $w(v_i) \geq \frac{\delta}{\beta}$.

If $\delta = \beta$, each type of buyer is promised different probability, with type $v_i = 1$ being the only type who receives delayed service with certainty. If $\delta > \beta$, then the buyers value the future service relatively lower than the seller and therefore the probabilities are relatively low. In particular, all promises are strictly lower probability than 1, and in the limit when β is very small, all promises are with probability $\hat{d}(v_i)$ very close to $F_2(p^*)$. On the other hand, if $\delta < \beta$, the buyers value the delayed service highly and in this case even some buyers with value strictly below 1 will receive delayed service with certainty.

Step 4: existence of $\pi(d)$. This step is exactly analogous to step 5 in the waiting list mechanism—we can again verify that Blackwell’s sufficient conditions are satisfied and therefore

the mapping we characterized above is a contraction (with rate δ). By the contraction mapping theorem there exists a fixed point, so that $\pi(d)$ is indeed well defined.

Implementation: The allocation rule described above is again implemented by a non-standard auction with scoring rule. The new aspect compared to the waiting list is that a promise of delayed service is personalized. It means that each type of buyer who gets delayed service pays a transfer that is a strictly increasing function of her own type¹⁶.

Similarly to the waiting list case the scoring rule favors close winners and punishes close losers. The reason for both effects is the same as with waiting list mechanisms—when deciding whether or not to serve a buyer from the previous period, the seller needs to compute the total distortion from doing so. This distortion is increasing in both new buyers' values, and therefore the allocation that a buyer gets also depends on the opponent's bid.

2.4 Partial commitment

Since this paper focuses on promises of delayed service, commitment power plays a crucial role. Whenever a promise is restrictive, the seller would have incentive to break it in the future. In this section and in the general model in the next section, we are extending the results from the three classes of mechanisms studied thus far to a general partial commitment assumption. The full commitment analysis we did above and no commitment are special cases of the general assumption.

There are many reasons why the promises are kept: legal rules and fines, reputation of the seller, or perhaps intrinsic motivation. In this paper we are taking a reduced-form approach and define partial commitment as an exogenous cost of breaking a promise to one buyer. This could be interpreted as a fine that has to be paid to a government agency. Of course, the model could be extended to make the commitment constant endogenous, for example by modeling reputation explicitly or allowing the seller to pay the compensation directly to the consumer, but this is not the focus of the current paper.

As always in mechanism design literature, we assume that the sets of states and promises are such that the seller is able to fully commit to any promise she has made. To study the lack of commitment, we are restricting the set of possible promises.

We define partial commitment with constant c as a restriction to possible promises and therefore states. A seller cannot give any promise with which at any future realization the benefit from breaking a promise to one buyer would increase revenue by more than c . Then c can be interpreted as a commitment cost or fine for breaking the promise. If $c = 0$ we say that the seller has no commitment power and if $c \rightarrow \infty$ she has full commitment power.

Benchmark: no commitment. As a benchmark, let's consider the case when the seller does not have any commitment power. This will be the limiting solution for both the waiting list and

¹⁶In particular, the transfer a buyer with value v_i pays when receiving delayed service with probability $\hat{d}(v_i)$ is $p_i(v) = \beta \hat{d}(v_i) v_i - \int_{p^*}^{v_i} \hat{d}(v'_i) dv'_i$ by the ex-post envelope condition

the unrestricted mechanism. Then the commitment constant is $c = 0$, which means that breaking any promise is costless. This is as if the next period seller is a separate player and the current seller does not have any power to enforce promises. In this situation, promising delayed service is limited. Since the transfers from the previous period are already sunk, the seller always gives the priority to new arrivals. Whenever a new buyer has positive virtual surplus, the seller would like to serve her as soon as possible. Feasibility implies that higher type is offered instant service and lower type is offered delayed service with highest feasible probability, whenever their types are above p^* . This puts an upper bound to the promises that can be given about the future service. The seller can promise service only in set $\hat{D} = [0, p^*]^2$, where p^* is defined so that virtual surplus $w(p^*) = 0$. Therefore if the buyer with lower type has positive virtual surplus, she is offered to be served with probability $F_2(p^*)$ and will be served when $v \in \hat{D} = [0, p^*]^2$. This allocation rule is illustrated by Figure 2.7.




Figure 2.7: Optimal no-commitment allocation. Intensity of shading color describes the discounted expected quantity for buyer 1, the area below dotted line is the set where the previous-period customer is served.

The optimal no-commitment mechanism can be implemented by the following indirect mechanism. First, both types are asked to pay $F_2(p^*)\beta p^*$ for a chance to be served tomorrow, if the seller has free time. After this, an upgrade to instant service is sold at a second price auction with reserve price $[1 - \beta F_2(p^*)]p^*$.

Pure calendar mechanism with partial commitment. As long as commitment constant c is higher than $1 - \delta$, commitment level is sufficient to sustain the optimal pure calendar mechanism described in Section 2.1. However, when $c < 1 - \delta$, the only possible calendar mechanisms are static. The discreteness of the result comes from the pure calendar mechanism assumption, where the seller either does or does not use unconditional promises of delayed service; there are no

intermediate cases. Promising delayed service with certainty decreases the revenue in the next period, which means that the pure calendar mechanism requires some commitment power. The question is how much commitment is sufficient.

The incentive to break a promise to a previous buyer would be the highest when both new arrivals have very high type, so the limiting case is when $v = (1, 1)$. In this case when keeping the promise, the seller would get $\beta w(1) = \beta$ from promising delayed service to one of the buyers and would have to refuse service to the other. When breaking the promise, he could instead get $1 + \beta$, with the same continuation value. This means that the maximal gain from breaking the promise in terms of normalized revenue is $(1 - \delta)w(1) = 1 - \delta$. This is the necessary commitment level c^* that ensures that the seller is able to keep any promises.

Waiting list mechanism with partial commitment. Above we computed the optimal no-commitment mechanism, which is a waiting list mechanism, since the delayed service is always with a fixed set $\hat{D} = [0, p^*]^2$ and therefore does not discriminate between buyers who are offered delayed service. As the opposite benchmark we have also characterized the optimal waiting list with full commitment in Section 2.2.

The optimal waiting list with full commitment at state D gave an optimal promise in the form $\hat{D} = \{v : \bar{\pi}(v) - \underline{\pi}(v) \leq \bar{b}\}$, where $b \in [0, 1 - \delta]$ is the maximum distortion that we know how to compute, and $\bar{\pi}(v) - \underline{\pi}(v)$ is the distortion from serving a buyer from the previous period to a current revenue when arriving buyers announce v (the difference between revenues when the instant service is and is not available). Therefore with arrivals v , if the seller breaks the promise to a previous buyer, the additional revenue is exactly $\bar{\pi}(v) - \underline{\pi}(v)$. This means that the maximal incentive to break a promise \hat{D} is the maximum distortion \bar{b} . If the maximum distortion \bar{b} is lower than commitment constant c , the seller can give promise \hat{D} , whereas when $\bar{b} > c$, it is known that the promise will be broken under some realizations and therefore this promise is not possible.

From this argument we see that finding the optimal waiting list mechanism for general partial commitment c is not very different from what we did with full commitment. The only difference is that maximum distortion \bar{b} now has an additional upper bound c .

It is illustrative to consider the two extremes. In the limit when commitment constant c converges to 0 we obviously get the no-commitment solution described above. However, the no-commitment solution is optimal only in the limit. Whenever $c > 0$, the seller would make the contract \hat{D} at least slightly larger than $[0, p^*]^2$. Increasing the probability of service is valuable for all buyers who receive delayed service and therefore strictly increases revenue. Near the no-commitment solution it is almost costless, since the first types added to \hat{D} have virtual surplus close to zero and therefore there is almost no distortion from not serving them.

The upper bound of \bar{b} in the full commitment analysis was $1 - \delta$. Therefore, if $c \geq 1 - \delta$, the commitment constraint is never binding. This is the same sufficient commitment level that was needed to commit to the optimal pure calendar mechanism. When $c < 1 - \delta$, we will sometimes get a corner solution where maximum distortion $\bar{b} = c$ and therefore the probability of service

\hat{d} is always strictly below 1. In this sense a waiting list can be seen as a relaxation of a pure calendar mechanism.

Unrestricted mechanism with partial commitment. The effect of adding partial commitment to the fully optimal mechanism (discussed in Section 2.3) is similar as to the waiting list. Again, it means that any promise of delayed service must be such that the distortion created by this promise to the next period's revenue is not larger than c . The only difference is that in this case we allow the promises to depend on the type of the buyer who receives the promise, so that the upper bound must hold for each type. Formally $\bar{b}(v_i) \in [0, \min\{1 - \delta, c\}]$. Again, the limit when $c \rightarrow 0$ is the no-commitment mechanism (but it is optimal only if $c = 0$) and when $c \geq 1 - \delta$, the commitment level is sufficient to sustain any mechanism.

3 General analysis

In this subsection we give formal results for arbitrary fixed number of service slots and random number of buyers. The main trade-offs are the same as with the illustrative example, but obviously the notation and characterizations of optimal mechanisms are somewhat more complex. In contrast to the previous section, we prove the results with the partial commitment assumption, so that full commitment is a special case.

3.1 Model

Seller has fixed number $m \in \mathbb{N}$ service slots each period $t \in \{0, 1, \dots\}$ and maximizes expected discounted revenue with discount factor $\delta \in (0, 1)$.

Every period random number n new buyers arrive such that the maximum number of buyers arriving with positive probability, $N > m$. Each buyer has unit demand and independent private value v_i which is i.i.d. draw from CDF F and PDF $f(v) > 0$. Let $F_{k:n}$ and $f_{k:n}$ denote the CDF¹⁷ and PDF¹⁸ of k 'th smallest value out of n (k 'th order statistic). Then for example, $F_{n:n}$ is the CDF of the maximum of the values of all arrived customers.

We maintain two important assumptions. First, each buyer interacts with the seller only at arrival. Buyer i who is promised service s_i periods after her arrival and pays p_i at arrival gets payoff of $\beta^{s_i} v_i - p_i$. Second, we concentrate on mechanisms where the seller never makes promises of delayed service more than one period ahead.

The promise of delayed service to buyer i is a set $D_i \subset \{(n, v) : n \in \{0, \dots, N\}, v \in [0, 1]^n\}$ and a state is a vector $D = (D_1, \dots, D_N)$, a collection of promises given to buyers from the previous period. If buyer i receives a promise of delayed service with a contract D_i , she will be served in the next period if and only if the number of buyers n and their announcements v are such that $(n, v) \in D_i$.

¹⁷ $F_{k:n}(v_i) = \sum_{j=k}^n \binom{n}{j} F(v_i)^j [1 - F(v_i)]^{n-j}$.

¹⁸ $f_{k:n}(v_i) = \frac{n!}{(k-1)!(n-k)!} F(v_i)^{k-1} [1 - F(v_i)]^{n-k} f(v_i)$.

Since buyers' utility functions are linear in value and quasilinear in transfers, the only variables affecting buyer i 's payoff are her transfer $p_i(v)$ and discounted expected quantity, denoted by $\phi_i(v)$, which is 1 if i receives instant service, 0 if refused service, and $\beta Pr(D_i)$ if is promised delayed service in set D_i . These variables are functions of the allocation the seller chooses as a response to vector of types v announced by the buyers in this period. Let $\tilde{p}_i(\hat{v}_i)$ and $\tilde{\phi}_i(\hat{v}_i)$ denote the expected values of the respective variables over the other buyers' types assuming they report their types truthfully.

