Influencing Waiting Lists

James Schummer*

May 16, 2016

Abstract

Scarce goods (donor organs, public housing, etc.) are often allocated through waiting lists, especially when monetary transfers are undesirable. Arriving objects are offered to priority-ordered agents who may *defer* the object to the next agent in line in order to wait for a better one. We consider the welfare implications of arbitrarily influencing such deferral decisions (by force, by "nudging," etc.). When agents are patient, uninfluenced (equilibrium) behavior is Pareto-dominant; this conclusion strengthens as agents' risk-aversion increases. When risk-neutral agents are impatient, however, influence results in a welfare tradeoff between earlier and later agents in the queue. The results have implications for the "organ spoilage" problem in waiting lists for donor organs, where useful lower-quality organs spoil in the time that it takes to process the deferrals by agents early in the queue. Removing the right to defer such organs appears to solve the wastage problem by increasing utilization. But for some parameters (e.g. extremely high risk-aversion) such solutions could paradoxically lower the welfare of some agents it is intended to help. Fortunately the parameters that result in this phenomenon appear to be atypical in real world settings.

*Kellogg School of Management, MEDS Department, Northwestern University. Email: schummer@kellogg.northwestern.edu. Part of this work was presented in 2014 under the title "Waiting for Prizes." I thank Hyoduk Shin for initial discussion on an earlier version of this model that eventually motivated this paper, and Philip Marx for research assistance. I also thank Nick Arnosti, Anna Bogomolnaia, Ozan Candogan, Tayfun Sönmez, and Kathy Spier for helpful comments.

1 Introduction

Many goods are allocated through priority-based waiting lists. Their use is especially prevalent when a society wishes to avoid the use of monetary transactions[1] such as in the assignment of transplant organs, public housing for the needy, openings in congested drug abuse treatment programs, vacancies in publicly run nurseries, etc. In such allocation problems we are concerned not only with efficiency but also with the distribution of total welfare. This motivates us to look more carefully at the distribution of welfare in such settings. We study how welfare is affected when we change the degree to which agents are permitted to make selfish choices in such settings.

We consider a planner who strictly adheres to the principal of *first-come-first-served*. The agents are exogenously ordered (e.g. by time of arrival to the queue, health condition, neediness, etc.), and the agent in first position (at the head of the queue) has first rights over the next object to arrive. This *right* is typically not an *obligation*, however, especially when objects are heterogeneous. The agent in position 1 might choose to defer the object to the agent in position 2—without losing his position in the queue—in order to wait for a better object. The agent in position 2 might do the same, and so on. The right to exercise such deferral options clearly benefits an agent starting in position 1. Agents in later positions also benefit conditional on reaching an earlier position, but they also must wait longer to be offered higher quality objects. The ambiguous total welfare effects lead to the question of whether deferral should be permitted. More generally, our objective is to describe the welfare implications of influencing deferral decisions in any arbitrary way, e.g. constraining deferrals by law, encouraging *more* deferrals, "nudging" otherwise rational behavior, etc. We show that under some conditions influence is unambiguously harmful, while in others it is typically beneficial to most of the agents in the queue.

To further illustrate the ideas, consider waiting lists for transplant organs, where deferral options not only affect welfare but can cause inefficient waste.[2] Patients on the waiting list are prioritized based on various characteristics.[3] When a donor organ arrives it is offered to the agents sequentially, each of whom may decline the current organ. Particularly in the case of cadaveric

[1] See Roth (2007).

[2] This application is analyzed specifically in Section 6.

[3] Queues for kidneys, livers, and other organs could be based on health, age, probability of survival, time of arrival to the queue, etc.

kidneys this results in waste: some kidneys that could have benefitted patients later in the waiting list end up spoiling due to the time spent processing the many deferrals by patients earlier in the list.[4] This leads to longer waiting times overall, contributing to the growing length of such waiting lists.

There are various potential solutions to this "spoilage problem." Imagine a policymaker evaluating the following proposals.

- "Should we disallow deferrals altogether, requiring the first patient in the list to take any acceptable organ?"
- "Should we nudge patients early in the list to accept organs more frequently than they otherwise would?"
- "Should we require patients later in the list to accept organs above some specific quality threshold?"

Each of these solutions increases the organ utilization rate by influencing agents to defer less often. The reduction in average waiting time is welfare enhancing, but each solution also impacts welfare by *altering the probability distribution of consumption bundles offered to the agents.* For example, by altering deferral decisions made by agents early in a waiting list, this impacts agents who start later in the waiting list through both (i) the distribution of waiting time and (ii) the correlation of waiting time and organ quality. As we show in our main results these latter effects often, but not always, can be negative. In extreme cases these negative effects can even overpower the positive effect of reduced waiting times. In such cases, a policymaker who attempts to improve efficiency by prohibiting deferral—thus increasing organ utilization—would perversely *lower* overall welfare! Fortunately it appears that such extreme cases are unlikely when using realistic parameters in our model.

There are two points we should emphasize about our approach. First, we do not consider *how* one might go about actually influencing the behavior of agents in the waiting list. We ignore this issue since it plays no role in our analysis; we describe how influence would affect welfare whether such influence is feasible or not. Depending on the application, deferral decisions could be influenced in a variety of ways. One direct way is through simple coercion such as legal constraints or disqualification from the queue, which impose prohibitive penalties for deferring an offered object. A more indirect

[4]See the NY Times article of Sack (2012).

method would be to conceal information about the quality of the object being offered, reducing the agents' ability to discriminate between objects and effectively reducing their ability to defer some objects but not others. Finally, agents might obtain advice (e.g. medical) on whether to defer an object, which provides another avenue for influence.[5] Regardless, to obtain the most general results we ignore any constraints that might bind how much an agent could be influenced.

The second point is to emphasize our objective of analyzing welfare *distribution* effects rather than welfare *creation* effects. As mentioned above, permitting one agent to defer objects has a welfare distribution effect in that it alters the distribution over consumption bundles for agents later in the queue. In some applications, however deferral options can create welfare in the following way. Suppose "A-type" and "B-type" agents respectively prefer a-type and b-type objects. Social efficiency is improved by correlating the types of agents and their assigned objects. An A-agent should be influenced to defer a b-type object to solve what would otherwise be a *coordination inefficiency*. Solutions to this kind of problem are presented by Leshno (2014) who, roughly speaking, rewards deferring agents by probabilistically advancing them in the queue. Our objective is to analyze the above welfare-distribution effect in isolation from the the issue of coordination inefficiency since these are two confounding issues. To separate these effects, we eliminate coordination inefficiencies from our model by removing the horizontal differentiation of objects (e.g. a/b types as in Leshno (2014)) and assume instead that agents have identical tastes over *vertically* differentiated objects (quality types).

1.1 Overview of Results

We begin with a benchmark case in which risk-neutral, homogeneous agents do not discount payoffs, absent any feasibility constraints (such as the organ spoilage problem mentioned earlier). This scenario yields a fairly striking "expected payoff equivalence" result stating that influencing the deferral decisions of agents in early queue positions has *no* effect on the expected payoffs to agents in later queue positions. This leads to the first main result stating that uninfluenced (equilibrium) behavior leads to a Pareto-dominant out-

[5] In the application of organ waiting lists, the use of these methods may also be bound by ethical constraints, which goes beyond the scope of this paper.

come: influencing the decisions made by agents in any one position in the queue cannot strictly improve the continuation payoff to any other position in the queue. Therefore without an efficiency-based reason to influence behavior (e.g. to solve the spoilage problem described earlier), a planner should allow agents to fully exercise deferral options in this scenario.

Despite our *expected* payoff equivalence result, a change in deferral decisions changes the *probability distribution* of agents' consumption, impacting risk-averse agents. We show that the *variance* of an agent's total waiting costs (time) is decreased by constraining deferral decisions by agents in earlier positions. This would benefit an entity that cares only about the waiting-cost component of an agent's payoff, e.g. an insurance company responsible for a patient's treatment costs before a transplant organ becomes available. However a change in the distribution of waiting time is typically accompanied by a change in the correlation between waiting time and object quality. Intuitively, *allowing* deferrals tends to create a negative correlation between waiting time and object quality, to the benefit of risk-averse agents. Despite the fact that there are these two confounding effects, Corollary 4 strengthens the conclusion drawn under risk-neutrality: for agents with constant absolute risk aversion, uninfluenced (equilibrium) behavior leads to a *strictly* Pareto-dominant outcome.

On the other hand agents' *impatience* erodes this kind of result. When agents discount future payoffs, we obtain an "opposed interests" result implying the following. If an object that is acceptable to an agent in a later queue position is instead assigned to agents in some earlier position, then an agent starting in the later position benefits from this if and only if it harms an agent starting in the earlier position. Roughly speaking, if we ask whether it is beneficial to influence the deferral decisions of agents early in the queue, we would typically receive opposite answers from agents who start early/late in the queue. A significant conclusion from this is that risk-aversion and impatience are two important parts of agents' preferences that determine, in opposite ways, the welfare consequences of influencing deferral decisions.

Finally we apply the above results to the organ spoilage problem. Interestingly the spoilage problem itself can be modeled as a form of influence: agents in sufficiently late queue positions are forced (by nature) to defer all objects (because they have spoiled by the time they are offered). This allows us to incorporate this problem into our model and obtain general qualitative statements. Should the spoilage problem be addressed by limiting agents' rights to defer organs? As we argue in Section 6 such a solution becomes more compelling as agents become more impatient, but could become less compelling as agents become more risk averse.

1.2 Related Literature

Our model is related to work that has been done both in operations and economics. One special case of our model reduces to the *parallel process problem* studied in a series of papers starting with Agrawala et al. (1984). A set of jobs, each with an unknown length, must be completed by a set of processors each with its own work rate. To minimize total job time (total waiting and processing time) it may be suboptimal to use the slowest of the processors. Agrawala et al. derive a simple, threshold-based policy for assigning jobs to available processors that minimizes total job time. They also show that this policy would result from equilibrium behavior if the jobs were individual agents.[6] Our results apply to their model by interpreting processors as objects and their speeds as quality. In particular, their policy is a special case of Equation 2. Their optimality result, which can be reinterpreted as a *utilitarian welfare* result, is a corollary of our Pareto-dominance result, Theorem 3.[7]

Su and Zenios (2004) consider a model very similar to ours: patients are offered arriving kidneys on a first-come-first-served basis and are free (uninfluenced) to defer kidneys of low quality. Among other results, they also show that equilibrium behavior maximizes total (utilitarian) welfare for the agents already in the queue.[8]

If we also consider the welfare of agents *yet to arrive* to a waiting list, however, it is known that equilibrium behavior does not maximize total (utilitarian) welfare. Naor (1969) observed that when agents are deciding *whether* to join a waiting list, decentralized behavior is socially suboptimal since the agent joining the queue fails to internalize the waiting cost he imposes on any future agents that might arrive. Using the same logic in our model, when the last agent in line defers a low-quality object—and thus fails to depart

[6]Kumar and Walrand (1985) extend this result to a more general setting.

[7]Coffman et al. (1987) consider the same setup with the object of minimizing the *makespan* (time elapsed until all jobs are complete), which can be reinterpreted as a *Rawlsian welfare* criterion.

[8]While we obtain the stronger conclusion of Pareto-optimality, Su and Zenios obtain their utilitarian result under more general conditions; hence there is no logical relation between our results.

the queue—he does disregards the effect it has on future agents who will arrive behind him in the queue. Expected total welfare could be improved by reducing the deferral agent of the last agent in line.