We are looking for (Bayesian) incentive-compatible and individually rational mechanisms. In this notation, the individual rationality constraint (IR) for buyer i with type v_i is $\tilde{\phi}_i(v_i)v_i - \tilde{p}_i(v_i) \geq 0$. The incentive-compatibility constraint (IC) is $\tilde{\phi}_i(v_i)v_i - \tilde{p}_i(v_i) \geq \tilde{\phi}_i(\hat{v}_i)v_i - \tilde{p}_i(\hat{v}_i)$ for all $\hat{v}_i \in [0, 1]$.

Using the standard methods developed by Myerson (1981) it is straightforward to verify that a mechanism satisfies IC and IR constraints if and only if (1) the discounted expected quantity $\tilde{\phi}_i$ is weakly increasing in buyer's own type, (2) buyer with type $v_i = 0$ gets payoff at least 0 (an in optimum exactly 0), (3) the transfers are computed from ex-ante envelope condition

$$\tilde{p}_i(v_i) = v_i \tilde{\phi}_i(v_i) - \int_0^{v_i} \tilde{\phi}_i(\hat{v}_i) d\hat{v}_i + \tilde{p}_i(0). \quad (3.1)$$

Although we are looking for optimal mechanisms that induce truthful behavior when buyers do not know the other buyers' types, it turns out that all the optimal mechanisms are implementable even when they know the other types. Let's define ex-post individual rationality constraint (EPIR) for buyer i when the type vector is $v = (v_1, \dots, v_n)$ by $\phi_i(v)v_i - p_i(v) \geq 0$. The dominant-strategy incentive-compatibility constraint (DSIC) is defined as $\phi_i(v)v_i - p_i(v) \geq \phi_i(\hat{v}_i, v_{-i})v_i - p_i(\hat{v}_i, v_{-i})$ for all i 's announcements $\hat{v}_i \in [0, 1]$ and vectors of opponents' types $v_{-i} \in [0, 1]^{n-1}$.

Again, it is straightforward to verify using standard methods that a mechanism satisfies DSIC and EPIR constraints if and only if (1) the ex-post discounted expected quantity ϕ_i is weakly increasing in buyer's own type v_i , (2) buyer with type $v_i = 0$ always gets a payoff at least 0 (and in optimum exactly 0), and (3) the transfers are computed from ex-post envelope condition

$$p_i(v_i) = v_i \phi_i(v_i) - \int_0^{v_i} \phi_i(\hat{v}_i) d\hat{v}_i + p_i(0). \quad (3.2)$$

We are using these results for two purposes. First, the ex-ante envelope condition (3.1) and the fact that $p_i(0)$ must be 0 are used to express the seller's expected revenue at state z as

$$\pi(z) = \max E_{n,v} \left[(1 - \delta) \sum_{i=1}^n \tilde{p}_i(v_i) + \delta \pi(z') \right] = \max E_{n,v} \left[(1 - \delta) \sum_{i=1}^n \phi_i(v) w(v_i) + \delta \pi(z') \right],$$

where the seller maximizes over feasible promises, $w(v_i) = v_i \frac{1-F(v_i)}{v_i}$ is the standard static virtual surplus, and z' is the continuation state. This expression is separable in (n, v) . That is, the seller can choose optimal promises for each possible combination (n, v) separately and the value of being in a particular state can then be computed as expectation over all realizations.

We will find the optimal allocation rule for different restrictive classes of mechanisms that restrict feasibility. In all the cases, the optimal allocation rule turns out to be monotone with respect to buyers' own types. Therefore, as the second use of the results above, we can apply ex-post envelope condition (3.2) to compute the transfers that implement this mechanism under even stronger conditions (EPIR, DSIC).

Before studying the dynamic mechanisms, there is one special mechanism in this model that serves as a useful benchmark and for which the allocation rule and implementation are well known. We call a mechanism static, if the seller never gives promises for delayed service. That is, at each period, each buyer either gets instant service or is refused service. This mechanism does use delayed service, so the restrictions for possible promises do not apply and it is therefore a feasible mechanism in all the cases studied below.

Analogously to Myerson (1981), the static optimum in this framework is the allocation rule that assigns objects to m buyers with highest types, given that their virtual valuations are positive. This is implemented under DSIC and EPIR constraints by a standard uniform price auction with reserve price p^* such that $w(p^*) = 0$.

3.2 Pure calendar mechanism

The pure calendar mechanism is a mechanism where delayed service is unconditional on new arrivals. Each buyer is either assigned instant service, refused service, or promised service in the next period with certainty. The general set of states $D = (D_1, \dots, D_N)$ can be simplified to a number $z \in \{0, \dots, m\}$. This is true since instead of a promise set we only have to keep track of which buyers received the delayed service and which did not. Moreover, from the seller's point of view, who did and who did not receive the delayed service is not payoff-relevant—the only relevant variable is how many people were promised delayed service. Therefore a state $z \in \{0, \dots, m\}$ denotes how many service slots are unavailable, and thus $m - z$ slots are still available.

The main result in this section is Theorem 3.1, which characterizes the optimal pure calendar mechanism (with partial commitment). The result has the same general trade-offs as the illustrative example in Section 2.1. First, delay is optimal only if buyer's discount factor β is high enough compared to δ . In particular, $\frac{\beta}{\delta}$ must be at least $E_{n,v}[\max\{w(v_{n:n}), 0\}]$, which is the average of expected revenue that can be extracted from the highest type buyer in a static mechanism. If β is strictly less than δ times this value, then it would not be optimal to give delayed service even to buyers with highest possible values of 1. The static mechanism would be strictly better. Note that the requirement is satisfied, for example, when $\beta \geq \delta$.

Second, the seller's commitment power must be high enough. The sufficient commitment power is the same as in the illustrative example, $c \geq 1 - \delta$ for the same reason. If this condition is satisfied, even if all new arrivals have types very close to 1, the seller would prefer not to break previous promises. But in general it is necessary only if $N \geq 2m$. When $N < 2m$, there can be only $N < 2m$ buyers with high positive surplus next period. This limits the cost of current promises, since the seller is never forced to refuse service to very high customers (he can extract

part of revenues by promises of future service).

The main trade-off is again between extracting revenue from current relatively high buyers and the opportunity cost of having one less service slot available next period. Giving away each additional unit of next period service will be associated with opportunity due to being in a more restrictive state next period. The optimal pure calendar mechanism balances extraction from current buyers by promising some of them delayed service and the opportunity costs of each additional slot that is promised away.

In particular, the optimal pure calendar mechanism is characterized by m thresholds, denoted by $\underline{w} = (\underline{w}_1, \dots, \underline{w}_m)$, where \underline{w}_k can be interpreted as the opportunity cost of the k 'th object. The actual choice of promises will be non-trivial. We will describe below how to make this choice and Corollary 3.3 describes a special case where the allocation rule is much simpler.

Theorem 3.1. *Suppose $m \in \mathbb{N}$ and $n \sim \Gamma$. Under sufficient conditions $\beta \geq \delta E_{n,v}[\max\{w(v_{n:n}), 0\}]$ and $c \geq 1 - \delta$, the optimal pure calendar mechanism with partial commitment uses the following allocation rule at each state $z = \{0, \dots, m\}$*

- (i) $m - z$ new buyers with highest strictly positive virtual surplus will get instant service.
- (ii) If there are more than $m - z$ buyers with positive virtual surplus, the next k highest of them will be promised service in the next period, where k is defined as

$$k = \arg \max_{k \in \{0, \dots, m\}} \sum_{i=1}^m [w(v_{n+1-(m-z)-i:n}) - \underline{w}_i]$$

where constants $\underline{w}_1, \dots, \underline{w}_m$ are uniquely determined by

$$\begin{aligned} \underline{w}_j &= \frac{\delta}{\beta} E_{n,v} \left[\max\{0, w(v_{n-(m-j):n})\} \right. \\ &\quad \left. + \beta \left(\max_k \sum_{i=1}^k [w(v_{n+1-(m-j)-i:n}) - \underline{w}_i] - \max_k \sum_{i=1}^k [w(v_{n-(m-j)-i:n}) - \underline{w}_i] \right) \right]. \end{aligned}$$

The result covers the situations where it is optimal to offer delayed service and commitment level is sufficient to do so. The following Corollary 3.2 shows that these sufficient conditions are also almost necessary. In particular, the condition relating β and δ is necessary, whereas the condition on commitment level is necessary when the probability that at least $2m$ new buyers arrive is positive.

Corollary 3.2. *If one of or neither of the sufficient conditions do not hold:*

- (i) When $\beta < \delta E_{n,v}[\max\{w(v_{n:n}), 0\}]$, the optimal mechanism is static.
- (ii) When $N \geq 2m$ and $c < 1 - \delta$, the optimal mechanism is static.
- (iii) When $m < N < 2m$ and $c < 1 - \delta$, depending on details the optimal mechanism could be static or with some delay.

The proofs for both are in Appendix A.1. Here is the overview of the proofs. We start by assuming that β and c are high enough, so that there is some optimal delay. We can then make some instant observations: (1) buyers with negative virtual surplus are never served, (2) if there are enough buyers with positive virtual surplus, all current objects will be allocated, and (3) promises are monotone in types (higher types will be served earlier). This gives an almost full description of optimal allocation rule.

The remaining question is for situations where there are more buyers with strictly positive virtual surplus than current service slots. The question is, to how many of them should the seller promise delayed service. To do this, we define constants $\underline{w}_1, \dots, \underline{w}_m$ where \underline{w}_k can be interpreted as the opportunity cost of giving away k th unit of tomorrow's service. Now, vector \underline{w} is defined in terms of continuation values and is therefore endogenous. To show that \underline{w} exists and is uniquely defined we write the equation system for \underline{w} in the form of $\underline{w} = \Phi(\underline{w})$ and show that there exists a fixed point. The solution—if it exists—will be unique because Φ will be strictly monotone.

The final part of the proof is about commitment power. We compute the cost of keeping a promise in all possible situations for each additional buyer. Value $1 - \delta$ will be an upper bound of these costs, therefore if $c \geq 1 - \delta$, commitment power is sufficient. When $N \geq 2m$, this upper bound is the lowest upper bound, since it is possible that $2m$ buyers arrive and all have values close to 1, which gives the bound. When $N < 2m$, this situation is not possible and therefore it is possible to have optimal delay even when c is slightly below $1 - \delta$.

Theorem 3.1 implies that the profit from calendar mechanism at state z can be written as

$$\pi(z) = E_{n,v} \left[(1 - \delta) \sum_{i=1}^{m-z} \max\{0, w(v_{n+1-i:n})\} + \delta\pi(m) + (1 - \delta)\beta \max_k \sum_{i=1}^k [w(v_{n+1-(m-z)-i:n}) - \underline{w}_i] \right].$$

The first two terms add up to the maximum revenue that can be extracted by static promises—the revenue from highest $m - z$ buyers plus continuation value of having all m object available next period. The last term is positive and captures the extra revenue from delayed service to buyers with lower valuations. In particular, it is the sum of virtual surpluses extracted from the k buyers who are promised delayed service minus the opportunity costs \underline{w}_i that are lost due to the fact that fewer objects are available in the next period.

Suppose (n, v) is such that there are enough buyers with positive surplus to allocate all current and next period objects, that is, $v_{n-2m:n} > p^*$. In this case, the number of buyers who are served instantly z , which is clearly increasing in z . In the appendix we prove Lemma A.1 which shows that the optimal number of buyers that are offered delayed service is weakly decreasing in $m - z$. This is true since with higher $m - z$ more of the high-value buyers will receive instant service and therefore the remaining buyers have lower values.

In general the optimal mechanism is non-monotone in terms of assigning delayed service. It is possible, for example, that at some state with some vector of values two buyers are offered delayed service, but after weakly increasing all the values only one buyer is offered delayed service. There is one case where the mechanism is relatively easier to interpret and implement—if the thresholds characterized by Theorem 3.1 happen to be ordered. In this case the assignment is

monotone and therefore each of the following buyers is served if and only if her value is higher than the corresponding threshold. Corollary 3.3 gives the formal result.

Corollary 3.3. *If the assumptions in Theorem 3.1 hold and the implied vector \underline{w} is such that $\underline{w}_1 \leq \dots \leq \underline{w}_m$, then the optimal mechanism uses the following allocation rule at each state z*

- (i) *$m - z$ buyers with highest strictly positive virtual surplus will get instant service.*
- (ii) *If there are more than z buyers with positive virtual surplus, k highest of them will be promised service in the next period, where k is such that $w(v_{n+1-(m-z)-i:n}) > \underline{w}_i$ if and only if $i \leq k$.*

The expected revenue at state is

$$\pi(z) = E_n \left[(1 - \delta) \sum_{i=1}^{m-z} \pi_{n+1-i}(0) + \delta \pi(0) + (1 - \delta) \beta \sum_{i=1}^m \pi_{n+1-(m-z)-i}(\underline{w}_i) \right].$$

where $\pi_k(\hat{w}) = \int_{\underline{w}^{-1}(\hat{w})}^1 [w(v_{k:n}) - \hat{w}] dF_{k:n}(v_{k:n})$.

Proof is in Appendix A.1.

3.3 No commitment

As a benchmark, we are considering the limiting case where the seller does not have any commitment power. That is, commitment constant $c = 0$. In this case the seller always gives priority to new arrivals, since the payments from previous buyers is sunk. That is, whenever a new buyer with strictly positive virtual surplus arrives, she will be served before all waiting buyers. This implies that promises can be made only for the cases when there are sufficiently few new buyers with positive virtual surpluses. The following Proposition 3.4 formalizes this.

Proposition 3.4. *Suppose $m \in \mathbb{N}$ and $n \sim \Gamma$. The optimal no-commitment mechanism is characterized as follows. Let k be the number of new buyers with strictly positive virtual surpluses.*

- (i) *If $k \geq m$, then m new buyers with highest strictly positive virtual surpluses are served instantly and no waiting buyers are served. If $k < m$ then all k new buyers with strictly positive virtual surplus as well as $m - k$ waiting buyers with highest positive surpluses are served.*
- (ii) *Buyers with strictly positive surpluses who were not served instantly are promised that they will be served if sufficiently few buyers with strictly positive surpluses arrive next period. In particular, $n + 1 - m - i$ th highest buyer will be served next period if and only if his virtual surplus is strictly positive and next period $k' \leq m - i$ new buyers with strictly positive virtual surplus arrive.*

Proof follows from the discussion above.

3.4 Waiting list

The waiting list mechanism is a mechanism where the promise depends only on the position and not the type of buyer who is in the position. Formally, at each state $D = (D_1, \dots, D_N)$, the waiting list is characterized by a vector of contracts $\hat{D} = (\hat{D}_1, \dots, \hat{D}_N)$ that are fixed before the arrival of buyers. Each contract \hat{D}_i is a set $\hat{D}_i \subset \{(n, v) : n \in \{0, \dots, N\}, v \in [0, 1]^n\}$ that describes under what conditions the receiver of the contract will be served.

In addition to contracts \hat{D} being fixed before the number of buyers and their values being observed, the waiting list assumption says that when k new buyers are promised delayed service, they receive contracts $\hat{D}_1, \dots, \hat{D}_k$. This is the “list” part of the assumption. For example if buyer i was promised set \hat{D}_k , it is as if he is added to some list in the position k , after $k - 1$ other buyers. He will be served if and only if n new buyers who arrive next period announce values v such that $(n, v) \in \hat{D}_k$.

Notice that we do not explicitly assume that the waiting list is a priority ranking in the sense that being higher in the list ensures more likely service than being in a lower position. This is a property of an optimal waiting list. In particular, by Theorem 3.5 the optimal waiting list is such that $\hat{D}_1 \supset \hat{D}_2 \supset \dots \supset \hat{D}_m$ and $\hat{D}_k = \emptyset$ for all $k > m$.

Another property of optimal waiting list will be that each contract is determined separately from all other contracts. It will balance the same trade-offs that the illustrative example. On one hand, the assignment of j th contract balances the marginal surplus extracted from a buyer who gets this contract (analogously to the pure calendar mechanism) and on the other hand the precise composition of the contract balances the trade-off between average surplus extracted from a buyer who gets the contact and increased distortion from this contract (which is analogous to the fully optimal case).

Theorem 3.5. *The optimal waiting list mechanism with partial commitment is such that at each state $D = (D_1, \dots, D_N)$, when n buyers arrive and announce values v , then*

- (i) $z = \#\{i : (n, v) \in D_i\}$ waiting buyers and $m - z$ new buyers with highest positive virtual surpluses are served instantly.
- (ii) The next m highest buyers with positive virtual surpluses are offered delayed services with sets $\hat{D} = (\hat{D}_1, \hat{D}_2, \dots, \hat{D}_m)$.

where contracts \hat{D} are characterized as follows.

$$\hat{D}_j = \{(n, v) : \pi^{j-1}(n, w) - \pi^j(n, w) \leq \bar{b}_j\},$$

$$\pi^j(n, w) = (1 - \delta) \sum_{i=1}^{m-j} \max\{0, w_{n+1-i}\} + \delta \pi(0, \dots, 0) + (1 - \delta) \beta \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k \hat{d}_i [w_{n+1-(m-j)-i} - \underline{w}_i],$$

and $\underline{w}_j \in [0, 1]$, $\bar{b}_j \in [0, \min\{1 - \delta, c\}]$ are characterized by

$$\underline{w}_j = \frac{\delta}{(1 - \delta)\beta} \int_0^{\bar{b}_j} b_j \frac{dG_{b_j}(b_j)}{G_{b_j}(\bar{b}_j)} \quad \text{and} \quad \frac{\delta \bar{b}_j}{(1 - \delta)\beta} \geq \int_{\underline{w}_j}^1 \hat{w}_j \frac{dG_{w_j}(\hat{w}_j)}{1 - G_{w_j}(\underline{w}_j)},$$

where $G_{b_j}(b_j) = Pr(\{(n, v) : \pi^{j-1}(n, w) - \pi^j(n, w) \leq \bar{b}_j\})$ and $G_{w_j}(\hat{w}_j) = \int_0^{\hat{w}_j} g_{w_j}(\hat{w}'_j) d\hat{w}'_j$, st

$$g_{w_j}(\hat{w}_j) = E_n \left[\mathbf{1}[D_1^c] f_{n+1-(m-j)}(\hat{w}_j) + \sum_{i=1}^{m-1} \mathbf{1}[D_i \setminus D_{i+1}] f_{n+1-(m-j)+i}(\hat{w}_j) + \mathbf{1}[D_m] f_{n+1-(m-j)+m}(\hat{w}_j) \right].$$

Proof is in Appendix A.2 and follows the same steps as the illustrative example. Proof structure:

- Step 0 We assume that functions $\pi(d)$ are well-defined and make some immediate observations regarding the contracts and the allocation rule.
- Step 1 For a fixed vector of contracts \hat{D} we characterize the optimal allocation rule. It will be characterized by a vector of constants $\underline{w} = (\underline{w}_1, \dots, \underline{w}_m)$. Constant \underline{w}_j can be interpreted as the opportunity cost of allocating the j th contract.
- Step 2 For a fixed vector of probabilities of delayed service $\hat{d} = (\hat{d}_1, \dots, \hat{d}_m)$ we show how to fulfill these promises with these probabilities optimally. In particular, each optimal contract \hat{D}_j minimizes the distortion from not having one extra object j available.
- Step 3 We derive the condition for optimal probability of delayed service for j th contract. To do this, we change the decision variables from probabilities of delay \hat{d} to maximum distortions \bar{b} and find the first order condition with respect to this. The optimum balances the revenue extracted from the average person who gets the contract and the marginal increase in the distortion from giving this contract.
- Step 4 We argue then that we have characterized the optimal contracts. At this point we have derived two equations and of two variables that each contract j has to satisfy. The two variables are opportunity cost of allocating j th contract, \underline{w}_j , and the maximum distortion characterizing the contract j , \bar{b}_j . The equations are the corresponding first order conditions. We argue that each contract is well defined, either it is in the interior or in the upper corner. The upper bound for \bar{b}_j is either $1 - \delta$ or commitment constant c .
- Step 5 Finally, the existence of $\pi(d)$ functions can be verified using Blackwell's sufficient conditions and the Contraction mapping theorem.

3.5 Fully optimal mechanism

In this section again, a state is a vector $D = (D_1, \dots, D_N)$, where D_k denotes the realizations in which the seller has to serve a particular buyer corresponding to k from the previous period. Promises of delayed service are again denoted by $\hat{D} = (\hat{D}_1, \dots, \hat{D}_N)$, but now this vector can depend on the announcements of current buyers.

The optimal mechanism has to fulfill all previous promises and assigns the remaining instant service slots to buyers with highest positive virtual values. If there still remains one or more buyers with positive surplus, the seller must decide what to promise. The result states that

the promises they get are ordered in terms of both probabilities and sets. The highest of the remaining buyers is promised a probability $\hat{d}_1 = Pr(\hat{D}_1)$ that depends only on her valuation and is increasing in her type.

Now, the seller would want to assign all the possible delayed promises to the same buyer, since she is the remaining buyer with the highest value, but she cannot use more than one service slot next period. Therefore the seller reserves the first slot for her and uses this first slot to serve only her or new buyers. The next slot is matched with the next highest remaining buyer with the positive valuation—he will receive a promise that is increasing in his type, but since his value is lower than the first buyer’s value, this probability is lower. The assignment continues until all m future slots are matched with buyers (some of whom may receive probability zero as a promise). The optimal way to fulfill the promise is such that the difference between maximized expected revenue when this slot is available relative to when it is unavailable is minimized.

Theorem 3.6. *The optimal unrestricted mechanism with partial commitment is such that at each state $D = (D_1, \dots, D_N)$, when n buyers arrive and announce values v , then*

(i) $z = \#\{i : (n, v) \in D_i\}$ waiting buyers and $m - z$ new buyers with highest positive virtual surpluses are served instantly.