In fact, Hassin (1985) concisely explains that this problem can be solved by switching the priority structure to *last*-come-first-served (LCFS), which prioritizes agents in the *reverse* order of their arrival. This forces agents to internalize the probabilistic arrival of future agents when making deferral decisions, inducing socially optimal decisions overall. Su and Zenios (2004) exploit this idea, estimating the hypothetical welfare gains from using LCFS for kidney allocation. At the same time they point out (as does Hassin) that LCFS methods are manipulable in various ways, increase risk to the agents, are inequitable, and are unlikely to be politically acceptable.

Within the economics literature there has been a recent surge of interest on dynamic matching and assignment. We have already mentioned the work of Leshno (2014) that is complementary to ours. With the discrete-time arrival of objects to agents with heterogeneous preferences, Leshno reduces the coordination inefficiency we mentioned earlier by optimizing a "buffer policy" which essentially probabilistically alters the waiting list position of agents who defer an object.

Bloch and Cantala (2014) also consider the discrete-time arrival of heterogeneous objects, but without the persistent preferences of Leshno. They analyze equilibria under a mechanism that assigns objects to a waiting list probabilistically, maintaining priority in the weaker sense that earlier agents in the list have a greater chance of receiving offers. By reducing the probability of making offers to earlier positions, the planner reduces the incentive for agents in later positions to defer objects in order to reach earlier positions.

With the motivation of a public housing application, Thakral (2015) examines a type of school choice model in which apartment types arrive in discrete time periods, each having its own priority order over the waiting agents. Thakral proposes a multiple wait list procedure, in which an arriving apartment "proposes" to its highest priority agent, which in turn gives that agent the option either to take the apartment or to join a specific waiting list for some other single apartment type. This idea, incorporated with the "you want my house I get your turn" concept in Abdulkadiroglu and Sönmez (1999), yields a strategyproof mechanism with a desirable efficiency property that also respects the apartments' priority orders.

In the kidney *exchange* model of Ünver (2010), agents *and* objects arrive in pairs, where the agents have implicit property rights over their initial endowments. With the motivation of organ compatibility, certain trades are feasible (or more desirable than others). The analysis covers both pairwise and multi-way exchanges.

If we generalize the concept of arriving *objects* to arriving *agents*, we obtain the related, burgeoning literature on dynamic 2-sided matching.[9] Doval (2014) considers stability in 2-sided matching when agents who arrive in different periods may postpone their arrivals. Akbarpour et al. (2014) analyze the limit behavior of a market in which agents randomly arrive (and depart) a market to be pairwise matched, comparing mechanisms that do and do not assign agents immediately upon arrival.

2 Model

We consider an ordered set of identical agents, each waiting to consume a single object. Due to the nature of our analysis and results, it turns out to be without loss of generality to suppose that the set of agents is the set of natural numbers $\mathbb{N} = \{1, 2, \ldots\}$.[10] Objects arrive randomly over time and waiting is costly.

There is a set of object types $\mathcal{O} = \{1, 2, \ldots, n\}$.[11] Objects of each type arrive according to a Poisson process independent of arrivals of other types. Formally, for each $i \in \mathcal{O}$, the time between arrivals of i-type objects is exponentially distributed with parameter $\mu_i > 0$ (mean $1/\mu_i$), and arrival times for i are independent of arrival times for types in $\mathcal{O} \setminus \{i\}$.[12] It is convenient to denote the arrival rate for a set of types as follows.

$$\forall O \subseteq \mathcal{O}, \quad \mu_O \equiv \sum_{i \in O} \mu_i$$

An agent who consumes an object of type $i \in \mathcal{O}$ receives a payoff of $v_i \in \mathbb{R}$. To simplify the statement of some results, we assume that values are distinct ($i \neq j$ implies $v_i \neq v_j$), though this does not affect our conclusions.

[9] See also Damiano and Lam (2005), Kadam and Kotowski (2014), Kennes et al. (2014), Kurino (2009), Pereyra (2013), and Marx and Schummer (2016).

[10] Results that apply to agent $k \in \mathbb{N}$ are independent of whether agents $\ell > k$ are present.

[11] The results easily extend to cases where $\mathcal{O}$ is infinite.

[12] An equivalent specification of the model is that objects in general arrive with exponentially distributed arrival times with parameter $\sum i \in \mathcal{O}\mu_i$, and an independent random process determines each object's type with probabilities $\mu_i / \sum_{j \in \mathcal{O}} \mu_j$.

Furthermore we order the types' labels from best to worst, so $v_1 > v_2 > \cdots > v_n$.

Agents incur a constant flow of costs while waiting for an object to consume, which we normalize to a cost of one unit per unit of time. That is, a (non-discounting) agent who is assigned an i-type object after waiting t units of time has a payoff of $v_i - t$. We initially assume that agents are risk-neutral, while in Section 4 we assume agents are risk-averse over such payoffs. In Section 5 we modify these payoffs to account for discounting.

2.1 Waiting List Policies

We now specify how objects may be allocated to agents by defining the class of *Waiting List Policies*. To motivate this definition, recall our objective to consider encouraging, restraining, or otherwise modifying the deferral decisions agents make when they are offered objects sequentially according to their priority order in a waiting list. When an agent is assigned an object, the agent leaves the queue, and all remaining agents move up one position (maintaining their relative original ordering).

When deferral decisions are unrestricted, equilibrium behavior of an agent in position 1 would be to accept any object from some "best" set of types, say $W^*(1)$, and to defer any others. Given this fact, the equilibrium behavior of an agent in position 2 would be to accept some (possibly empty) set of second-best object types, say $W^*(2) \setminus W^*(1)$ and to defer the rest. Similar sets would describe the remaining positions.[13]

Now suppose that we can influence the deferral decisions of any agent who reaches position 1. If we consider any *arbitrary* modification of agents' behavior (which we do for the sake of generality), then an agent in position 1 could be influenced to accept some arbitrary set of object types $W(1)$ and to defer any others. For example if $W(1) \supseteq W^*(1)$ then this would represent a restriction on deferral decisions in position 1; $W(1) \subseteq W^*(1)$ would represent encouragement of deferrals; in general there could be no inclusion relation.

Similarly we can imagine influencing the acceptance/deferral choices of any agent in position 2 so that such an agent accepts object types in $W(2) \setminus W(1)$ and defers the rest. In general we imagine modifying the behavior of any agent who reaches position k so that they accept only object

[13] Recall the homogeneity of agents, so these sets need not depend on the identities of the agents.

types belonging to $W(k) \setminus W(k-1)$. In this sense a *Waiting List Policy* W describes a consistent way of influencing agents' behavior anonymously: only as a function of their position in the queue, and not of their identity. This captures the kinds of policy proposals suggested in Section 1, e.g. disallowing deferrals by agents in some position k, where values of k could be chosen by a policymaker.

Definition 1. A **waiting list policy** is a correspondence $W\colon \mathbb{N} \to 2^{\mathcal{O}}$ that is monotonic and nonempty: $k < \ell$ implies $W(k) \subseteq W(\ell)$, and $W(1) \neq \emptyset$.

The definition is interpreted as follows. An agent in position 1 of the waiting list is allocated the next object to arrive whose type belongs to $W(1)$. When this happens, this agent leaves and every other agent moves up one position in the waiting list. In general when an object arrives whose type belongs to $W(k) \setminus W(k-1)$, the object is assigned to the agent currently in position k, and any agent in some later position $\ell > k$ move to position $\ell - 1$.[14] If no such k exists the object is discarded. Nonemptiness guarantees that each agent eventually receives an object.[15]

Our definition ignores whether it would be *feasible* to influence agents in a way that implements any particular policy W. We are treating the vague concept of *influence* as a black box. This is intentional in the interest of generality. Feasibility constraints (including ethical ones) can be imposed on the set of waiting list policies without affecting our welfare comparisons.

Despite this generality, it is also worth pointing out limitations of our approach. The definition implies that a planner must influence agents' decisions *anonymously*, i.e. solely as a function of their position in the queue, but not of their identity. In one sense this good: it rules out cases in which a planner treats one agent better than he would treat another, given their identical place in the queue. On the other hand this rules out mechanisms that create incentives by dynamically *changing* the agents' relative priorities based on past deferral decisions (e.g. as in Leshno (2014)). Second, our definition disallows the planner from using information about the *current queue length* when influencing behavior, e.g. constraining only the last agent

[14] In this sense, W is the cumulative function of a social choice function that maps object types to queue positions. We use this cumulative form because it allows us to state our results much more clearly.

[15] This is without loss of generality since, if $W(1) = \emptyset$, we can simply ignore the agent in position 1 who will be stuck in the queue forever and apply our results to the remaining (relabeled) positions.

in line. Such methods can overcome inefficiencies when agents arrive dynamically over time (Naor (1969)), but could also treat two agents differently based on the presence of lower-priority agents. This example illustrates the degree to which we are specifically analyzing problems in which the planner is committed to priority-based allocation. Among other examples, organ waiting lists appear to adhere to this principle.

Examples

The most natural example of a waiting list policy is the one that would result when deferral options are completely uninfluenced. When rational, homogenous agents can foresee the actions of others, their behavior would result in consumption described by a waiting list policy W^* alluded to at the beginning of Subsection 2.1. This policy turns out to play a central role for us beginning with Theorem 3.[16]

At another extreme, consider the complete removal of deferral options so that the agent in position 1 must accept whichever object arrives next. In this case, $W(1) = \mathcal{O}$, and hence $W(k) = \mathcal{O}$ for all $k \in \mathbb{N}$. More generally consider a planner who disallows all deferral options but also withholds (discards) certain object types from the agents. It turns out that our technical analysis makes use of the class of such policies.

Definition 2. The waiting list policy W is the **no-deferrals policy with acceptable set** $\hat{\mathcal{O}}$ when, for all $k \in \mathbb{N}$, $W(k) = \hat{\mathcal{O}}$.

Many other ideas can be expressed as waiting list policies. For instance a class of object types $\mathcal{O}' \subseteq \mathcal{O}$ can be reserved for agents whose position *exceeds* some value k. by requiring $W(\ell) \cap \mathcal{O}' = \emptyset$ for any $\ell \leq k$. In general, the set $W(1)$ need not be "better" than every other $W(k)$.

Finally in Section 6 we exploit the fact that the organ spoilage problem itself can be modeled as a constraint on waiting list policies by requiring that $W(k)$ be constant for all sufficiently large k. This captures the idea that agents cannot be offered objects if they are sufficiently late in the queue.

[16] The derivation of this equilibrium appears in Su and Zenios (2004) and has counterparts in earlier literature cited in Section 1.

3 Baseline Setting

3.1 Equivalence results

We begin with some fundamental results on the expectations of payoffs for an agent starting from a given position in the waiting list. The results have implications for the baseline case of non-discounting, risk-neutral agents. In particular, there is a (typically unique) waiting list policy that Pareto-dominates every other policy.

To begin, we calculate the expected (continuation) payoff to an agent starting in any kth position of the waiting list under an arbitrary policy W, showing that it depends only on the set $W(k)$. Policies W and W' are shown to be payoff-equivalent for position k whenever $W(k) = W'(k)$, i.e. whenever positions 1 through k *collectively* consume the same types of objects under the two policies. This result also implies that we can determine the set of all achievable expected payoffs for position k even restricting attention to the simple class of no-deferrals.