(ii) The next m highest buyers with positive virtual surpluses are offered delayed services with sets $\hat{D}(v) = (\hat{D}_1(v_{n+1-(m-j)-1:n}), \hat{D}_2(v_{n+1-(m-j)-2:n}), \dots, \hat{D}_m(v_{n+1-(m-j)-m:n}))$,

where contracts $\hat{D}(v)$ are characterized as follows.

$$\begin{aligned} \hat{D}_j(v_i) &= \{(n, v) : \pi^{j-1}(n, v) - \pi^j(n, v) \leq \bar{b}(w_i)\} \\ \pi^j(n, v) &= (1 - \delta) \sum_{i=1}^j \max\{0, w(v_{n+1-i:n})\} + \max_d \left\{ (1 - \delta)\beta \sum_{i=1}^m \hat{d}_i w(v_{n+1-(m-j)-i:n}) + \delta\pi(\hat{d}) \right\}, \\ \pi(d) &= E_{n,v}[\pi^0] - \sum_{j=1}^m E_{n,v} \left[\mathbf{1}[\pi^{j-1}(n, w) - \pi^j(n, w) \leq \bar{b}_j] [\pi^{j-1}(n, w) - \pi^j(n, w)] \right], \\ \bar{b}(w_i) &= \begin{cases} 0 & \forall w_i \leq 0, \\ \frac{(1-\delta)\beta}{\delta} w_i & \forall 0 < w_i < \frac{\delta}{(1-\delta)\beta} \min\{1 - \delta, c\}, \\ \min\{1 - \delta, c\} & \forall w_i \geq \frac{\delta}{(1-\delta)\beta} \min\{1 - \delta, c\}. \end{cases} \end{aligned}$$

Proof is in Appendix A.3. The structure is the same as in the illustrative example and parts of it are analogous to the waiting list. Proof structure:

Step 0 We assume $\pi(d)$ is well-defined function and make the same immediate observations regarding the allocation rule and the optimal contracts as in the waiting list case.

Step 1 We derive the optimal allocation rule. In this case optimal delayed services are personalized, so we get first-order conditions relating probabilities of service \hat{d}_k and buyer’s virtual valuation who gets the promise.

Step 2 Next we describe how to fulfill promises optimally. The solution minimizes distortions as in the waiting list case.

Step 3 We can now analyze the first-order conditions we got in Step 1. We show that the optimal probability of delayed service depends only of buyer's own virtual surplus and is strictly increasing in the interior.

Step 4 Finally to prove the existence of $\pi(d)$ we use the Contraction mapping theorem.

4 Discussion

I studied the problem of selling service slots. From previous literature we already knew the solution to the problem in two special cases. First, when the capacity is high, the optimal mechanism is a constant posted price mechanism. Second, when the buyers' discount rate and exit rate are constant, the fully optimal allocation rule always assigns the objects to the highest-valuing¹⁹ buyers among all remaining buyers. This paper shows how to solve the optimal mechanism design problem in the remaining case, where rationing is necessary and buyers are heterogeneous not only in their valuations for service, but also in their exit rates, so that the buyers from different periods should be treated differently. This introduces a new set of trade-offs which are balanced optimally by a new type of auction mechanism.

In addition to the fully optimal mechanism, we analyzed two classes of restricted mechanisms where the possible contracts of delayed service are limited: pure calendar mechanisms and waiting list mechanisms. The optimal pure calendar mechanism was a simple mechanism, characterized by a vector of opportunity costs. The qualitative properties of waiting lists were similar to the fully optimal mechanisms. Finally, we showed how the optimal mechanisms change in all three classes when the seller has less than perfect commitment power.

We made several simplifying assumptions that could be relaxed without changing the conclusions. We assumed that the arrivals of buyers are observable both for the other buyers and for the seller. This assumption is without loss in the sense that if buyers had a chance to choose whether to enter the mechanism at their arrival or later—without knowing anything about the other buyers—they would always enter immediately.

It would also be relatively straightforward to introduce more general arrival processes. We assumed that the number of objects m is constant and the number of buyers n is an independently and identically distributed random variable. We could assume instead that m and n are independent draws from some joint distribution, where the distribution can be different each period. This is a generalization that includes, for example, seasonal fluctuations in demand and supply.

In real life applications the most common mechanisms to allocate service are posted price mechanisms. As we argued, whenever there is a need for rationing, the optimal mechanism would

¹⁹If the goal is to maximize revenue, then the mechanism assigns service only to buyers with positive virtual surplus.

involve an auction. The assumption that the seller has to use a posted price mechanism would be another restriction to the allocation rule. Some related papers²⁰ have shown that if the arrival process is continuous, then the optimal mechanism would be implementable with posted prices, since two buyers never arrive at the same time. The same could be done with the model here, by assuming arrival in continuous time and leaving the production unchanged. Alternatively, we could assume that buyers in one period arrive sequentially and the seller has to assign the allocation to each buyer at the moment of their arrival, before observing the other new buyers' arrivals and their types.

Perhaps the most restrictive assumption was concentrating on contracts only one period ahead. A generalization of current models would involve a common discounting process β , where payoff from being served eventually at date s_i is $\beta(s_i)v_i - p_i$, such that $\beta(s_i)$ is decreasing in s_i and $\lim_s \beta(s_i)/\delta^{s_i} \rightarrow 0$. Special cases of this model would be $\beta(s_i) = \beta^{s_i}$ which follows from Said (2008) and $\beta(0) = 1, \beta(1) = \beta, \beta(s_i) = 0$, which was studied in this paper.

References

- BERGEMANN, D., AND J. VÄLIMÄKI (2010): "The Dynamic Pivot Mechanism," *Econometrica*, 78(2), 771–789.
- BOARD, S., AND A. SKRZYPACZ (2010): "Optimal dynamic auctions for durable goods: Posted prices and fire-sales," Discussion paper, Stanford University.
- GERSHKOV, A., AND B. MOLDOVANU (2009): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*, 1(2), 168–198.
- HART, O. D., AND J. TIROLE (1988): "Contract Renegotiation and Coasian Dynamics," *The Review of Economic Studies*, 55(4).
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–73.
- PAI, M., AND R. VOHRA (2008): "Optimal Dynamic Auctions," Discussion Papers 1461, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2011): "Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure," Discussion paper, Northwestern University.
- RILEY, J. G., AND W. F. SAMUELSON (1981): "Optimal Auctions," *The American Economic Review*, 71(3), 381–392.

²⁰See Gershkov and Moldovanu (2009) and Board and Skrzypacz (2010) for example.

SAID, M. (2008): “Auctions with Dynamic Populations: Efficiency and Revenue Maximization,” *Social Science Research Network Working Paper Series*.

STOKEY, N. L. (1979): “Intertemporal Price Discrimination,” *The Quarterly Journal of Economics*, 93(3), 355–371.

Appendices

A Proofs

To shorten the notation, for a given realization (n, v) we denote the k th lowest virtual surplus by $w_k = w(v_{k:n})$, so that w is always the ordered virtual surplus vector corresponding to (n, v) . We denote its distribution function by F_k , so that $F_k(w) = Pr(w_k \leq w) = Pr(v_{k:n} \leq w^{-1}(w)) = F_{k:n}(w^{-1}(w))$.

A.1 Calendar mechanism

Proof of Theorem 3.1 Remember that state $z \in \{0, \dots, m\}$ denotes how many service slots are already promised to buyers in the previous period. Then at state z the seller can serve up to $m - z$ new buyers instantly. Let $\pi(z)$ denote maximum revenue from state z . It is clear that $\pi(z)$ is strictly decreasing in z .

We will start by assuming that c is large enough so that the seller can commit to all promises. Fix state z , number of buyers n and the corresponding ordered virtual surplus vector w . Let $\{I, D, R\}$ be a partition of buyers $\{1, \dots, n\}$, such that I is the set of buyers who will be served instantly, D receive delayed service, and R are refused service. Feasibility requires that $\#I \leq z$, $\#R \leq m$. We can express the maximum revenue from state z and realization (n, w) by

$$\pi^z(n, w) = (1 - \delta) \sum_{i \in I} w_i + (1 - \delta) \sum_{i \in D} \beta w_i + \delta \pi(\#D)$$

We can make some instant observations. First, when $w_k \leq 0$, buyer k will not get the object, since refusing him service would be feasible and would increase instant revenue without decreasing continuation payoff. Second, if $z + \#I < m$, so that some of the current objects remain unallocated, it must be because there were exactly $\#I$ buyers with $w_i > 0$. This is true, since by replacing the policy that assigns $w_{n-(\#I+1)}$ current object, the seller would strictly increase flow payoff while still satisfying feasibility and not decreasing continuation payoff. Finally, it must be the case that if $i \in I, d \in D$, and $r \in R$, then $w_i \geq w_d \geq w_r$. That is, we can express I, D , and R as $I = \{n, \dots, n + 1 - \#I\}$, $D = \{n - \#I, \dots, n - \#I + 1 - \#D\}$, $R = \{n - \#I - \#D, \dots, n - \#I - \#D + 1 - \#R\}$. If this is not true, the seller could swap the service dates of the buyers for whom the order does not hold, thus preserving feasibility and continuation payoffs, but increasing flow revenue.

This means that the optimal policy such that up to $m - z$ buyers with highest positive virtual surpluses will receive the current service. If there are less than $m - z$ buyers with positive surpluses, the rest of the buyers are refused service and the continuation state will be m (and $k = 0$).

Consider the situation where there are more than $m - z$ buyers with strictly positive virtual surpluses, so that $w_{n-(m-z)} > 0$. Then $\#I = \{n, \dots, n + 1 - (m - z)\}$ and the seller chooses optimal number k of buyers with highest positive virtual surpluses among remaining buyers to whom to offer delayed service, so that $D = \{n - (m - z), \dots, n + 1 - (m - z) - k\}$.

Denote the revenue difference between having all m service slots available and having one slot unavailable by $\Delta_1 = \pi(0) - \pi(1) > 0$. Similarly, let Δ_k denote the revenue decrease from having k th slot unavailable, so that $\Delta_k = \pi(k - 1) - \pi(k)$. Analogously to the illustrative example, the normalized discounted version of this difference is the opportunity cost, which we denote by $\underline{w}_k = \frac{\delta}{1-\delta} \frac{\Delta_k}{\beta}$. Now, k is optimal if it maximizes

$$\pi^z(n, w) = (1 - \delta) \sum_{i=n+1-(m-z)}^n w_i + (1 - \delta)\beta \sum_{i=1}^k w_{n+1-(m-z)-i} + \delta\pi(k).$$

Since $\pi(k) = \pi(0) - \pi(0) + \pi(1) - \dots + \pi(k) = \pi(0) - \sum_{i=1}^k \Delta_i$, we can rewrite the objective as

$$\pi^z(n, w) = (1 - \delta) \sum_{i=n+1-(m-z)}^n w_i + \delta\pi(0) + (1 - \delta)\beta \sum_{i=1}^k [w_{n+1-(m-z)-i} - \underline{w}_i]. \quad (\text{A.1})$$

The first part, $(1 - \delta) \sum_{i=n+1-(m-z)}^n w_i + \delta\pi(0)$, is the maximum static revenue that the seller can achieve, whereas the last part is the dynamic revenue, which in optimum is always non-negative.

To compute the expected revenue in state z , we have to take expectation of $\pi^z(n, w)$ with respect to the number of buyers n and valuations v . This gives

$$\pi(z) = E_n \left[(1 - \delta) \sum_{i=1}^{m-z} \int_0^1 w dF_{n+1-i}(w) + \delta\pi(0) + (1 - \delta)\beta E_v W((w_{n-(m-z)}, \dots, w_{n+1-(m-z)-m}), \underline{w}) \right],$$

where $W((w_{n-(m-z)}, \dots, w_{n+1-(m-z)-m}), \underline{w}) = \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k [w_{n+1-(m-z)-i} - \underline{w}_i]$. Therefore for each $j \in \{1, \dots, m\}$,

$$\underline{w}_j = \frac{\delta}{1 - \delta} \frac{\pi(j - 1) - \pi(j)}{\beta} = \frac{\delta}{1 - \delta} \frac{E_{n,w}[\pi^{j-1}(n, w) - \pi^j(n, w)]}{\beta} = \Psi_j(\underline{w}),$$

where

$$\begin{aligned} \Psi_j(\underline{w}) = & \frac{\delta}{\beta} E_n \left[\int_0^1 w dF_{n-(m-j)}(w) \right. \\ & \left. + \beta E_v \left(W((w_{n-(m-j)}, \dots, w_{n+1-(m-j)-m}), \underline{w}) - W((w_{n-(m-j)-1}, \dots, w_{n-(m-j)-m}), \underline{w}) \right) \right] \end{aligned}$$

To show that vector $\underline{w} \in [0, 1]^m$ exists and is uniquely defined, we show that Ψ is continuous and strictly increasing and $\Psi(\underline{w}) \in [0, 1]^m$ for all $\underline{w} \in [0, 1]^m$. Then \underline{w} exists by Brouwer's fixed point theorem and is unique because of strict monotonicity. We will need some properties of W functions that are proven separately in Lemmas A.1 to A.3.

Continuity comes directly from the fact that expectation operator is continuous and the function W is continuous (Lemma A.1). For boundedness we look at $\underline{w} = \mathbf{1} = (1, \dots, 1)$ and $\underline{w} = \mathbf{0} = (0, \dots, 0)$.

At $\underline{w} = \mathbf{0}$, $\max_{k \in \{0, \dots, m\}} \sum_{i=1}^k w_{n+1-(m-j)-i}$ includes $w_{n+1-(m-j)-i}$ if and only if $w_{n+1-(m-j)-i} > 0$. Therefore²¹

$$E_v[W((w_{n-(m-j)}, \dots, w_{n+1-(m-j)-m}))] = \sum_{i=1}^m \int_0^1 [1 - F_{n+1-(m-j)-i}(w)] dw.$$

We get an analogous expression for the other W term. Plugging this into Ψ_i and rearranging gives

$$\Psi_j(\mathbf{0}) = \frac{\delta}{\beta} E_n \left[1 - \int_0^1 \left[(1 - \beta) F_{n-(m-j)}(w) + \beta F_{n-(m-j)-m}(w) \right] dw \right] \geq 0.$$

Next, at $\underline{w} = \mathbf{1}$ note that $[w_k - 1] \leq 0$ for any $k > 0$, therefore $W(\cdot, \mathbf{1}) = 0$. Therefore

$$\Psi_j(\mathbf{1}) = \frac{\delta}{\beta} E_n \left[\int_0^1 [1 - F_{n-(m-j)}(w)] dw \right] \leq \frac{\delta}{\beta} E_n \left[\int_0^1 [1 - F_n(w)] dw \right] \leq 1,$$

where the first inequality comes from the stochastic dominance of higher order statistics and the second inequality from the assumptions. When $\frac{\delta}{\beta} E_n \left[\int_0^1 [1 - F_n(w)] dw \right] > 1$, not only is the existence not guaranteed, but also optimal calendar mechanism with positive probability of delay cannot exist, since at any (n, v, z) , the surplus extracted from delay is less than the expected value from having one more object available next period. Therefore in this case the optimal mechanism would be the optimal static mechanism.

Finally, strict monotonicity of $\Psi_j(\underline{w})$ with respect to \underline{w} follows from Lemma A.3. Therefore $\Psi_j(\underline{w}) \in [\Psi_j(\mathbf{0}), \Psi_j(\mathbf{1})] \subset [0, 1]$ and so we have continuous $\Psi : [0, 1]^m \rightarrow [0, 1]^m$ which means that by Brouwer's fixed point theorem \underline{w} exists and by strict monotonicity it is unique.

To complete the proof we have to verify that when $c \geq 1 - \delta$, the optimal mechanism described above satisfies commitment conditions. Fix any state $z < m$ (at state m there are obviously no commitment issues), realization (n, w) . By fulfilling all promises, the seller gets revenue $\pi^z(n, w)$, whereas leaving a promise to one previous customer unfulfilled will give one additional instant service spot, which therefore gives revenue $\pi^{j-1}(n, w)$. Therefore the cost of fulfilling the promise is

$$\pi^{j-1}(n, w) - \pi^j(n, w) = (1 - \delta)w_{n-(m-z)}$$

²¹Integration by parts gives $\int_0^1 w dF_k(w) = \int_0^1 [1 - F_k(w)] dw$.

$$+(1 - \delta)\beta \left[\max_k \sum_{i=1}^k [w_{n+1-(m-z)-i} - \underline{w}_i] - \max_k \sum_{i=1}^k [w_{n-(m-z)-i} - \underline{w}_i] \right].$$

By Lemma A.1, $\max_k \sum_{i=1}^k [w_{n+1-(m-z)-i} - \underline{w}_i] \leq \max_k \sum_{i=1}^k [w_{n-(m-z)-i} - \underline{w}_i]$, and by assumptions $w_{n-(m-z)} \leq w(1) = 1$, therefore maximum possible value for $\pi^{j-1}(n, w) - \pi^j(n, w)$ is $1 - \delta$. This means that $c \geq 1 - \delta$ is sufficient for the commitment.

It is also necessary if at each z there exists (n, w) such that $\pi^{j-1}(n, w) - \pi^j(n, w)$ achieves value $1 - \delta$. When $N \geq 2m$, we have that $n \geq z + m$ at each $z \in \{0, \dots, m\}$, and therefore a possible realization is such that $N = 2m$, $w_{n-(m-z)} = \dots = w_{n-(m-z)-m} = 1$, which gives the equality. When $N \in \{m + 1, \dots, 2m - 1\}$ the necessary condition (especially for low states) is even weaker and needs to be computed directly from the expression above. \square

Lemma A.1. *Let $W(w, \underline{w}) = \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k [w_i - \underline{w}_i]$ for all $w, \underline{w} \in [0, 1]^m$ such that $w_1 \geq w_2 \geq \dots \geq w_m$. Let $k^*(w, \underline{w}) = \min \arg \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k [w_i - \underline{w}_i]$. Then*

- (i) W continuous, non-decreasing in w , non-increasing in \underline{w} , and convex.
- (ii) k^* is non-decreasing in w and non-increasing in \underline{w} .

Remark: k^* is defined as the smallest maximizer simply to avoid the complications that arise when there are multiple maximizers and the changes do not affect the maximum. In this case the set of maximizers may change (monotonically), but the maximum does not. Since we are not interested in characterizing all possible tie-breaking cases, we will pick just the smallest one.

Proof Since W is the maximum of linear non-decreasing functions of $(w, -\underline{w})$, it is continuous, non-decreasing in each argument (thus non-increasing in \underline{w}), and convex.

Suppose k^* is not non-decreasing in $(w, -\underline{w})$. Then there exist (w, \underline{w}) and $(\hat{w}, \hat{\underline{w}})$ such that $\hat{w} \geq w$ and $-\hat{\underline{w}} \geq -\underline{w}$, but $k^*(\hat{w}, \hat{\underline{w}}) < k^*(w, \underline{w})$. Optimality of $k^*(w, \underline{w})$ and $k^*(\hat{w}, \hat{\underline{w}})$ requires that

$$\sum_{i=1}^{k^*(\hat{w}, \hat{\underline{w}})} [\hat{w}_i - \hat{\underline{w}}_i] > \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \underline{w}_i] \Rightarrow 0 \geq \sum_{i=k^*(\hat{w}, \hat{\underline{w}})}^{k^*(w, \underline{w})} [w_i - \underline{w}_i]$$

$$\sum_{i=1}^{k^*(w, \underline{w})} [w_i - \underline{w}_i] > \sum_{i=1}^{k^*(\hat{w}, \hat{\underline{w}})} [w_i - \underline{w}_i] \Rightarrow \sum_{i=k^*(\hat{w}, \hat{\underline{w}})}^{k^*(w, \underline{w})} [w_i - \underline{w}_i] > 0.$$

Since $[\hat{w}_i - \hat{\underline{w}}_i] \geq [w_i - \underline{w}_i]$ for all i , this is a contradiction. \square

Lemma A.2. *Let $W(w, \underline{w}) = \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k [w_i - \underline{w}_i]$ for all $w, \underline{w} \in [0, 1]^m$ such that $w_1 \geq w_2 \geq \dots \geq w_m$. Fix $w_0 \geq w_1 \geq \dots \geq w_m$ and denote $w = (w_1, \dots, w_m)$ and $\hat{w} = (w_0, w_1, \dots, w_{m-1})$. Fix $\underline{w} \in [0, 1]^m$, $l \in \{1, \dots, m\}$, $\varepsilon > 0$ and let $\hat{\underline{w}}$ be such that $\hat{\underline{w}}_l = \underline{w}_l + \varepsilon$ and $\hat{\underline{w}}_i = \underline{w}_i$ for all $i \neq l$. Then*

$$W(w, \hat{\underline{w}}) - W(\hat{w}, \hat{\underline{w}}) \geq W(w, \underline{w}) - W(\hat{w}, \underline{w}).$$

Proof First note that by Lemma A.2 $k^*(w, \underline{w}) \leq k^*(\hat{w}, \underline{w})$ and $k^*(w, \underline{\hat{w}}) \leq k^*(\hat{w}, \underline{\hat{w}})$. Moreover, by the same result $k^*(w, \underline{\hat{w}}) \leq k^*(w, \underline{w})$ and $k^*(\hat{w}, \underline{\hat{w}}) \leq k^*(w, \underline{\hat{w}})$.

Suppose $l > k^*(w, \underline{w})$. Then, since the relevant terms in the summation are not affected by ε , we have $W(w, \underline{\hat{w}}) = W(w, \underline{w})$. Formally,

$$W(w, \underline{\hat{w}}) \geq \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \hat{w}_i] = \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \underline{w}_i] = W(w, \underline{w}) \geq W(w, \underline{\hat{w}}).$$

Since we have $W(\hat{w}, \underline{\hat{w}}) \leq W(\hat{w}, \underline{w})$, Lemma A.1 implies that $W(w, \underline{\hat{w}}) - W(\hat{w}, \underline{\hat{w}}) \geq W(w, \underline{w}) - W(\hat{w}, \underline{w})$.

Suppose now that $l \leq k^*(w, \underline{w})$. Then $l \leq k^*(\hat{w}, \underline{w})$. Now observe that either $k^*(w, \underline{\hat{w}}) = k^*(w, \underline{w})$ or $k^*(w, \underline{\hat{w}}) < l$, because if this were not the case, it would imply that

$$W(w, \underline{\hat{w}}) = \sum_{i=1}^{k^*(w, \underline{\hat{w}})} [w_i - \hat{w}_i] \geq \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \hat{w}_i] \Rightarrow \sum_{i=1}^{k^*(w, \underline{\hat{w}})} [w_i - \underline{w}_i] - \varepsilon \geq \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \underline{w}_i] - \varepsilon,$$

but this contradicts the optimality of $k^*(w, \underline{w})$ at (w, \underline{w}) . Similarly either $k^*(\hat{w}, \underline{\hat{w}}) = k^*(\hat{w}, \underline{w})$ or $k^*(\hat{w}, \underline{\hat{w}}) < l$. This means that we have four remaining potential cases to consider.