To establish intuition, consider an agent starting in position $k = 2$ evaluating his potential payoffs under W vs. W' where

$$\begin{array}{ll} W(1) \subsetneq W(2) = \mathcal{O} & \text{(position 1 defers objects accepted by position 2),} \\ W'(1) = W'(2) = \mathcal{O} & \text{(position 1 defers no objects).} \end{array}$$

With what probability will the agent ultimately consume an object type in $W(1)$? To do so under policy W, he must first reach position 1 (which happens with probability $\mu_{W(1)}/\mu_{\mathcal{O}}$), after which he is guaranteed to do so. To do so under policy W', he must first reach position 1 (which happens with probability 1), after which he consumes an object from $W(1) \subsetneq \mathcal{O}$ with probability $\mu_{W(1)}/\mu_{\mathcal{O}}$. Thus we obtain the same probability under both policies. The same argument generalizes to individual object types $i \in \mathcal{O}$ and to any position k. Therefore any two policies for which $W(k) = W'(k)$ induce equivalent distributions over objects to the agent in position k.

Theorem 1 (Object equivalence)**.** *For any policy W, the probability that an agent in position k ultimately consumes an object of type $i \in W(k)$ is $\mu_i/\mu_{W(k)}$.*

The straightforward proof is in the Appendix, as are all other omitted proofs.

Continuing with our example above, consider now the amount of time that the agent in position 2 will wait before accepting an object. Under W' the agent must wait for the arrival of two objects (of any type), which on average takes $2/\mu_{\mathcal{O}}$ units of time. Under W he first must wait an average $1/\mu_{\mathcal{O}}$ units of time for the next object to arrive. If it belongs to $W(2) \setminus W(1)$ then he is finished waiting, but otherwise (with probability $\mu_{W(1)}/\mu_{\mathcal{O}}$) waits an average $1/\mu_{W(1)}$ additional units of time, for a total expectation of $2/\mu_{\mathcal{O}}$ as under W'. This argument also generalizes to any two waiting list policies that coincide at some position k. Therefore the total expected payoff to an agent in position k can be described by referring only to the set $W(k)$. In the Appendix we prove the latter result directly, yielding the following.

Theorem 2 (Expected-payoff equivalence)**.** *Under any policy W, the expected payoff to an agent starting in position $k \in \mathbb{N}$, denoted $\Pi(k; W)$, is*

$$\Pi(k;W) = \frac{\sum_{j \in W(k)} \mu_j v_j - k}{\mu_{W(k)}} \tag{1}$$

In particular $\Pi(k; W)$ is a function only of $W(k)$.

With Theorem 1 this implies the expected waiting time equivalence of the kind we described earlier for the special case of position $k = 2$. We let t_k^W denote the random variable describing the amount of time it takes for an agent starting in position k to receive an object under policy W.

Corollary 1 (Expected-waiting-time equivalence)**.** *For any policy W, an agent starting in position $k \in \mathbb{N}$ has an expected waiting time (cost) of $E(t_k^W) = k/\mu_{W(k)}$.*

A second consequence of Theorem 2 is that, in order to determine which expected payoffs are achievable for any given position k, it is sufficient to consider only the class of no-deferrals policies.

Corollary 2. *Under any policy W, the expected payoff to an agent starting in position k is the same as the expected payoff under the no-deferrals policy W' defined by $W'(\ell) = W(k)$ for all $\ell \in \mathbb{N}$.*

3.2 No influence and Pareto-dominance

Consider which policies maximize the expected payoff to some fixed position $k \in \mathbb{N}$. To construct such a policy W it is necessary and sufficient to specify a set $W(k)$ that maximizes Equation 1. The remaining sets $W(\ell)$, $\ell \neq k$, can be specified arbitrarily (subject to Definition 1). Not surprisingly such a set $W(k)$ is a "threshold set" of the form $W^*(k) = \{1, 2, \ldots, i^*(k)\}$,[17] where the threshold $i^*(k)$ is the highest index (i.e. lowest quality) whose value exceeds the expected payoff $\Pi(k; W^*)$.[18]

Lemma 1 (k's favorite policy)**.** *For any $k \in \mathbb{N}$, $\Pi(k; \cdot)$ is maximized by any policy W that satisfies $W(k) = W^*(k) \equiv \{1, 2, \ldots, i^*(k)\}$, where*

$$i^*(k) \equiv \max \left\{ i \in \mathcal{O} : v_i \geq \frac{\sum_{j=1}^{i-1} \mu_j v_j - k}{\sum_{j=1}^{i-1} \mu_j} \right\} \tag{2}$$

The thresholds are increasing: $\ell < \ell'$ implies $i^*_\ell \leq i^*_{\ell'}$ and $W(\ell) \subseteq W(\ell')$.

The monotonicity of thresholds is intuitive. Agents in later positions are willing to accept lower-value objects in order to avoid greater expected wait times. The monotonicity also means that it is feasible to *simultaneously* provide each position k with its favorite set $W^*(k)$. We say that a policy W **(position-wise) Pareto dominates** a policy W' if, for all $k \in \mathbb{N}$, $\Pi(k; W) \geq \Pi(k; W')$.

Theorem 3 (Pareto-dominance)**.** *Let W^* be the policy where each $W^*(k)$ is defined as in Lemma 1. Policy W^* Pareto dominates every other policy.*

Proof. Since $k < \ell$ implies $W(k) \subseteq W(\ell)$, W^* is a feasible policy. Pareto dominance is immediate from Lemma 1. □

In a related model, Su and Zenios (2004) show that a policy analogous to W^* is optimal for a utilitarian planner: it maximizes the *sum* of agents' expected payoffs.[19] Agrawala et al. (1984) also provide a utilitarian result for the parallel processor problem which is a special case of our model. Theorem 3 strengthens these observations: since W^* is Pareto-dominant, it is the optimal policy under any reasonable welfare objective.

Finally we observe that W^* represents how objects would be allocated if agents could selfishly—and without influence—decide what objects to defer

[17] Recall objects are ordered in decreasing order of value v_i.

[18] Such threshold results are well-known, e.g. see Agrawala et al. (1984) or Su and Zenios (2004).

[19] They also point out that this result no longer holds when patients are arriving over time; see Subsection 1.2 and Naor (1969).

based on their own position in the queue. This follows from Equation 2. First, once any agent reaches position 1 in the queue, it is clearly optimal to accept an object if and only if it belongs to $W^*(1)$ since the behavior of other agents becomes irrelevant at this point. Given this, while occupying position 2 in the queue it is optimal for an agent to accept an object if and only if it belongs to $W^*(2)$. The agent cannot improve by accepting any other object, while deferring such an object cannot change the behavior of the agent in position 1.[20]

Continuing such arguments leads to a conclusion analogous to one the related model of Su and Zenios (2004): W^* would describe subgame-perfect equilibrium consumption in a dynamic "waiting list game" where agents make uninfluenced deferral decisions. Because of the closeness of this observation to the result of Su and Zenios, we omit further formalization.[21] Nevertheless we shall refer to W^* as the **no-influence equilibrium policy**.

3.3 Waiting time distributions

Fixing a policy W, Corollary 1 describes the *expected* waiting time from any position k as a function only of the set $W(k)$. This does not tell us anything else about the *distribution* of agents' waiting times. In the application of organ waiting lists, for example, doctors and patients might care about predictability. Though we do not model this in our payoffs, improved forecasting of the timing of a transplant can aid in the doctor's choice of treatment. Both patients and care provides might benefit from a reduction in the variance of patient waiting times. Even an insurance company responsible for waiting costs (e.g. dialysis) might have a financial preference to see lower variance in waiting costs.

For these reasons we are interested in the distribution of waiting time from any position k under an arbitrary policy W. Let t_k^W denote the random variable representing position k's waiting time under policy W: the time that an agent who starts in position k must wait before receiving an object. We have shown that $W(k) = W'(k)$ implies $E(t_k^W) = E(t_k^{W'})$. It turns out,

[20] There is a minor technical point here, which is that an agent in position 1 could be indifferent about accepting type $i^*(1)$. However then all agents would be indifferent about this decision, using arguments as in Theorem 2.

[21] Su and Zenios restrict attention to threshold strategies, partly because they have a continuum of object types. Our finite model demonstrates that this assumption is unnecessary.

however, that even when $W(k) = W'(k)$, two distinct policies W and W' typically yield different distributions for t_k^W and $t_k^{W'}$. While the distribution of t_k^W may be difficult to describe in general, the class of no-deferrals policies yield waiting times that follow an Erlang distribution—a sum of i.i.d. exponential distributions, implying the following.

Lemma 2 (Waiting times for no-deferrals policies)**.** *Consider a no-deferrals policy W, i.e. where $W(k) \equiv \hat{O} \subseteq \mathcal{O}$ for all $k \in \mathbb{N}$. For any $k \in \mathbb{N}$ the waiting time t_k^W has an Erlang distribution with mean $E(t_k^W) = k/\mu_{\hat{O}}$ and variance $\mathrm{Var}(t_k^W) = k/\mu_{\hat{O}}^2$.*

Proof. The waiting time for an arrival of a single object from $\hat{O}$ is exponentially distributed with parameter $\mu_{\hat{O}}$. An agent in position k receives such an object precisely after k i.i.d. such arrivals, hence the mean and variance calculations follow directly. Furthermore the sum of k i.i.d. exponentially distributed variables yields an Erlang distribution. □

Despite the difficulty of describing waiting time distributions in general, it is possible to describe the variance of t_k^W for an arbitrary policy W. To understand the idea behind the proofs, consider two object types $\mathcal{O} = \{1, 2\}$ and a policy W where $W(1) = \{1\}$ and $W(2) = \{1, 2\}$. Position 2's waiting time decomposes into $t_2^W = t' + t''$ as follows. First is a wait of t' units of time for the first arrival of an object from the set $\{1, 2\}$; note that t' is exponentially distributed with parameter $\mu_1 + \mu_2$. Conditional on that first object being of type 1, there is a wait of t'' for the arrival of a new object of type $1 \in W(1)$, and t'' is exponentially distributed with parameter μ_1. Otherwise, however, the agent departs immediately with the type 2 object, waiting $t'' = 0$ additional units of time. So t'' is either distributed $\exp(\mu_1)$ with probability $p_1 \equiv \mu_1/(\mu_1 + \mu_2)$, or is identically zero with the remaining probability.

The variance of t'', which can be computed in a few different ways[22] is

$$\mathrm{Var}(t'') = \frac{2}{\mu_1(\mu_1 + \mu_2)} - \frac{1}{(\mu_1 + \mu_2)^2}$$

[22] Being a weighted density of exponentials, t'' follows a hyper-exponential distribution which has a known expression for variance. It can also be computed using Theorem 4 or its proof.

Since t' is exponentially distributed and independent of t'' we have

$$\begin{aligned}\mathrm{Var}(t_2^W) = \mathrm{Var}(t' + t'') &= \frac{1}{(\mu_1+\mu_2)^2} + \left(\frac{2}{\mu_1(\mu_1+\mu_2)} - \frac{1}{(\mu_1+\mu_2)^2}\right) \\ &= \frac{2}{\mu_1(\mu_1+\mu_2)}\end{aligned}$$

This exceeds the variance under policy W' where $W'(2) = W'(1) = \{1, 2\}$ (see Lemma 2) which is $\mathrm{Var}(t_2^{W'}) = 2/(\mu_1 + \mu_2)^2$. Thus W and W' yield different variances in waiting time even though both policies provide the same expected waiting time (and expected payoff) to an agent in position 2.

In the proof of Theorem 4, the two "objects" used above represent the sets of objects $W(1)$ and $W(2) \setminus W(1)$, or more generally $W(k-1)$ and $W(k) \setminus W(k-1)$. In addition t' represents the wait for the first object from $W(k)$ and t'' represents the *entire* continuation waiting time (either zero or a continued wait from position $k-1$). This gives us a recursive expression for the variance of t_k^W which reduces to the following.

Theorem 4 (Waiting time variance)**.** *For any policy W, the waiting time from position k has a variance of*

$$\mathrm{Var}(t_k^W) = \frac{1}{\mu_{W(k)}}\left(\left(\sum_{\ell=1}^{k} \frac{2\ell}{\mu_{W(\ell)}}\right) - \frac{k^2}{\mu_{W(k)}}\right). \tag{3}$$

Hence for positions $\ell < k$, $\mathrm{Var}(t_k^W)$ *decreases* as we expand the set $W(\ell)$ (increasing $\mu_{W(\ell)}$).[23] In other words, imagine that position $\ell < k$ defers some object type $i \in W(k)$. If we remove position ℓ's ability to defer this object type, then this necessarily reduces the variance in waiting time from position k. At the extreme, the variance of t_k^W is minimized by using a no-deferrals policy.

Corollary 3 (Waiting time variance decreases in W)**.** *Fix a position $k \in \mathbb{N}$ and W and W' such that $W'(k) = W(k)$ and, for all $\ell < k$, $W(\ell) \subseteq W'(\ell)$. Then $Var(t_k^{W'}) \leq Var(t_k^W)$. Hence, subject to the constraint that $W(k) = \hat{\mathcal{O}}$ for some $\hat{\mathcal{O}} \subseteq \mathcal{O}$, the policy that minimizes* $\mathrm{Var}(t_k^W)$ *is the no-deferrals policy $W(\ell) \equiv \hat{\mathcal{O}}$.*

[23] Of course feasibility requires that we maintain $W(\ell) \subseteq W(\ell+1)$.

Thus there is a tradeoff between Pareto-optimality and predictability. The former is obtained with the the no-influence equilibrium policy W^*. By influencing agents early in the queue to exercise fewer deferrals, Corollary 3 shows that we can lower the variance of waiting times for agents later in the queue.

Finally we observe that expanding the set $W(k)$ itself could increase or decrease $\mathrm{Var}(t_k^W$ depending on the parameters of the problem. For example, fix $\mathcal{O} = \{1, 2\}$, $\mu_1 = 1$, and consider the no-deferrals policy $W(\ell) = \{1\}$ for all $\ell \in \mathbb{N}$ for which $Var(t_\ell^W) \equiv \ell$ (Lemma 2). Consider a switch to policy W' where some position k no longer defers type 2, i.e. $W'(\ell) = \{1\}$ for $l < k$ and $W'(\ell) = \{1, 2\}$ otherwise. This switch has two effects. First, intuitively, the increased object consumption rate lowers average waiting times for positions $\ell \geq k$, which tends to lower variance. Second however is that an agent in, say, position k now either leaves the queue somewhat quickly (if the next object to arrive is type $i = 2$) or must wait to reach position 1 before receiving an object. Either effect could dominate.[24]

4 Risk-averse agents

We now consider how our results are impacted by agents' risk-aversion. Risk-averse agents prefer lower variability in their payoffs. As we showed in Subsection 3.3, variability in waiting time can be reduced by reducing the rate at which deferrals are made, naively suggesting that influence could be used to benefit risk-averse agents. This ignores another effect, however. Payoff variability is lowered by increasing the correlation between waiting time (cost) and object value. When agents make uninfluenced deferral decisions, agents beginning late in the queue wait longer conditional no receiving higher-valued objects. This argument suggests that risk-averse agents also would prefer to have uninfluenced deferral decisions, as was the case for risk-neutral agents (Theorem 3). Strikingly, the latter argument is the only one that matters. Specifically, we extend Theorem 3 to the case of agents with constant relative risk aversion. In the process we prove an "aligned interests" result, showing that a kind of marginal policy change at some position k benefits an agent in that position if and only if it benefits an agent in position $k-1$.

[24]The latter effect indeed dominates unless μ_2 is large or k is small. When $k \leq 3$ the former effect always dominates.

Formally, we suppose in this section that agents have constant absolute risk aversion, i.e. utility from payoff $v-t$ is $u(v-t) = -e^{-\alpha(v-t)}$ with common risk parameter $\alpha > 0$. We assume that $\alpha < \mu_1$ in order to bound the expected utility from waiting only for the best object.[25]

The following two observations are straightforward. First, if (waiting time) t is exponentially distributed with parameter μ, then

$$E(u(-t)) = \int_0^\infty -e^{-\alpha(-t)} df(t) = -\frac{\mu}{\mu - \alpha} \tag{4}$$

Second, if (payoff components) $x_1, \ldots, x_k$ are independent random variables, then

$$\begin{aligned} E(u(\sum x_i)) &= \int \cdots \int -e^{-\alpha(\sum x_i)} df(x_1) \cdots df(x_k) \\ &= -\prod \int e^{-\alpha(x_i)} df(x_i) = -\prod -E(u(x_i)) \end{aligned} \tag{5}$$

Fixing a policy W, denote the **expected utility** to an agent starting in position k as U_k^W. It is not difficult to derive U_1^W: An agent in position 1 receives an object from $W(1)$ (with uncertain value v) after waiting t units of time (which is exponentially distributed with parameter $\mu_{W(1)}$). Thus the expected utility to position 1 is

$$\begin{aligned} U_1^W \equiv E(u(v-t)) &= -(-E(u(v)))(-E(u(-t))) \\ &= E(u(v)) \frac{\mu_{W(1)}}{\mu_{W(1)} - \alpha} \\ &= \frac{\mu_{W(1)}}{\mu_{W(1)} - \alpha} \sum_{i \in W(1)} \frac{\mu_i}{\mu_{W(1)}} (-e^{-\alpha v_i}) \end{aligned}$$

It is more tedious to describe U_2^W. An agent in position 2 ultimately receives an object from $W(2)$, but his waiting time (distribution) depends on whether or not that object belongs to $W(1) \subseteq W(2)$. Nevertheless Equation 4 and Equation 5 allow us to write a recursive expression for U_k^W as follows.

$$U_{k+1}^W = \frac{\mu_{W(k+1)}}{\mu_{W(k+1)} - \alpha} \left(\frac{\mu_{W(k)}}{\mu_{W(k+1)}} U_k^W + \sum_{i \in W(k+1) \setminus W(k)} \frac{\mu_i}{\mu_{W(k+1)}} (-e^{-\alpha v_i}) \right) \tag{6}$$

[25] The weaker assumption $\alpha < \sum_{\mathcal{O}} \mu_i$ would suffice, but would require additional care in some statements.

To prove Equation 6 observe that an agent in position $k+1$ must (i) endure the waiting time for an object from $W(k+1)$, and then (ii) either experience the additional (continuation) payoff of being in position k, or immediately receive an object from $W(k+1) \setminus W(k)$. Since (i) and (ii) are independent, the total expected utility of (i) and (ii) is a product of two terms (Equation 5) the first of which is given by Equation 4. The second (parenthetical) term is the expected utility of the payoffs described in (ii). The recursive expression yields the following result (though we prove it non-recursively in the Appendix).

Theorem 5. *Fix a policy W and suppose agents have CARA utility functions with parameter α. The expected utility of an agent in position $k \in \mathbb{N}$ is*

$$U_k^W = \sum_{i \in W(k)} \frac{\mu_i}{\mu_{W(k)}}(-e^{-\alpha v_i}) \cdot \prod_{\ell \leq k : W(\ell) \ni i} \frac{\mu_{W(\ell)}}{\mu_{W(\ell)} - \alpha} \tag{7}$$

Theorem 5 allows us to determine how a "marginally influence" in deferral decisions impacts U_k^W. Fix positions $\ell < k$ and consider removing some object $i \in W(\ell) \setminus W(\ell - 1)$ only from $W(\ell)$ (so that position ℓ now defers type i instead of consuming it, and it is now accepted at position $\ell + 1$). This has two effects on the RHS expression in Equation 7.

- First, the ith term in the summation is subjected to one less "penalty term" $\mu_{W(\ell)}/(\mu_{W(\ell)} - \alpha) > 1$, which increases U_k^W. This is interpreted as follows: the agent in position k "gains in the short run" when an position ℓ must now defer object i to a later position, increasing the speed with which object i might be consumed.

- Second, however, $\mu_{W(\ell)}$ is *decreased*; this increases the $\mu_{W(\ell)}/(\mu_{W(\ell)} - \alpha)$ "penalty term" which applies to all (remaining) objects in $W(\ell) \setminus \{i\}$. The intuition for this change is that the agent in position k "loses in the long run" since, conditional on reaching position ℓ there is a longer expected wait to consume an object.

The sum of these two effects depends on the magnitude of v_i. The first (beneficial) effect is *lower* for higher values of v_i. Intuitively, this captures an agent's preference to correlate high object values with long waiting times. Roughly speaking, an agent in position k would prefer that earlier positions defer only the relatively lower-value objects in $W(k)$. In fact, imagine restricting the ability of an agent in position $k-1$ to defer an object $i \in W(k)$.

Strikingly, we show that this improves U_k^W *if and only if* it improves U_{k-1}^W. In that sense agents (positions) have aligned interests in deferral decisions.

Theorem 6 (aligned interests)**.** *Fix policy W, position $k \geq 2$, and (if one exists) object type $j \in W(k) \setminus W(k-1)$. Let $W'(k-1) = W(k-1) \cup \{j\}$, and $W'(\ell) = W(\ell)$ for all $\ell \neq k-1$, i.e. W' is obtained from W by allocating j to $k-1$ instead of to k. Then $U_k^W \geq U_k^{W'}$ if and only if $U_{k-1}^W \geq U_{k-1}^{W'}$.*

Thus positions $k-1$ and k "agree" on whether object j should be allocated to position $k-1$ or k. It is intuitive then that Theorem 6 leads to an analog of Theorem 3 under CARA utility. Namely there is a "no-influence equilibrium" policy W_α^* that Pareto-dominates every other policy. It is constructed by sequentially finding sets $W_\alpha^*(k)$ that maximize the expected utility to positions $k = 1, 2, 3, \ldots$, subject to the constraint of the earlier positions' "choices."

Corollary 4 (Pareto-dominance)**.** *Suppose agents' preferences are described by a CARA utility function with (common) parameter α. Consider a (generically unique) policy W_α^* defined by sequentially maximizing U_k^W (using Equation 7) for $k = 1, 2, 3, \ldots$. Then W_α^* Pareto-dominates every other policy.*

Proof. By construction no policy can improve upon $U_1^{W_\alpha^*}$. Consider position 2. Fixing $W_\alpha^*(1)$, by construction we cannot improve upon $U_2^{W_\alpha^*}$ by adding any object $i \notin W_\alpha^*(2)$ to $W_\alpha^*(2)$. Fixing $W_\alpha^*(2)$, Theorem 6 implies that we cannot improve upon $U_2^{W_\alpha^*}$ by changing $W_\alpha^*(1)$ to any arbitrary $W(1) \subseteq W_\alpha^*(2)$. Sequentially repeating these arguments for positions $k = 3, 4, \ldots$ proves the result. $\square$

5 Discounted payoffs

5.1 Intuition

Lastly we consider agents who discount future payoffs.[26] With the introduction of discounting, our earlier results begin to break down. For instance

[26]We separate the analyses of risk-aversion and discounting for two reasons. First the combination of assumptions yields to intractable payoff expressions. Secondly our objective is to demonstrate the contradictory effects of these two assumptions. This is most easily done by considering them as two separate departures from the baseline case of risk-neutral, non-discounting agents.

in this scenario, no single policy Pareto-dominates all others. In fact a new effect emerges in which agents have *opposed* interests in influencing deferral decisions, opposite of the effect described in Theorem 6.