1. When $k^*(w, \underline{\hat{w}}) = k^*(w, \underline{w})$, $k^*(\hat{w}, \underline{\hat{w}}) = k^*(\hat{w}, \underline{w})$, it implies that $W(w, \underline{\hat{w}}) = W(w, \underline{w}) - \varepsilon$ and $W(\hat{w}, \underline{\hat{w}}) = W(\hat{w}, \underline{w}) - \varepsilon$, so $W(w, \underline{\hat{w}}) - W(\hat{w}, \underline{\hat{w}}) = W(w, \underline{w}) - W(\hat{w}, \underline{w})$.
2. When $k^*(w, \underline{\hat{w}}) < l$, $k^*(\hat{w}, \underline{\hat{w}}) = k^*(\hat{w}, \underline{w})$, we know that $W(\hat{w}, \underline{\hat{w}}) = W(\hat{w}, \underline{w}) - \varepsilon$ and

$$W(w, \underline{\hat{w}}) \geq \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \hat{w}_i] = \sum_{i=1}^{k^*(w, \underline{w})} [w_i - \underline{w}_i] - \varepsilon = W(w, \underline{w}) - \varepsilon.$$

Moreover, the inequality has to be strict, since otherwise $k^*(w, \underline{\hat{w}})$ would also be optimal in (w, \underline{w}) . Therefore $W(w, \underline{\hat{w}}) - W(\hat{w}, \underline{\hat{w}}) > W(w, \underline{w}) - \varepsilon - W(\hat{w}, \underline{w}) + \varepsilon = W(w, \underline{w}) - W(\hat{w}, \underline{w})$.

3. Case $k^*(w, \underline{\hat{w}}) = k^*(w, \underline{w})$, $k^*(\hat{w}, \underline{\hat{w}}) < l$ is not possible, since it would mean that $k^*(w, \underline{\hat{w}}) = k^*(w, \underline{w}) \geq l > k^*(\hat{w}, \underline{\hat{w}}) \geq k^*(w, \underline{\hat{w}})$.
4. Finally, when $k^*(w, \underline{\hat{w}}) < l$, $k^*(\hat{w}, \underline{\hat{w}}) < l$, we can write

$$\begin{aligned} & [W(w, \underline{\hat{w}}) - W(\hat{w}, \underline{\hat{w}})] - [W(w, \underline{w}) - W(\hat{w}, \underline{w})] \\ &= - \sum_{i=k^*(w, \underline{\hat{w}})+1}^{k^*(\hat{w}, \underline{\hat{w}})} [w_i - \underline{w}_i] + \sum_{i=k^*(\hat{w}, \underline{\hat{w}})+1}^{k^*(w, \underline{w})} [\hat{w}_i - w_i] + \sum_{i=k^*(w, \underline{w})+1}^{k^*(\hat{w}, \underline{w})} [\hat{w}_i - \underline{w}_i] > 0, \end{aligned}$$

because (1) $\sum_{i=k^*(w, \underline{\hat{w}})+1}^{k^*(\hat{w}, \underline{\hat{w}})} [w_i - \underline{w}_i] \leq 0$, since otherwise $k^*(w, \underline{\hat{w}})$ would not be optimal at $(w, \underline{\hat{w}})$, (2) $\hat{w}_i \geq w_i$ for each i , (3) $\sum_{i=k^*(w, \underline{w})+1}^{k^*(\hat{w}, \underline{w})} [\hat{w}_i - \underline{w}_i] > 0$, since otherwise $k^*(\hat{w}, \underline{w})$ would not be optimal at (\hat{w}, \underline{w}) .

We have shown that the weak inequality holds for all possible arguments. \square

Lemma A.3. Let $W(w, \underline{w}) = \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k [w_i - \underline{w}_i]$ for all $w, \underline{w} \in [0, 1]^m$ such that $w_1 \geq w_2 \geq \dots \geq w_m$. Let $w_{n-j+1} \geq \dots \geq w_{n-j+1-m}$ be generated so that $\Pr(n = \hat{n}) = \gamma_{\hat{n}}$, $w_k = w(v_{k:n})$ when $k > 0$ and $w_k = 0$ otherwise, and $\Pr(v_{k:n} \leq \hat{v}) = F_{k:n}(v)$. Fix $\underline{w} \in (0, 1)^m$, $l \in \{1, \dots, m\}$, $\varepsilon \in (0, 1 - \underline{w}_l)$ and let $\hat{\underline{w}}$ be such that $\hat{\underline{w}}_l = \underline{w}_l + \varepsilon$ and $\hat{\underline{w}}_i = \underline{w}_i$ for all $i \neq l$. Denote $\underline{w} = (\underline{w}_m, \dots, \underline{w}_1)$, $w = (w_{n-j}, \dots, w_{n-j+1-m})$, $\hat{w} = (\hat{w}_{n-(j-1)}, \dots, \hat{w}_{n-(j-1)-m})$. Then

$$E_{n,v}[W(w, \hat{\underline{w}}) - W(\hat{w}, \hat{\underline{w}})] > E_{n,v}[W(w, \underline{w}) - W(\hat{w}, \underline{w})].$$

Proof Lemma A.2 gives weak inequality for any realization of (n, v) . Moreover, the proof shows that whenever $k^*(w, \hat{\underline{w}}) < l \leq k^*(w, \underline{w})$, the inequality is strict. Therefore it suffices to construct a positive measure set of w_i 's, where $k^*(w, \hat{\underline{w}}) < l \leq k^*(w, \underline{w})$. We construct the set as follows:

- (i) $w_{n-j+1} \geq \dots \geq w_{n-j+1-(l-1)} > \max\{\underline{w}_m, \dots, \underline{w}_{l+1}\}$
- (ii) $w_{n-j+1-l} \in (\underline{w}_l, \underline{w}_l + \varepsilon)$
- (iii) $w_{n-j+1-m} \leq \dots \leq w_{n-j+1-l-1} < \min\{\underline{w}_1, \dots, \underline{w}_{l-1}\}$

A combination of w_i 's satisfying these restrictions occur with strictly positive probability. By construction,

- $[w_{n-j+1-i} - \underline{w}_{m+1-i}] > 0$ for all $i \leq l$ and $[w_{n-j+1-i} - \underline{w}_{m+1-i}] < 0$ for all $i > l$, therefore $k^*(w, \underline{w}) = l$.
- $[w_{n-j+1-i} - \hat{\underline{w}}_{m+1-i}] > 0$ for all $i < l$ and $[w_{n-j+1-i} - \hat{\underline{w}}_{m+1-i}] < 0$ for all $i \geq l$, therefore $k^*(w, \hat{\underline{w}}) = l - 1 < l$.

Therefore under all these realizations the inequality is indeed strict. \square

Proof of Corollary 3.3 We apply Theorem 3.1 to get \underline{w} and the general structure of the optimal mechanism. Suppose $\underline{w}_1 \leq \dots \leq \underline{w}_m$. The claim is that $(n + 1 - (m - z) - i)$ th highest customer with virtual surplus $w_{n+1-(m-z)-i}$ is offered delayed service if and only if $w_{n+1-(m-z)-i} > \underline{w}_i$. To see this, let

$$k^* = \arg \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k [w_{n+1-(m-z)-i} - \underline{w}_i],$$

$$\hat{k} = \max \left\{ k \in \{0, \dots, m\} : [w_{n+1-(m-z)-k} - \underline{w}_k] > 0 \right\}.$$

For each $i \leq \hat{k}$ we have $w_{n+1-(m-z)-i} \geq w_{n+1-(m-z)-\hat{k}}$ and $\underline{w}_i \leq \underline{w}_{\hat{k}}$, so $[w_{n+1-(m-z)-i} - \underline{w}_i] \geq [w_{n+1-(m-z)-\hat{k}} - \underline{w}_{\hat{k}}] \geq 0$. For each $i > \hat{k}$ we have by definition that $[w_{n+1-(m-z)-i} - \underline{w}_i] \leq 0$. Therefore we must have $k^* = \hat{k}$. \square

A.2 Waiting list

Proof of Theorem 3.5 By assumptions, allocation assigns each buyer i either instant service, refusal of service, or promise of delayed service at some set \hat{D}_i , so that this buyer will be

served next period if and only if $(n, v) \in \hat{D}_i$. The waiting list assumption says that vector $\hat{D} = (\hat{D}_1, \dots, \hat{D}_N)$ must be fixed before observing (n, v) . Feasibility requires that the total number of buyers served in the current period must be at most m , and promises of delayed service are such that $\#\{i : (n, v) \in \hat{D}_i\} \leq m$ for each (n, v) . Remember that we denote state by $D = (D_1, \dots, D_N)$, the vector of contracts by $\hat{D} = (\hat{D}_1, \dots, \hat{D}_N)$, and the corresponding probabilities $d = (d_1, \dots, d_N)$, $\hat{d} = (\hat{d}_1, \dots, \hat{d}_N)$ st $d_i = Pr(D_i)$ and $\hat{d}_i = Pr(\hat{D}_i)$.

Let $\pi(D)$ denote the maximized revenue at state D and with some abuse of notation let $\pi(d)$ be the maximized revenue for given probabilities d . That is, $\pi(d) = \max_{D: Pr(D_i) \geq d_i} \pi(D)$. We assume that the functions $\pi(d)$ are well-defined and will verify it later.

After the standard manipulations we can express the revenue recursively as

$$\pi(d) = \max E_{n,v} \left[(1 - \delta) \sum_{i=1}^n \phi_i w_i + \delta \pi(\hat{d}) \right], \quad (\text{A.2})$$

where $\phi_i = 1$ if i is assigned instant service, $\phi_i = 0$ if refused service, and $\phi_i = \beta \hat{d}_i$ when promised delayed service with probability \hat{d}_i .

We can make some immediate observations. It is clear that giving instant service or promising delayed service to buyers with negative virtual surpluses is dominated by not serving them. Moreover, it is optimal to serve higher-valued buyers before lower valued buyers. That is, if buyer i is served, buyers j and j' are offered delayed service with respective sets \hat{D}_j and $\hat{D}_{j'}$ such that $\hat{d}_j > \hat{d}_{j'}$, and r is refused service, then it must be that $w_i \geq w_j \geq w_{j'} \geq w_r$, since otherwise the seller could strictly increase revenue by swapping promises given to these buyers. Finally, the seller will never assign positive probability of delayed service to more than m buyers. Then we could take the lowest-valued buyer who receives positive probability of service and reassign his probability to higher-valued buyers without changing the feasibility or continuation states, but increasing the flow revenue.

The final immediate observation is that the monotonicity of promises also holds in terms of sets. In particular, contracts are ordered so that²² $\hat{D}_1 \supset \dots \supset \hat{D}_m$. This is a priority ranking—the keeper of the first contract will always be served before the keeper of the second contract. If the ordering were not true for some j , the seller could replace promises for \hat{D}_{j-1} and \hat{D}_j by new promises $\hat{D}'_{j-1} = \hat{D}_{j-1} \cup \hat{D}_j$ and $\hat{D}'_j = \hat{D}_j - 1 \cup \hat{D}_j$ while keeping other contracts unchanged. The new contracts do not violate feasibility but extract more revenue from the current buyers.

Step 1: allocation rule. Suppose first that the contracts \hat{D} are already optimally designed. Fix realization (n, v) and the corresponding ordered virtual surplus vector w and denote the number of slots unavailable by $z = \#\{i : (n, v) \in D_i\}$. By the observations above, it is clear that $m - z$ highest-valuing buyers with strictly positive surplus will receive instant service and buyers with negative virtual surplus are never served. If there are more than $m - z$ buyers with strictly positive virtual surplus, k highest of them will receive contracts $\hat{D}_1, \dots, \hat{D}_k$. The remaining question is the choice of optimal $k \in \{0, \dots, m\}$.