To establish some intuition for how discounting plays a role, compare the following two policies defined for some fixed set $\hat{O} \subseteq \mathcal{O}$.

$$W(1) = \cdots = W(k) = \hat{O}$$
$$W'(1) \subsetneq \cdots \subsetneq W'(k) = \hat{O}$$

Policy W is a no-deferrals policy—agents in positions 1 through k wait until reaching position 1 before receiving the next object from $\hat{O}$)—while W' could potentially assign the next object from $\hat{O}$ to any position 1 through k. These policies offer position k identical expected waiting time (Corollary 1) and distribution over allocated object (Theorem 1). However they differ in both the probability distribution of waiting costs and the correlation between consumed object value and waiting time.

An agent who discounts waiting costs would prefer to reduce the variance of those costs.[27] If we could hold everything else constant, this would suggest a preference for W over W' (see Corollary 3). Of course we cannot hold everything else constant because, as we vary the policy, this also changes the timing with which an agent might receive different kinds of objects. To consume an object from the set $W'(1)$, more time (on average) would pass under policy W' than under W. Similarly an object from $W'(k) \setminus W'(k-1)$ is consumed more quickly under W' than W. If W' offers a lottery over either a good object sooner or a worse object later, then W' could be preferable to W. Conversely if W' is a lottery resulting either in a long wait for better objects or a shorter wait for worse objects, the tends to be worse than the lottery induced by W.

This intuition underlies our result that agents have "opposed interests" in roughly the following sense. Consider a policy that changes only by assigning some object type $\alpha \in \mathcal{O}$ to position ℓ that was previously assigned to some position $k > \ell$. *An agent in any position $k' \geq k$ gains from this change if and only if an agent in position ℓ is harmed by this policy change.* This result is fairly striking. It implies that an agent in position k benefits from a policy change when positions ahead of him are no longer permitted to defer an

[27] E.g. the discounted cost of waiting for two periods is lower than the expected discounted cost of waiting for either one or three periods with equal probability.

object that otherwise would have been consumed at position k, *even though he will be subject to this same constraint* should he reach the earlier position.

5.2 Results

We suppose that agents continuously discount the future at a nominal interest rate r per unit of time: The present value of receiving payoff of x at t units of time in the future is $x \cdot e^{-rt}$. Observe that if an agent incurs waiting costs for t units of time, where t is exponentially distributed with parameter μ, then the **expected NPV** (ENPV) of these costs is

$$1/(r+\mu). \tag{8}$$

Likewise if an agent is to receive an object of value v after waiting t such units of time then the ENPV of this payoff is

$$v\mu/(r+\mu). \tag{9}$$

Using these two facts, it is straightforward to derive the expected NPV of an agent's payoff from position $k = 1$ under policy W, denoted $ENPV_1^W$. It is the above ENPV of payoff from the eventual arrival of an object from $W(1)$ minus that of the waiting time for that object to arrive.

$$\begin{aligned} ENPV_1^W &= \left(\sum_{i \in W(1)} \frac{\mu_i}{\mu_{W(1)}} \cdot \frac{v_i \mu_{W(1)}}{r + \mu_{W(1)}} \right) - \frac{1}{r + \mu_{W(1)}} \\ &= \frac{\left(\sum_{i \in W(1)} \mu_i v_i\right) - 1}{r + \mu_{W(1)}} \end{aligned} \tag{10}$$

since the arrival time of an object from $W(1)$ is exponentially distributed with parameter $\mu_{W(1)}$.

While Equation 10 is a simple generalization of the risk-neutral expected payoff $\Pi(1; W)$ (Equation 1), the general description of $ENPV_k$ is not as elegant. The simplest way to think of $ENPV_2$ is to consider the agent in position 2 acquiring payoffs as follows. First the agent waits for an arrival of some type $i \in W(2)$. That object is allocated either to the agent in position 1 (if $i \in W(1)$) or to the agent in position 2. In the former case, we can think of the agent in position 2 also being immediately allocated a pseudo-object with value $v = ENPV_1$.

That is, define a list of object values v' such that $v'_i = ENPV_1$ for $i \in W(1)$ and $v'_i = v_i$ otherwise. Now $EPNV_2$ satisfies an analog of Equation 10 with respect to v' and $W(2)$, which we then rewrite in terms of v.

$$\begin{aligned} ENPV_2^W &= \frac{(\sum_{i \in W(2)} \mu_i v'_i) - 1}{r + \mu_{W(2)}} \\ &= \frac{1}{r + \mu_{W(2)}} \left[\left(\mu_{W(1)} \cdot ENPV_1 + \sum_{i \in W(2) \setminus W(1)} \mu_i v_i \right) - 1 \right] \end{aligned}$$

This recursion relation generalizes to

$$ENPV_k^W = \frac{1}{r + \mu_{W(k)}} \left[\mu_{W(k-1)} \cdot ENPV_{k-1} + \sum_{i \in W(k) \setminus W(k-1)} \mu_i v_i - 1 \right] \quad (11)$$

which we use to derive the following theorem.

Theorem 7 (discounted payoffs)**.** *Suppose agents continuously discount payoffs at nominal rate r per period and fix a policy, W. The expected NPV of the payoff to an agent in position k is*

$$ENPV_k^W = \frac{1}{\mu_{W(k)}} \sum_{\ell=1}^{k} \left[\left(\sum_{i \in W(\ell) \setminus W(\ell-1)} \mu_i v_i - 1 \right) \prod_{j=\ell}^{k} \frac{\mu_{W(j)}}{r + \mu_{W(j)}} \right] \quad (12)$$

This expression generalizes Equation 1 (the case $r = 0$). Equation 12 can be interpreted as follows. The ENPV to position k decomposes into k components corresponding to the possible positions ($\ell = 1, \dots, k$) the agent might yet reach before receiving an object. Each (ℓth) component has

- a "value part" $(\sum_{i \in W(\ell) \setminus W(\ell-1)} \mu_i v_i)$, the expected value of objects assigned to that position, and
- a "cost part" (-1) the passing through that position.

More subtle is the final product term, whose interpretation differs across these two parts. Applied to the value part, the product term is the amount by which the future object value is discounted to the present (see Equation 9). Applied to the cost part, the product term's denominator is the amount by which future flow costs are discounted to the present (see Equation 8), while its numerator is the the probability of reaching any such position $\ell \leq k$ (and thus incurring those waiting costs).

Strikingly, this payoff expression can be used to show Theorem 8, stating that agents in consecutive queue positions have "opposed interests" in the following sense. Fix a policy W and object type α assigned to position $k+1$ (so $\alpha \in W(k+1) \setminus W(k)$), and consider a policy change in which α is instead assigned to position k (so $\alpha \in W'(k) \setminus W'(k-1)$). It turns out that this change would make the agent in position $k+1$ strictly better off *if and only if* it would make the agent in position k strictly worse off. Recall by Theorem 2 that when $r = 0$, the agent in position $k+1$ necessarily would be indifferent about such a change, making the statement a counterfactual. When $r > 0$, however, the agent in position $k+1$ is typically not indifferent.

Theorem 8 (Opposed interests)**.** *Fix a policy W, position k, and object type $\alpha \in W(k+1) \setminus W(k)$. Let policy W' be identical to W except that $\alpha \in W'(k) \setminus W'(k-1)$. Then*

$$ENPV_{k+1}^{W'} > ENPV_{k+1}^{W} \iff ENPV_{k}^{W'} < ENPV_{k}^{W} \tag{13}$$

The "opposed interests" described in Theorem 8 actually occur more broadly than described. Specifically, moving object α into $W(k)$ as described in the theorem clearly has no impact on earlier positions (1 through $k-1$). However, the agent in position $k+1$ strictly prefers this change *if and only if* the agent in position $k+2$ strictly prefers it as well. This can be seen via Equation 11: since $W(k+2)$ and $W(k+1)$ are being held constant, $ENPV_{k+2}$ is affected only by changes in the term $ENPV_{k+1}$. Iterating this argument, position $k+1$ strictly prefers the change if and only if any position $\ell \geq k+1$ prefers it. To summarize, Equation 13 can be strengthened to

$$\forall \ell \geq k+1 \; [ENPV_{\ell}^{W'} > ENPV_{\ell}^{W} \iff ENPV_{k}^{W'} < ENPV_{k}^{W}] \tag{14}$$

Thus there is typically no single policy that emerges as Pareto-dominant for discounting agents. A Pareto-dominant policy would have the threshold structure of policies constructed in Section 3 or Section 4, where later positions are assigned relatively worse objects that are deferred by earlier positions. By Theorem 8, an agent in a later position would prefer a policy in which earlier agents cannot defer such objects, contradicting Pareto-dominance. This observation is relevant in our next section, where we turn to the organ spoilage problem discussed in Section 1. When discounting is a significant part of agents' preferences, agents later in the queue unambiguously benefit from certain restrictions on earlier positions' deferrals.

6 Application: Organ Spoilage

We turn to our main motivating application, the organ spoilage problem described in Section 1. Current policy requires the planner to adhere to first-come-first-served principles, i.e. organs are to be offered according to patients according to an exogenous priority order. This leads to a well documented inefficiency. In the time it takes for, say, $\hat{k}$ agents to reject the offer of a (lower quality) organ, the organ spoils and cannot be used, even though an agent in position $\ell > \hat{k}$ might have been willing to accept it.

There are various ways to possibly address the spoilage problem. A planner could go so far as to abandon the principle of priority-based allocation, e.g. offering lower-quality organs directly to agents later in the queue. While such a solution seems to make sense in theory, it also has its drawbacks. First it requires the planner to evaluate the quality of each arriving organ rather than the patient, when such an evaluation might best be done by the agent (e.g. the patient's own doctor). Second this solution invites legal challenges from agents who believe that they were wrongly denied access to an organ that went to a lower-priority agent. E.g. if an organ is offered first to an agent in position k in the queue and accepted, an agent in position $k-1$ would justifiably complain that the organ was not offered to him.

A less extreme solution is to (partially) limit (influence) the degree to which agents defer organs. This solution preserves the spirit of first-come-first-served (priority-based) allocation, and is less susceptible to the criticisms above. We show that the desirability of such a solution depends (positively) on the degree to which agents discount the future, and (negatively) on the level of agents' risk aversion.

6.1 A simple model of spoilage

To illustrate these two effects we consider a scenario in which an object spoils after it has been offered to and deferred by the agents in positions up to and including $\hat{k}$; thus it is unavailable to the agents in positions $\ell \geq \hat{k}$. Observe that this model of spoilage can be expressed simply as a "spoilage constraint" on waiting list policies: we require that for all $\ell > \hat{k}$, $W(\ell) = W(\hat{k})$. No arriving object can be assigned directly to a position $\ell > \hat{k}$. To demonstrate our main points it is sufficient to consider the case $\hat{k} = 1$, i.e. an object spoils immediately after the first agent defers it. This assumption merely simplifies exposition, as the effects we demonstrate can also occur for $\hat{k} > 1$. Similarly

$k = 1$	$k \geq 2$
$W^S(1) = \{1\}$	$W^S(k) = \{1\}$
$W^I(1) = \{1, 2\}$	$W^I(k) = \{1, 2\}$
$W^*(1) = \{1\}$	$W^*(k) = \{1, 2\}$

Figure 1. Policies representing consumption in three scenarios: (W^S) spoilage occurs; (W^I) deferrals are disallowed; (W^*) spoilage does not occur.

it is sufficient to consider only two object types.

Formally, fix two object types, $\mathcal{O} = \{1, 2\}$ with values $v_1 > v_2$, and arrival rates $\mu_1, \mu_2 > 0$. Along with the baseline case of preferences ($\alpha = r = 0$), we alternately consider risk-averse agents ($\alpha > 0$) or discounting agents ($r > 0$). Once we fix all of these parameters of the model, we say that they exhibit a **spoilage problem** for position 2 when

(I) an agent in position $\hat{k} = 1$ would prefer to defer (only) type 2 objects, and

(II) an agent in position $k = 2$ would prefer to receive a type 2 object immediately rather than wait to receive the *second* arrival of a type 1 object.

Condition (I) says that in a decentralized (uninfluenced) equilibrium, type 2 objects would spoil. Condition (II) says that this spoilage causes inefficiency. Combined, these two conditions imply that v_2 is neither very high (the first condition would fail) nor very low (the second condition would fail). We can more generally define the weaker condition of a "spoilage problem for position $k > 2$" in which an agent in position k prefers instant consumption of v_2 over waiting for the kth arrival of a type 1 object.

We now imagine a planner attempting to solve a spoilage problem by removing the ability of agents to defer objects. To determine whether such a solution makes sense from a welfare perspective, we need to compare agents' payoffs across three different scenarios summarized in Figure 1.

First imagine that agents are uninfluenced, selfishly optimizing their deferral decisions. By condition (I) above, type 1 objects are assigned to position 1 and type 2 objects *spoil*. Thus agents' consumption is described by policy W^S (see Figure 1).

Second imagine the planner exerting *influence* in order to prevent deferrals.[28] In this case, the agent in position 1 must consume whichever object arrives next, resulting in policy W^I.

Third consider the alternate scenario in which spoilage no longer occurs (and deferrals are permitted). By conditions (I) and (II), type 2 objects now go to the agent in position 2, resulting in the policy we label W^*.[29]

To consider the implications of disallowing position 1's deferrals, we wish to compare payoffs under W^S and W^I. However we need to perform this comparison only when a **spoilage problem** exists in the first place. Condition (I) states that an agent in position 1 prefers W^S (or W^*) to W^I. Condition (II) states that an agent in position 2 prefers W^* to W^S. Given these two conditions, we also ask whether

(III) an agent in position $k = 2$ would prefer policy W^I to W^S.

Condition (III) would imply that, indeed, the agent in position 2 benefits when any agent who reaches position 1 is prevented from deferring an organ.

It would be significant if we face a spoilage problem in which Condition (III) fails. This means that, in an attempt to improve welfare by reducing deferrals, a planner could paradoxically harm agents in later positions *despite* the statistical evidence that spoilage rates decreased. It turns out that such spoilage problems do exist, though they appear to be "rare." Using our earlier results we can show that Condition (III) has to do with the agents' risk-aversion (α), and not their impatience (r) as we now explain.

Patient, risk-neutral agents. Consider our simplest setup of risk-neutral, non-discounting agents. By Theorem 3, W^* Pareto-dominates all other policies. By Theorem 2, any agent starting in position $k \geq 2$ is indifferent between W^* and W^I, and hence prefers W^I to W^S. Thus Condition (III) is always satisfied. In fact the more striking consequence of Theorem 2 is that *removing position 1's deferral option makes all other positions as well off as if the spoilage problem never existed in the first place.*

Impatient agents. Discounting merely serves to amplify the conclusion of the baseline case. By assumption, the agent in position 1 prefers W^* to W^I.

[28]As discussed earlier this could be done through indirect influence (doctors advising patients), or directly by penalizing agents who defer.

[29]This notation is chosen intentionally: ignoring the spoilage constraint it is the Pareto-dominant policy for non-discounting agents. See Corollary 4.

With discount rate $r > 0$, Theorem 8 and Condition (I) imply that an agent in position 2 prefers W^I over W^* (and hence over W^S by Condition (II)).[30] Again Condition (III) is always satisfied. Now, however, position 2 can even be *strictly* better off under W^I than when the spoilage problem does not exist at all. Repeated application of Equation 11 extends this conclusion to positions $k = 3, 4, \ldots$ also prefer W^I to W^*.

Risk-averse agents. Risk-aversion can cause Condition (III) to fail. Fortunately it appears that such cases are atypical, requiring either extreme levels of risk-aversion or that the cost of the spoilage problem is low. To establish an intuition, observe that restricting position 1's option to defer type 2 objects (moving from W^* to W^I) strictly lowers position 1's payoff, so by Theorem 6 position 2 also strictly prefers W^* to W^I. If the "cost of spoilage" is low to position 2 (i.e. W^* is only slightly better than W^S) then an agent in position 2 could prefer W^S to W^I: he would prefer suffering the spoilage problem over the "solution" of requiring any agent in position 1 to accept lower quality objects.

Indeed if v_2 is sufficiently close to v_1, then this is intuition is always accurate. For low levels of risk-aversion, however, v_2 needs to be very close to v_1 in order for this to happen. The following result formalizes this.

Proposition 1. *Fix two object types and define policies W^*, W^S, and W^I as in Figure 1. Fix values $v_1 > v_2$, arrival rates $\mu_1, \mu_2 > 0$, and CARA parameter $0 < \alpha < \mu_1$. Then*

1. $U_1^{W^S} \geq U_1^{W^I}$ *(Condition (I) is satisfied),*
2. $U_2^{W^*} \geq U_2^{W^S}$ *(Condition (II) is satisfied), and*
3. $U_2^{W^S} \geq U_2^{W^I}$ *(Condition (III) fails)*

if and only if

$$\frac{1}{\alpha} \log \left(\frac{\mu_1(\mu_1^2 - \alpha^2 + \mu_2\mu_1)}{(\mu_1 + \mu_2)(\mu_1 - \alpha)^2} \right) \leq v_1 - v_2 \leq \frac{2}{\alpha} \log \left(\frac{\mu_1}{\mu_1 - \alpha} \right) \tag{15}$$

[30] More generally, that theorem implies that an agent in position k is better off if we force an agent in position $k - 1$ to accept an object that the agent in position k *would have been willing to accept* in a scenario where spoilage does not occur.

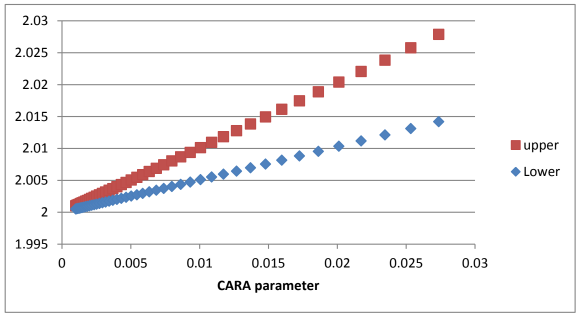

Figure 2. An example of the bounds of Equation 15, as a function of α, for the case $\mu_1 = \mu_2 = 1$. For instance when $\alpha = .02$, $v_1 - v_2$ must lie roughly between 2.01 and 2.02. between

It can be checked that the width of the band described by Equation 15 (i.e. the upper bound minus the lower bound) is

$$\frac{1}{\alpha}\log\left(\frac{\mu_1(\mu_1+\mu_2)}{\mu_1(\mu_1+\mu_2)-\alpha^2}\right)$$

This is positive for any $\alpha > 0$, so for any arrival rates $\mu_1 > \alpha$ and μ_2 there exist object values v_1, v_2 jointly satisfying the inequalities. The width of this range decreases in α, shrinking to zero as α converges to zero; see Figure 2. This is not surprising since Condition (III) cannot fail under risk-neutrality. More generally both the upper and lower bounds in Equation 15 converge to $2/\mu_1$ as α converges to zero. Intuitively, as the agent in position 2 becomes risk-neutral he becomes indifferent between receiving a type 2 object immediately versus paying an expected $2/\mu_1$ additional units of waiting costs to receive an object worth $v_1 - v_2$ additional units of value. One can also verify that the width of the band is decreasing in both arrival rates μ_1, μ_2.

7 Conclusion

We have considered the welfare implications of arbitrarily influencing or constraining the deferral decisions of prioritized agents who are in a waiting list for randomly arriving objects. Payoff expressions for both risk-averse agents and for impatient (discounting) agents show that these two characteristics of preferences lead to different prescriptions for the planner. Particularly, our first main results (Theorem 3 and Corollary 4) show that such influence has an unambiguously negative welfare effect for risk-averse (but patient) agents. On the other hand when (risk-neutral) agents are impatient, agents later in the waiting list would typically benefit when deferral decisions are constrained for any agents who reach earlier positions in the waiting list (Theorem 8).

We then apply these results to the application of the organ spoilage problem. The spoilage problem can be modeled in our setup as a constraint where agents sufficiently late in the waiting list are forced to defer all objects. First, when agents are impatient (and sufficiently risk-neutral), it is generally welfare-improving to at (least partially) limit the deferral options of agents early in the waiting list. At the time such a policy change is implemented, the agents early in the waiting list would indeed suffer a welfare loss from this change, but all remaining agents in the queue would benefit.

On the other hand, risk-aversion plays the opposite role. While a reduction in deferrals partially "solves" the spoilage problem by decreasing wastage and increasing the organ yield rate, this efficiency gain is partially offset by the welfare loss incurred when we prevent risk-averse agents from selfishly optimizing their deferral decisions. Any policy change that hurts agents in earlier queue positions typically also has a negative welfare effect on agents in later positions, i.e. risk-averse agents have what we called "aligned interests" (Theorem 6). We even show that, under some parameter values, this negative welfare effect can dominate the efficiency gains for at least some late positions in the waiting list. Fortunately such parameter values seem to be the exception rather than the norm. Nevertheless this suggests that welfare analyses in such settings must carefully consider the degree to which agents are risk-averse and/or impatient. Under risk-aversion, the planner must measure more than an organ yield rate in order to determine the general welfare implications from constraining the agents' rights to defer organs of marginal quality.

References

Agrawala, A.K., E.G. Coffman, M.R. Garey, and S.K. Tripathi, 1984. "A Stochastic Optimization Algorithm Minimizing Expected Flow Times on Uniform Processors." *IEEE Transactions on Computers*, 33(4):351–356.

Akbarpour, Mohammad, Shengwu Li, and Shayan O. Gharan, 2014. "Dynamic Matching Market Design." mimeo.

Bloch, Francis and David Cantala, 2013. "Markovian assignment rules." *Social Choice and Welfare*, 40:1–25.

Bloch, Francis and David Cantala, 2014. "Dynamic Allocation of Objects to Queuing Agents: The Discrete Model." Mimeo.

Coffman, E.G., L. Flatto, M.R. Garey, and R.R. Weber, 1987. "Minimizing Expected Makespans on Uniform Processor Systems." *Advances in Applied Probability*, 19:177–201.

Damiano, Ettore, and Ricky Lam, 2005. "Stability in dynamic matching markets." *Games and Economic Behavior*, 52(1):34–53.

Doval, Laura, 2014. "A Theory of Stability in Dynamic Matching Markets." mimeo.

Hassin, Refael, 1985. "On the Optimality of First Come Last Served Queues." *Econometrica*, 53(1):201–202.

Kadam, Sangram V., and Maciej H. Kotowski, 2014. "Multi-period Matching," mimeo.

Kennes, John, Daniel Monte, and Norovsambuu Tumennasan, 2014. "The day care assignment: A dynamic matching problem." *American Economic Journal: Microeconomics*, 6(4): 362-406.

Kumar, P.R. and J. Walrand, 1985. "Individually Optimal Routing in Parallel Systems." *Journal of Applied Probability*, 22(4):989–995.

Kurino, Morimitsu, 2009. "Credibility, efficiency, and stability: A theory of dynamic matching markets." Jena economic research papers.

Leshno, Jacob, 2014. "Dynamic Matching in Overloaded Systems." Mimeo, Columbia University.

Marx, Philip and James Schummer, 2016. "Revenue from Matching Platforms." Mimeo, Northwestern University.

Naor, P., 1969. "The Regulation of Queue Size by Levying Tolls." *Econometrica*, 37:15–24.

Sack, Kevin, 2012. "In Discarding of Kidneys, System Reveals Its Flaws," *The New York Times*, Sept 19, 2012.