²²More precisely this relation must hold with probability one, that is $Pr(\hat{D}_j \setminus \hat{D}_{j-1}) = 0$ for all j .

Let Δ_1 be the difference between revenue when no promises are given compared to revenue when one buyer received promise \hat{D}_1 , that is, $\Delta_1 = \pi(0, \dots, 0) - \pi(\hat{D}_1, 0, \dots, 0)$. More generally, let the revenue effect of having assigned the contract k by

$$\Delta_k = \pi(\hat{D}_1, \dots, \hat{D}_{k-1}, 0, \dots, 0) - \pi(\hat{D}_1, \dots, \hat{D}_k, 0, \dots, 0).$$

Again, it will serve as the opportunity cost of allocating k th contract \hat{D}_k in the allocation problem. Define $\underline{w}_k = \frac{\delta}{1-\delta} \frac{\Delta_k}{\beta \hat{d}_k}$.

Let $\pi^j(n, w)$ denote the maximum revenue when j instant service slots are unavailable at realization (n, w) . Then we can express it as

$$\pi^j(n, w) = (1-\delta) \sum_{i=1}^{m-j} \max\{0, w_{n+1-i}\} + \delta \pi(0, \dots, 0) + (1-\delta)\beta \max_{k \in \{0, \dots, m\}} \sum_{i=1}^k \hat{d}_i [w_{n+1-(m-j)-i} - \underline{w}_i].$$

Again, the sum of the first two terms is the maximal static revenue that the seller can extract without giving any promises of delayed service. The last term is the extra revenue extracted through delay and always non-negative.

Step 2: fulfilling promises optimally. The next question is how to choose particular contracts \hat{D} if for fixed probabilities of service \hat{d} . Using the fact that $\hat{D}_1 \supset \dots \supset \hat{D}_m$ we can express the expected revenue from a given \hat{d} as

$$\pi(\hat{d}) = \max_{\hat{D}: Pr(\hat{D}_i) \geq \hat{d}_i} \left\{ E_{n,v} \left[\mathbf{1}[\hat{D}_1^c] \pi^0(n, w) + \mathbf{1}[\hat{D}_1 \setminus \hat{D}_2] \pi^1(n, w) + \dots + \mathbf{1}[\hat{D}_m] \pi^m(n, w) \right] \right\}.$$

With some straightforward manipulations, we can rewrite it as

$$\pi(\hat{d}) = E_{n,v}[\pi^0(n, v)] - \sum_{j=1}^m \min_{\hat{D}_j: Pr(\hat{D}_j) \geq \hat{d}_j} E_{n,v} \left[\mathbf{1}[\hat{D}_j] [\pi^{j-1}(n, w) - \pi^j(n, w)] \right]. \quad (\text{A.3})$$

Note that the problem in (A.3) is separable in j , since both the objective and the constraint only depend on \hat{D}_j and \hat{d}_j . Now, the difference in revenues, $\pi^{j-1}(n, w) - \pi^j(n, w)$, is equal to 0 whenever there is less than $m - j$ new buyers with strictly positive virtual valuations, and is strictly increasing otherwise. That is, it is strictly positive and strictly increasing if and only if $w_{n-(m-j)} > 0$.

This means that we have two cases to consider. When (n, v) such that $w_{n-(m-j)} \leq 0$, then $\pi^{j-1}(n, w) - \pi^j(n, w) = 0$. Therefore if $\hat{d}_j < Pr(w_{n-(m-j)} \leq 0) = E_n[F_{n-(m-j)}(p^*)]$, then the problem is simple to solve, any subset of $\{(n, v) : w_{n-(m-j)} \leq 0\}$ works. However, it is not optimal, since by increasing the probability \hat{d}_j to $E_n[F_{n-(m-j)}(p^*)]$ the seller can increase the revenue while not violating the feasibility and not decreasing the continuation value.

Therefore the optimum must be in the second case where $\hat{d}_j \geq E_n[F_{n-(m-j)}(p^*)]$. In this case the contract set \hat{D}_j has to include some points where $w_{n-(m-j)} > 0$ so that $\pi^{j-1}(n, w) - \pi^j(n, w)$ is positive. Since then $\pi^{j-1}(n, w) - \pi^j(n, w)$ is strictly increasing in $w_{n-(m-j)}$, we get that each

contract \hat{D}_j includes exactly measure \hat{d}_j of points where $\pi^{j-1}(n, w) - \pi^j(n, w)$ is the smallest.

This means that we can characterize contract \hat{D}_j by the upper bound to the distortion $\pi^{j-1}(n, w) - \pi^j(n, w)$. We will denote the maximum distortion level by \bar{b}_j . Then

$$\hat{D}_j = \{(n, v) : \pi^{j-1}(n, w) - \pi^j(n, w) \leq \bar{b}_j\}$$

for uniquely determined maximum distortion level \bar{b}_j such that $Pr(\hat{D}_j) = \hat{d}_j$.

Step 3: optimal probabilities of delayed service. Since \bar{b}_j is strictly increasing with \hat{d}_j in the relevant region, there is a one-to-one mapping between $\hat{d}_j \in [E_n[F_{n-(m-j)}(p^*)], 1]$ and $\bar{b}_j \in [0, 1 - \delta]$.

The choice of optimal contracts takes the form

$$\max_{\bar{b}_j} \left\{ E_{n,v} \left[\mathbf{1}[\hat{D}_1^c] \pi^0(n, w) + \mathbf{1}[\hat{D}_1 \setminus \hat{D}_2] \pi^1(n, w) + \dots + \mathbf{1}[\hat{D}_m] \pi^m(n, w) \right] \right\},$$

subject to $0 \leq \bar{b}_j \leq 1 - \delta$ and $\bar{b}_j \leq c$ (commitment constraint). Notice that here is when the commitment level affects the decision problem. So far we analyzed allocation rule for a given contract and choice of contracts for a given probability (which was not necessarily possible under partial commitment). Now the choice of probability of delay or equivalently the maximum distortion \bar{b}_j must be such that when contract j is assigned, then under any realization (n, v) the maximum benefit from not fulfilling contract j is less than c , which is exactly what $\bar{b}_j \leq c$ guarantees.

We will now define two new distribution functions. First, $G_{b_j}(b_j) = Pr(\{(n, v) : \pi^{j-1}(n, w) - \pi^j(n, w) \leq b_j\})$. This allows us to compute the probability of delayed service from its corresponding maximum distortion \bar{b}_j by $\hat{d}_j = G_{b_j}(\bar{b}_j)$ as well as the average distortion from promising the optimal contract corresponding to the probability \hat{d}_j by $\Delta_j = \int_0^{\bar{b}_j} b_j dG_{b_j}(b_j)$.

The second new distribution will be the distribution of the virtual surplus of an average person who gets the contract \hat{D}_j . Let's define

$$g_{w_j}(\hat{w}_j) = E_n \left[\mathbf{1}[D_1^c] f_{n+1-(m-j)}(\hat{w}_j) + \sum_{i=1}^{m-1} \mathbf{1}[D_i \setminus D_{i+1}] f_{n+1-(m-j)+i}(\hat{w}_j) + \mathbf{1}[D_m] f_{n+1-(m-j)+m}(\hat{w}_j) \right]$$

and $G_{w_j}(\hat{w}_j) = \int_0^{\hat{w}_j} g_{w_j}(\hat{w}'_j) d\hat{w}'_j$. Remember that $w_k = w(0)$ whenever $k \leq 0$. The distribution G_{w_j} is now a mixture of distributions of virtual surpluses $w(v_{n+1-(m-j):n}), \dots, w(v_{n+1-(m-j)+m:n})$ with partition $\{D_1^c, D_1 \setminus D_2, \dots, D_{m-1} \setminus D_m, D_m\}$.

We can now compute the marginal effect of \bar{b}_j to the revenue $\pi^i(n, w)$. For all (n, w) and all i ,

$$\frac{\partial \pi^i(n, w)}{\partial \bar{b}_j} = \begin{cases} 0 & \forall w_{n+1-(m-i)-j} \leq \underline{w}_j, \\ (1 - \delta) \beta g_{b_j}(\bar{b}_j) w_{n+1-(m-i)-j} - \delta \bar{b}_j g_{b_j}(\bar{b}_j) & \forall w_{n+1-(m-i)-j} > \underline{w}_j. \end{cases}$$

Therefore the first order condition from the expected revenue at state D with respect to \bar{b}_j is

$$0 = E_{n,v} \left[\mathbf{1}[D_1^c] \frac{\partial \pi^0}{\partial \bar{b}_j} + \mathbf{1}[D_1 \setminus D_2] \frac{\partial \pi^1}{\partial \bar{b}_j} + \cdots + \mathbf{1}[D_{m-1} \setminus D_m] \frac{\partial \pi^m}{\partial \bar{b}_j} + \mathbf{1}[D_m] \frac{\partial \pi^m}{\partial \bar{b}_j} \right],$$

which we can rewrite as

$$\frac{\delta \bar{b}_j}{(1-\delta)\beta} = \int_{\underline{w}_j}^1 \hat{w}_j \frac{dG_{w_j}(\hat{w}_j)}{1 - G_{w_j}(\underline{w}_j)} = E_{G_{w_j}} [\hat{w}_j | \hat{w}_j > \underline{w}_j]. \quad (\text{A.4})$$

As in the illustrative example, we get that the maximum distortion from j th contract—adjusted with discount factor differences and de-normalized—must be equal to the average virtual surplus of a customer whom the seller expects to allocate this contract.

Step 4: characterization of optimal contracts By definition of \underline{w}_j we can now write

$$\underline{w}_j = \frac{\delta \int_0^{\bar{b}_j} b_j dG_{b_j}(b_j)}{(1-\delta)\beta G_{b_j}(\bar{b}_j)} = \frac{\delta E_{G_{b_j}}[b_j | b_j \leq \bar{b}_j]}{(1-\delta)\beta}. \quad (\text{A.5})$$

For any $j \in \{1, \dots, m\}$ have now two equations that relate \underline{w}_j and \bar{b}_j , equations (A.5) and (A.4). Both describe a strictly increasing continuous relationships between \bar{b}_j and \underline{w}_j . Let us first look at (A.4). At the lower bound $\underline{w}_j = 0$ we get that $\frac{\delta \bar{b}_j}{(1-\delta)\beta} = E_{G_{w_j}}[\hat{w}_j | \hat{w}_j > 0] \in (0, 1)$ and near the upper bound $\underline{w}_j = 1$ we have $\frac{\delta \bar{b}_j}{(1-\delta)\beta} = E_{G_{w_j}}[\hat{w}_j | \hat{w}_j \geq 1] = 1$. Equation (A.5) gives that at the lower bound $\bar{b}_j = 0$ we have $\underline{w}_j = \frac{\delta E_{G_{b_j}}[b_j | b_j \leq 0]}{(1-\delta)\beta} = 0$ and at the upper bound $\bar{b}_j = 1 - \delta$ we have $\underline{w}_j = \frac{\delta E_{G_{b_j}}[b_j]}{(1-\delta)\beta} \in \left(0, \frac{\delta}{\beta}\right)$.