Pereyra, Juan, 2013. "A dynamic school choice model." *Games and Economic Behavior*, 80:100–114.

Roth, Alvin E. 2007. “Repugnance as a Constraint on Markets.” *Journal of Economic Perspectives*, 21(3):37–58.

Su, Xuanming and Stefanos Zenios, 2004. “Patient Choice in Kidney Allocation: The Role of the Queueing Discipline.” *Manufacturing & Service Operations Management*, 6(4):280-âĂŞ301.

Thakral, Neil, 2015. “Matching with Stochastic Arrival.” mimeo.

Ünver, Utku, 2010. “Dynamic kidney exchange.” *The Review of Economic Studies*, 77(1): 372–414.

8 Appendix

8.1 Proofs

Proof of Theorem 2. The proof is by induction. For $k = 1$, the agent consumes the first arrival from $W(1)$, so the expected object value minus the expected waiting time is

$$\Pi(1;W) = \frac{\sum_{j\in W(1)} \mu_j v_j}{\mu_{W(1)}} - \frac{1}{\mu_{W(1)}} = \frac{\sum_{j\in W(1)} \mu_j v_j - 1}{\mu_{W(1)}}$$

consistent with Equation 1.

Fix $k \in \mathbb{N}$ and suppose that Equation 1 holds for $k-1$. The next object-type to arrive that belongs to $W(k)$ either belongs to $W(k-1)$ or to $W(k) \setminus W(k-1)$. In the former case the agent in position k moves to position $k-1$ and continues with an additional expected continuation payoff $\Pi(k-1;W)$. In the latter case the agent is assigned the object, receiving payoff v_j. Accounting for these two possibilities, along with the expected waiting time for the arrival from $W(k)$, we have the following.

$$\begin{aligned}
\Pi(k;W) &= \frac{\mu_{W(k-1)} \cdot \Pi(k-1;W)}{\mu_{W(k)}} + \frac{\sum_{W(k)\setminus W(k-1)} \mu_j v_j}{\mu_{W(k)}} - \frac{1}{\mu_{W(k)}} \\
&= \frac{\mu_{W(k-1)} \cdot \left(\frac{\sum_{W(k-1)} \mu_j v_j - (k-1)}{\mu_{W(k-1)}}\right) + \sum_{W(k)\setminus W(k-1)} \mu_j v_j - 1}{\mu_{W(k)}} \\
&= \frac{\sum_{W(k-1)} \mu_j v_j - (k-1) + \sum_{W(k)\setminus W(k-1)} \mu_j v_j - 1}{\mu_{W(k)}} \\
&= \frac{\sum_{W(k)} \mu_j v_j - k}{\mu_{W(k)}}
\end{aligned}$$

proving the result. $\square$

Proof of Lemma 1. Fix k, and for any subset of types $C \subseteq \mathcal{O}$, consider the no-deferrals policy W defined by $W(\ell) \equiv C \neq \emptyset$. Rather than writing $\Pi(k;W)$, let $\pi(C)$ denote the expected payoff to position k under such a policy, since we consider varying C.

From Theorem 2,

$$\pi(C) = \frac{\sum_{j\in C} \mu_j v_j - k}{\mu_C}$$

and for any $i \in \mathcal{O} \setminus C$, adding i to C yields a payoff of

$$\pi(C \cup \{i\}) = \frac{\sum_{j \in C} \mu_j v_j - k + \mu_i v_i}{\mu_C + \mu_i}$$

which (weakly) improves on $\pi(C)$ if and only if $v_i \geq (\sum_{j \in C} \mu_j v_j - k)/\mu_C$. Since object types are in decreasing order of the v_i's, any W^* defined via Equation 2 maximizes $\Pi(k; \cdot)$.[31]

Finally, observe that the right-hand side of the inequality within Equation 2 is decreasing in k. Therefore the type index i_k^* is indeed increasing in the position index k. □

Proof of Theorem 1. The statement is obviously true when $k = 1$. Inductively, fix a k and suppose that the statement is true for any $k' < k$. Nothing happens for the agent in position k until the arrival of some object type in $W(k)$. Upon the arrival of such an object, the probability it is of type $i \in W(k)$ is $\mu_i/\mu_{W(k)}$. If $i \in W(k) \setminus W(k-1)$ then the agent consumes that object (and otherwise cannot consume that object type), proving the claim for $i \in W(k) \setminus W(k-1)$.

Otherwise $i \in W(k-1)$, so the agent moves into position $k-1$; that is, the total probability of moving into position $k-1$ is $\sum_{j \in W(k-1)} \mu_j/\mu_{W(k)}$. By the induction assumption, the probability of eventually consuming any $j \in W(k-1)$ given that that the agent starts in position $k-1$ is $\mu_j/\mu_{W(k-1)}$. Hence the probability of ultimately consuming $j \in W(k-1)$ conditional on starting in position k is

$$\frac{\mu_{W(k-1)}}{\mu_{W(k)}} \cdot \frac{\mu_j}{\mu_{W(k-1)}} = \frac{\mu_j}{\mu_{W(k)}}$$

proving the claim for $j \in W(k-1)$. □

Proof of Theorem 4. The wait time t_k^W is the sum of two independent random variables: the initial wait t' until the arrival of the next object $i \in W(k)$, and the remaining wait t'', which either has the same distribution as t_{k-1}^W (if $i \in W(k-1)$) or is degenerately $t'' = 0$ (if $i \in W(k) \setminus W(k-1)$).

Since t' is exponentially distributed,

$$\mathrm{Var}(t') = 1/\mu_{W(k)}^2.$$

[31](Ties are irrelevant.) In the nongeneric case that $v_{i_k^*} = \pi(W^*(k))$, it is easy to see that $W'(k) \equiv W^*(k) \setminus \{i_k^*\}$ also maximizes k's payoff. This impacts neither the Lemma nor any other results of the paper.

To consider the variance of t'', we first recall the following easily proven fact. Let a random variable Y equal the value of some r.v. X with probability p and be degenerately $Y = 0$ with probability $1 - p$. Then

$$\mathrm{Var}(Y) = p\mathrm{Var}(X) + (p - p^2)E(X)^2$$

Here,

$$\mathrm{Var}(t'') = \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t^W_{k-1}) + \left(\frac{\mu_{W(k-1)}}{\mu_{W(k)}} - \left(\frac{\mu_{W(k-1)}}{\mu_{W(k)}}\right)^2\right) E(t^W_{k-1})^2$$

By Corollary 1, $E(t^W_{k-1}) = (k-1)/\mu_{W(k-1)}$. Therefore

$$\begin{aligned}
\mathrm{Var}(t^W_k) &= \mathrm{Var}(t') + \mathrm{Var}(t'') \\
&= \frac{1}{\mu^2_{W(k)}} + \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t^W_{k-1}) + \left(\frac{\mu_{W(k-1)}}{\mu_{W(k)}} - \frac{\mu^2_{W(k-1)}}{\mu^2_{W(k)}}\right) \frac{(k-1)^2}{\mu^2_{W(k-1)}} \\
&= \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t^W_{k-1}) + \frac{(k-1)^2}{\mu_{W(k)}\mu_{W(k-1)}} - \frac{(k-1)^2 - 1}{\mu^2_{W(k)}} \qquad (16)
\end{aligned}$$

which we can solve recursively.

For any policy, t^W_1 is exponentially distributed with variance of $1/\mu^2_{W(1)}$ which coincides with Equation 3. We show that if Equation 3 holds for some arbitrary $k-1$ then it holds for k. Substituting into Equation 16,

$$\begin{aligned}
\mathrm{Var}(t^W_k) &= \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t^W_{k-1}) + \frac{(k-1)^2}{\mu_{W(k-1)}\mu_{W(k)}} - \frac{k^2 - 2k}{\mu^2_{W(k)}} \\
&= \frac{\mu_{W(k-1)}}{\mu_{W(k)}} \frac{1}{\mu_{W(k-1)}} \left(\sum_{\ell=1}^{k-2} \frac{2\ell}{\mu_{W(\ell)}} + \frac{2(k-1) - (k-1)^2}{\mu_{W(k-1)}}\right) + \frac{(k-1)^2}{\mu_{W(k-1)}\mu_{W(k)}} - \frac{k^2 - 2k}{\mu^2_{W(k)}} \\
&= \frac{1}{\mu_{W(k)}} \left(\sum_{\ell=1}^{k-2} \frac{2\ell}{\mu_{W(\ell)}} + \frac{2(k-1)}{\mu_{W(k-1)}} - \frac{k^2 - 2k}{\mu_{W(k)}}\right) \\
&= \frac{1}{\mu_{W(k)}} \left(\sum_{\ell=1}^{k-1} \frac{2\ell}{\mu_{W(\ell)}} + \frac{2k - k^2}{\mu_{W(k)}}\right) = \frac{1}{\mu_{W(k)}} \left(\sum_{\ell=1}^{k} \frac{2\ell}{\mu_{W(\ell)}} + \frac{-k^2}{\mu_{W(k)}}\right)
\end{aligned}$$

proving the result. □

Proof of Theorem 5. Fix W, α, and a position k. By Theorem 1 an agent in position k ultimately consumes object $i \in W(k)$ with probability $\mu_i/\mu_{W(k)}$.

Conditional on consuming $i \in W(k)$, the agent's waiting time is $t_k + t_{k-1} + \cdots + t_\ell$ where $i \in W(\ell) \setminus W(\ell - 1)$, and where t_j is exponentially distributed with parameter $\mu_{W(j)}$. This is because, in order to consume such an i, the agent must first advance to position ℓ in the queue and then receive an object, requiring waits for objects from $W(k)$, $W(k-1)$, ..., $W(\ell)$.

Denoting t as the total (unconditional) waiting time and v as the value of the received object, we have

$$\begin{aligned}
U_k^W \equiv E(u(v-t)) &= \sum_{\ell=1}^{k} \sum_{i\in W(\ell)\setminus W(\ell-1)} \frac{\mu_i}{\mu_{W(k)}} E(u(v_i - \tau_i)) \\
&= \sum_{\ell=1}^{k} \sum_{i\in W(\ell)\setminus W(\ell-1)} \frac{\mu_i}{\mu_{W(k)}} E(u(v_i))(-E(u(-\tau_i))) \\
&= \sum_{\ell=1}^{k} \sum_{i\in W(\ell)\setminus W(\ell-1)} \frac{\mu_i}{\mu_{W(k)}} u(v_i) \prod_{j=\ell}^{k} -E(u(-t_j)) \\
&= \sum_{\ell=1}^{k} \sum_{i\in W(\ell)\setminus W(\ell-1)} \frac{\mu_i}{\mu_{W(k)}} (-e^{-\alpha v_i}) \prod_{j=\ell}^{k} \frac{\mu_{W(j)}}{\mu_{W(j)} - \alpha}
\end{aligned}$$

where the second and third lines follow from Equation 5, and the last from Equation 4. For each i, the $\mu_{W(j)}/(\mu_{W(j)} - \alpha)$ term appears for each position $j \le k$ satisfying $i \in W(j)$, so the last line yields Equation 7. □

Proof of Theorem 6. Observe that $U_{k-1}^W \ge U_{k-1}^{W'}$ if and only if $U_{k-1}^W \ge u(v_j) = -e^{-\alpha v_j}$, i.e. $k-1$ prefers to defer j whenever the utility from v_j does not exceed the expected utility of continuing to wait. This follows intuitively but can also be derived from Equation 7. Therefore we need to show that $U_k^W \ge U_k^{W'}$ if and only if $U_{k-1}^W \ge u(v_i) = -e^{-\alpha v_i}$.

Observe that $\mu_{W'(k)} = \mu_{W(k)}$ and that $W(k) \setminus W(k-1) = \{j\} \cup (W'(k) \setminus W'(k-1))$. This cancels some terms in Equation 6, so that

$$U_k^W \ge U_k^{W'} \Leftrightarrow \frac{\mu_{W(k-1)}}{\mu_{W(k)}} U_{k-1}^W + \frac{\mu_j}{\mu_{W(k)}} u(v_j) \ge \frac{\mu_{W'(k-1)}}{\mu_{W'(k)}} U_{k-1}^{W'}$$

Since $\mu_{W'(k-1)} = \mu_{W(k-1)} + \mu_j$ the latter inequality becomes

$$\begin{aligned}
&\frac{\mu_{W(k-1)}}{\mu_{W(k)}} U_{k-1}^W + \frac{\mu_j}{\mu_{W(k)}} u(v_j) \ge \frac{\mu_{W(k-1)} + \mu_j}{\mu_{W(k)}} U_{k-1}^{W'}, \text{ or} \\
&\frac{\mu_{W(k-1)}}{\mu_{W(k-1)} + \mu_j} U_{k-1}^W + \frac{\mu_j}{\mu_{W(k-1)} + \mu_j} u(v_j) \ge U_{k-1}^{W'} \qquad (17)
\end{aligned}$$

Next we express $U^{W'}_{k-1}$ in terms of U^{W}_{k-1}. The following equation can be derived (tediously) from Equation 7; however it can be understood as follows. After adding j to $W(k-1)$, with probability $\frac{\mu_{W(k-1)}}{\mu_{W(k-1)}+\mu_j}$ the agent receives the payoff he would have received under W, and with the remaining probability he receives $u(v_j)$. In both cases the term $\frac{\mu_{W(k-1)}+\mu_j}{\mu_{W(k-1)}+\mu_j-\alpha}$ represents the waiting cost utility as in Equation 4. However in the former case, U^{W}_{k-1} is corrected for the fact that the waiting cost utility $\frac{\mu_{W(k-1)}}{\mu_{W(k-1)}-\alpha}$ no longer applies. In summary, we have

$$\begin{aligned}
U^{W'}_{k-1} &= \frac{\mu_{W(k-1)}}{\mu_{W(k-1)}+\mu_j} U^{W}_{k-1} \left[\frac{\mu_{W(k-1)}-\alpha}{\mu_{W(k-1)}} \frac{\mu_{W(k-1)}+\mu_j}{\mu_{W(k-1)}+\mu_j-\alpha} \right] \\
&\quad + \frac{\mu_j}{\mu_{W(k-1)}+\mu_j} u(v_j) \frac{\mu_{W(k-1)}+\mu_j}{\mu_{W(k-1)}+\mu_j-\alpha} \\
&= \frac{\mu_{W(k-1)}-\alpha}{\mu_{W(k-1)}+\mu_j-\alpha} U^{W}_{k-1} + \frac{\mu_j}{\mu_{W(k-1)}+\mu_j-\alpha} u(v_j)
\end{aligned}$$

Now Equation 17 becomes

$$\begin{aligned}
&\frac{\mu_{W(k-1)}}{\mu_{W(k-1)}+\mu_j} U^{W}_{k-1} + \frac{\mu_j}{\mu_{W(k-1)}+\mu_j} u(v_j) \\
&\qquad \geq \frac{\mu_{W(k-1)}-\alpha}{\mu_{W(k-1)}+\mu_j-\alpha} U^{W}_{k-1} + \frac{\mu_j}{\mu_{W(k-1)}+\mu_j-\alpha} u(v_j) \qquad (18)
\end{aligned}$$

which is true precisely when $U^{W}_{k-1} \geq u(v_i)$. □

Proof of Theorem 7. Equation 10 proves the case $k = 1$. Supposing Equation 12 holds for some k, we show it to hold for $k + 1$.

Upon the arrival of an object $i \in W(k + 1)$, the agent in position $k + 1$ either receives the object, or moves into position k. Conditional on the latter event (moving into position k), that agent's eventual (continuation) payoff has an expected NPV of $ENPV_k$ by definition. Hence, starting from position $k+1$, the agent incurs waiting costs until seeing an arrival of $i \in W(k+1)$ and then faces two possible lump sum payoffs: receiving v_i if $i \in W(k+1) \setminus W(k)$ or otherwise "receiving" $ENPV_1$ as an expected continuation payoff.

The expected NPV of waiting costs for an arrival from $W(k + 1)$ is

$$1/(r + \mu_{W(k+1)})$$

as described earlier. The expected NPV of the lump sum payoff is

$$\frac{\mu_{W(k)}}{\mu_{W(k+1)}} \cdot \frac{ENPV_k \cdot \mu_{W(k+1)}}{r + \mu_{W(k+1)}} + \sum_{i \in W(k+1) \setminus W(k)} \frac{\mu_i}{\mu_{W(k+1)}} \cdot \frac{v_i \mu_{W(k+1)}}{r + \mu_{W(k+1)}}$$

Combining these terms and substituting, we have

$$\begin{aligned}
ENPV_{k+1} &= \frac{\mu_{W(k)} \cdot ENPV_k}{r + \mu_{W(k+1)}} + \sum_{i \in W(k+1) \setminus W(k)} \frac{\mu_i v_i}{r + \mu_{W(k+1)}} - \frac{1}{r + \mu_{W(k+1)}} \\
&= \frac{1}{r + \mu_{W(k+1)}} \left[\frac{\mu_{W(k)}}{\mu_{W(k)}} \sum_{\ell=1}^{k} \left(\left(\sum_{i \in W(\ell) \setminus W(\ell-1)} \mu_i v_i - 1 \right) \prod_{j=\ell}^{k} \frac{\mu_{W(j)}}{r + \mu_{W(j)}} \right) \right. \\
&\quad \left. + \left(\sum_{i \in W(k+1) \setminus W(k)} \mu_i v_i - 1 \right) \right] \\
&= \frac{1}{\mu_{W(k+1)}} \left[\sum_{\ell=1}^{k+1} \left(\left(\sum_{i \in W(\ell) \setminus W(\ell-1)} \mu_i v_i - 1 \right) \prod_{j=\ell}^{k+1} \frac{\mu_{W(j)}}{r + \mu_{W(j)}} \right) \right]
\end{aligned}$$

yielding Equation 12 for $k+1$. □

Proof of Theorem 8. Fix notation as in the statement of the theorem. It should be intuitively clear that

$$ENPV_k^{W'} < ENPV_k^{W} \iff v_\alpha < ENPV_k^{W}$$

(Alternatively one can derive this fact using the same approach we use below to rewrite the inequality $ENPV_{k+1}^{W'} > ENPV_{k+1}^{W}$.)

Define $\mu_{W(j)}$'s with respect to W. Observe that they are the same for W' except at k, where it becomes $\mu_{W(k)} + \mu_\alpha$. Denote

$$X = \sum_{\ell=1}^{k-1} \left[\left(\sum_{i \in W(\ell) \setminus W(\ell-1)} \mu_i v_i - 1 \right) \prod_{j=\ell}^{k-1} \frac{\mu_{W(j)}}{r + \mu_{W(j)}} \right] + \sum_{i \in W(k) \setminus W(k-1)} \mu_i v_i - 1$$

Thus Equation 12 becomes $ENPV_k^W = X/(r + \mu_{W(k)})$. Similarly,

$$ENPV_{k+1}^{W} = \frac{1}{r + \mu_{W(k+1)}} \left[X \frac{\mu_{W(k)}}{r + \mu_{W(k)}} + \sum_{i \in W(k+1) \setminus W(k)} \mu_i v_i - 1 \right]$$

and, since W' moves α into $W'(k)$,

$$ENPV_{k+1}^{W'} = \frac{1}{r+\mu_{W(k+1)}}\left[(X+\mu_\alpha v_\alpha)\frac{\mu_{W(k)}+\mu_\alpha}{r+\mu_{W(k)}+\mu_\alpha} + \sum_{i\in W(k+1)\setminus W(k)} \mu_i v_i - \mu_\alpha v_\alpha - 1\right]$$

Canceling some terms, we have $ENPV_{k+1}^{W'} > ENPV_{k+1}^{W}$ being true if any only if

$$\begin{aligned}
(X+\mu_\alpha v_\alpha)\frac{\mu_{W(k)}+\mu_\alpha}{r+\mu_{W(k)}+\mu_\alpha} - \mu_\alpha v_\alpha &> X\frac{\mu_{W(k)}}{r+\mu_{W(k)}}\\
X\left(\frac{\mu_{W(k)}+\mu_\alpha}{r+\mu_{W(k)}+\mu_\alpha} - \frac{\mu_{W(k)}}{r+\mu_{W(k)}}\right) &> \mu_\alpha v_\alpha\left(1-\frac{\mu_{W(k)}+\mu_\alpha}{r+\mu_{W(k)}+\mu_\alpha}\right)\\
X\left(\frac{\mu_\alpha r}{(r+\mu_{W(k)}+\mu_\alpha)(r+\mu_{W(k)})}\right) &> \mu_\alpha v_\alpha\left(\frac{r}{r+\mu_{W(k)}+\mu_\alpha}\right)\\
ENPV_k^W = \frac{X}{(r+\mu_{W(k)})} &> v_\alpha
\end{aligned}$$

proving the theorem. $\square$

Proof of Proposition 1. By Equation 7,

$$\begin{aligned}
U_1^{W^*} = U_1^{W^S} &= \frac{\mu_1(-e^{-\alpha v_1})}{\mu_1-\alpha}\\
U_1^{W^I} &= \frac{\mu_1(-e^{-\alpha v_1})+\mu_2(-e^{-\alpha v_2})}{\mu_1+\mu_2-\alpha}
\end{aligned}$$

and for $k \geq 2$,

$$\begin{aligned}
U_k^{W^*} &= \frac{\mu_1(-e^{-\alpha v_1})}{\mu_1+\mu_2}\left(\frac{\mu_1}{\mu_1-\alpha}\right)\left(\frac{\mu_1+\mu_2}{\mu_1+\mu_2-\alpha}\right)^{k-1} + \frac{\mu_2(-e^{-\alpha v_2})}{\mu_1+\mu_2}\left(\frac{\mu_1+\mu_2}{\mu_1+\mu_2-\alpha}\right)^{k-1}\\
U_k^{W^S} &= (-e^{-\alpha v_1})\left(\frac{\mu_1}{\mu_1-\alpha}\right)^k\\
U_k^{W^I} &= \frac{\mu_1(-e^{-\alpha v_1})+\mu_2(-e^{-\alpha v_2})}{\mu_1+\mu_2}\cdot\left(\frac{\mu_1+\mu_2}{\mu_1+\mu_2-\alpha}\right)^k
\end{aligned}$$

Therefore we have

$$\begin{aligned}
U_1^{W^S} \geq U_1^{W^I} &\iff v_1 - v_2 \geq \frac{1}{\alpha} \log\left(\frac{\mu_1}{\mu_1 - \alpha}\right) \\
U_2^{W^*} \geq U_2^{W^S} &\iff v_1 - v_2 \leq \frac{2}{\alpha} \log\left(\frac{\mu_1}{\mu_1 - \alpha}\right) \\
U_2^{W^S} \geq U_2^{W^I} &\iff v_1 - v_2 \geq \frac{1}{\alpha} \log\left(\frac{\mu_1(\mu_1^2 - \alpha^2 + \mu_2\mu_1)}{(\mu_1 + \mu_2)(\mu_1 - \alpha)^2}\right)
\end{aligned}$$

Since $\alpha < \mu_1 + \mu_2$, the third condition implies the first one. Hence all three conditions are satisfied when (15) is true. □