Therefore we have either that there exists interior solution $(\underline{w}_j, \bar{b}_j)$, so that the probability $\hat{d}_j < 1$ or that the first order condition (A.4) always holds as strict equality and therefore the optimal $\bar{b}_j = 1 - \delta$ which means $\hat{d}_j = 1$ or $\bar{b}_j = c$ which means $\hat{d}_j < 1$ because promising delayed service with certainty is impossible due to partial commitment.

Step 5: existence of $\pi(d)$ To complete the analysis, we need to argue that the mechanism description given above gives well-defined functions $\pi(d)$ which we initially assumed to exist. The characterization above is a mapping from a continuation value function $\pi(\hat{d})$ to current expected revenue $\pi(d)$. We denote the mapping by $T\pi$. We will first use Blackwell's sufficient conditions to show that the mapping T is a contraction with speed of convergence δ . Then we can apply the contraction mapping theorem to show that $\pi(d)$ exists.

From the recursive formulation (A.2) it is clear that since $\sum_i \phi_i w_i \in [0, N]$, mapping T maps bounded functions $\pi(\hat{d})$ to bounded functions. If we have two functions, $\pi, \hat{\pi}$ such that $\hat{\pi}(\hat{d}) \geq \pi(\hat{d})$ for all \hat{d} , then $T\hat{\pi}(\hat{d}) \geq T\pi(\hat{d})$ for all \hat{d} , since at each state with any promises the continuation revenue is increased and the flow revenue unchanged. Therefore the optimal promises cannot lead to lower revenue. This guarantees monotonicity. To verify discounting, suppose $\hat{\pi}(\hat{d}) = \pi(\hat{d}) + a$.

Then each $\pi^j(n, w)$ will simply have an extra δa at the end. In the dynamic part of the problem δa terms cancel in $\pi^{j-1}(n, w) - \pi^j(n, w)$ and therefore $T\hat{\pi}(\hat{d}) = E_{n,v}[\pi^0(n, w) + \delta a] = T\pi(\hat{d}) + \delta a$. \square

A.3 Fully optimal mechanism

Proof of Theorem 3.6 We can make the same immediate observations as in the proof of Theorem 3.5: (1) all buyers with negative virtual surplus are never served; (2) the allocation is monotone in values in the sense that the highest types get instant service and the next highest will get delayed service; (3) only up to m of the new buyers are offered delayed service; (4) the optimal promises are ordered both in probabilities $\hat{d}_1 \geq \dots \geq \hat{d}_m$ and also in sets²³ $\hat{D}_1 \supset \dots \supset \hat{D}_m$; and (5) the recursive definition of the revenue function is again

$$\pi(d) = \max E_{n,v} \left[(1 - \delta) \sum_{i=1}^n \phi_i w_i + \delta \pi(\hat{d}) \right], \quad (\text{A.6})$$

where $\phi_i = 1$ if i receives instant service, $\phi_i = 0$ if is refused service, and $\phi_i = \beta \hat{d}_i$ if is promised delayed service with probability \hat{d}_i .

Step 1: allocation rule. Let $\pi^j(n, w)$ denote the maximum revenue that the seller can achieve when reserving exactly j current service slots for buyers from the previous period (for example $j = \#\{i : (n, v) \in D_i\}$). That is, the seller can serve only up to $m - j$ of the new buyers instantly and gives optimal promises for next period service to m buyers who were not served instantly. Because of monotonicity in the order of the promises, the buyers with virtual surpluses $w_n, \dots, w_{n+1-(m-j)}$ receive instant service (if their virtual surpluses are positive) and buyers with virtual surpluses $w_{n-(m-j)}, \dots, w_{n+1-(m-j)-m}$ receive the corresponding promises such that $\hat{d}_1 \geq \dots \geq \hat{d}_m$. Then $\pi^j(n, w)$ can be computed as

$$\pi^j(n, w) = \max_{\hat{d}} (1 - \delta) \sum_{i=1}^{m-j} \max\{0, w_{n+1-i}\} + (1 - \delta) \beta \sum_{i=1}^m \hat{d}_i w_{n+1-(m-j)-i} + \delta \pi(\hat{d})$$

such that $\hat{d}_j \in [0, 1]$ for each $j \in \{1, \dots, m\}$ and the commitment constraint is satisfied. Differentiation with respect to the probability of delay \hat{d}_i gives

$$\frac{d\pi^j(n, w)}{d\hat{d}_i} = (1 - \delta) \beta w_{n+1-(m-j)-i} + \delta \frac{\partial \pi(\hat{d})}{\partial \hat{d}_i}. \quad (\text{A.7})$$

The m first order conditions define vector of optimal probabilities \hat{d} as a function of virtual surpluses $w_{n-(m-j)}, \dots, w_{n+1-(m-j)-m}$, which we will discuss in detail later when we have determined the properties of $\frac{\partial \pi(\hat{d})}{\partial \hat{d}_i}$.

²³With the exception of zero-measure sets. The precise statement is: $Pr(\hat{D}_j \setminus \hat{D}_{j-1}) = 0$ for all $j \in \{2, \dots, m\}$.

Step 2: fulfilling promises optimally. We have determined the optimal probabilities of service, \hat{d} . Next we will find the precise optimal promises in terms of sets \hat{D} , where these promises are fulfilled. We argued above that the promises are ordered in the sense that higher types are ordered higher probability of service. Analogously to the waiting list mechanism case we can write

$$\pi(\hat{d}) = E_{n,v}[\pi^0(n, w)] - \sum_{j=1}^m \min_{\hat{D}: Pr(\hat{D}_j) \geq \hat{d}_j} E_{n,v} \left[\mathbf{1}[\hat{D}_j] [\pi^{j-1}(n, w) - \pi^j(n, w)] \right]. \quad (\text{A.8})$$

Since both the objective and the constraints are separable in j , we can again optimize point-wise. That is, to choose each contract \hat{D}_j such that $\pi^{j-1}(n, w) - \pi^j(n, w)$ is the smallest, conditional on $Pr(\hat{d}_j) \geq \hat{d}_j$.

Again, when $w_{n-(m-j)} \leq 0$, then all the buyers starting from $n - (m - j)$ will be refused service, so $\pi^{j-1}(n, w) - \pi^j(n, w) = 0$, whereas when $w_{n-(m-j)} > 0$, this difference is strictly positive. Therefore the binding promise constraint is always binding in the optimum, $\hat{d}_j = Pr(\hat{D}_j) \geq E_n[F_{n-(m-j)}(p^*)]$. Finally, note that $\pi^{j-1}(n, w) - \pi^j(n, w)$ is still strictly increasing in $w_{n-(m-j)}$.

Now, for each j , the optimal contract is $\hat{D}_j = \{(n, v) : \pi^{j-1}(n, w) - \pi^j(n, w) \leq \bar{b}_j\}$ for uniquely defined \bar{b}_j such that $Pr(\hat{D}_j) = \hat{d}_j$.

Step 3: optimal probabilities of delayed service. We can again change the optimization variable, instead of choosing the optimal probability of delayed service \hat{d}_j we can choose the optimal maximal distortion \bar{b}_j . Using this we get that $\frac{\partial \pi(\hat{d})}{\partial \hat{d}_j} = -\bar{b}_j$, which is independent of \hat{d}_i for $i \neq j$.

Moreover, if $\hat{d}_j = 1$ (seller is promised service with certainty), then we know that \hat{D}_j must include all pairs (n, w) , therefore $\bar{b}_j = 1 - \delta$ and so $\frac{\partial \pi(\cdot, 1, \cdot)}{\partial \hat{d}_j} = -(1 - \delta)$. If $\hat{d}_j \leq E_n[F_{n-(m-j)}(p^*)]$, the constraint is not binding and therefore $\frac{\partial \pi(\cdot, \hat{d}_j, \cdot)}{\partial \hat{d}_j} = 0$, whereas if $\hat{d}_j > E_n[F_{n-(m-j)}(p^*)]$ we have that $\bar{b}_j > 0$ and so $\frac{\partial \pi}{\partial \hat{d}_j} = -\bar{b}_j < 0$.

Whenever $E_n[F_{n-(m-j)}(p^*)] < \hat{d}_j < \hat{d}_{j'}$ we must have that $\hat{D}_j \subsetneq \hat{D}_{j'}$ and therefore $\bar{b}_j < \bar{b}_{j'}$. Therefore $\frac{\partial \pi(\cdot, \hat{d}_j, \cdot)}{\partial \hat{d}_j} = -\bar{b}_j > -\bar{b}_{j'} = \frac{\partial \pi(\cdot, \hat{d}_{j'}, \cdot)}{\partial \hat{d}_{j'}}$, therefore $\frac{\partial^2 \pi}{\partial \hat{d}_j^2} > 0$.

We can now study the first order condition (A.7). Suppose for a moment that commitment level c is high. Let's consider the upper corner solution first. Suppose $\hat{d}_j = 1$. This is optimal only if $(1 - \delta)\beta w_{n+1-(m-z)-j} \geq \delta(1 - \delta)$, which is equivalent to $w_{n+1-(m-z)-j} \geq \frac{\delta}{\beta}$. Now, let's take the lower corner solution. We know that $\frac{\partial \pi}{\partial \hat{d}_j} = 0$ for all $\hat{d}_j \leq E_n[F_{n-(m-j)}(p^*)]$, so $\hat{d}_j \leq E_n[F_{n-(m-j)}(p^*)]$ implies that $(1 - \delta)\beta w_{n+1-(m-z)-j} \leq 0$, which means that $w_{n+1-(m-z)-j} \leq 0$. This is what we already argued above—whenever the buyer to whom the seller is supposed to promise delayed service (according to her position in the ordered values) has negative virtual surplus, it will be optimal to promise service with zero probability. Finally, interior solutions, $\hat{d}_j \in \left(E_n[F_{n-(m-j)}(p^*)], 1 \right)$ have to satisfy $(1 - \delta)\beta w_{n+1-(m-z)-j} = \delta \frac{\partial \pi}{\partial \hat{d}_j} = \delta \bar{b}_j$. Therefore \hat{d}_j is

strictly increasing in $w_{n+1-(m-z)-j}$ for the interior solutions. If commitment level is not high, it gives a tighter upper bound, $\bar{b}_j \leq c$. The lower bound and the interior points are not affected.

We can summarize the choice of \bar{b}_j for a buyer²⁴ i who gets the promise by a function $\bar{b}_j = \bar{b}(w_i)$, where

$$\bar{b}(w_i) = \begin{cases} 0 & \forall w_i \leq 0, \\ \frac{(1-\delta)\beta}{\delta} w_i & \forall 0 < w_i < \frac{\delta}{(1-\delta)\beta} \min\{1-\delta, c\}, \\ \min\{1-\delta, c\} & \forall w_i \geq \frac{\delta}{(1-\delta)\beta} \min\{1-\delta, c\}. \end{cases}$$

Therefore we showed that the optimal probability of delayed service $\hat{d}_j(w_i)$ is equal to 0 when $w_i \leq 0$, equal to the upper bound $G_{b_j}(\min\{1-\delta, c\}) \leq 1$ when $w_i \geq \frac{\delta}{(1-\delta)\beta} \min\{1-\delta, c\}$ and is strictly increasing function $G_{b_j}\left(\frac{(1-\delta)\beta}{\delta} w_i\right) \in (0, 1)$ in the interior.

Step 4: existence of $\pi(d)$. The proof that functions $\pi(d)$ are well defined uses the Contraction mapping theorem and is analogous to the proof in the waiting list case. \square

²⁴This is to simplify the notation, in particular the contract j goes to i th highest valuing buyer with $i = n + 1 - (m - z) - j$.