Influencing Waiting Lists

James Schummer*

April 20, 2021

Abstract

Stochastically arriving objects (e.g. transplant organs, public housing units) often are allocated via waiting lists exhibiting *deferral rights*: agents may decline offers, keeping their position in line. We consider the welfare implications of bestowing or constraining such rights, concluding that their desirability depends—in opposite ways—on agents' risk-aversion and impatience. Under risk-aversion, uninfluenced deferral rights typically enhance welfare. Under discounting some restrictions on deferral rights can benefit all agents joining the list. In a stylized "organ spoilage" model our results demonstrate that policy evaluations should not be based solely on throughput metrics (e.g. organ utilization rates) that ignore such preference characteristics.

Keywords: waiting lists, queueing, dynamic assignment.
JEL numbers: C78, D47.

*MEDS, Kellogg School of Management, Northwestern University, Evanston IL 60201, USA. Email: schummer@kellogg.northwestern.edu.

1 Introduction

Waiting lists are used to allocate various kinds of objects that arrive over time, e.g. transplant organs, public housing units, openings in substance abuse treatment programs, and vacancies in child care facilities. As objects become available, they are offered to priority-ordered agents who may not be obligated to accept the offer. In applications where agents are granted *deferral rights*, they may decline offered objects while maintaining their relative position in the list. A deferred object is offered to the next agent in line, who also may have deferral rights. We highlight novel welfare implications that arise from granting or restricting such rights.

Conditional on reaching an early position of the waiting list, an agent clearly benefits from deferral rights: the freedom to choose whether to wait for a better offer. For agents in later positions the welfare implications are ambiguous without further assumptions on preferences. Previous work shows that welfare can be improved by allowing or even encouraging deferrals when agents have heterogeneous preferences (Bloch and Cantala (2013), Leshno (2019)) or have heterogeneous outside options (Arnosti and Shi (2020). Our contribution is to show that, even when agents are completely homogeneous, the desirability of influencing deferral rights is tied to preference characteristics that so far have not been focal in the study of wait list mechanisms—risk tolerance and discounting.

We consider the following kind of thought experiment. Imagine asking an agent already occupying some position in a waiting list to selfishly decide whether, as a matter of policy, deferral rights should be granted to all agents or to none. How would the agent's answer depend on her current position in the queue or on preference characteristics? More generally, would the agent prefer that deferral rights only be *partially* granted, e.g. only over certain types of objects or from certain positions of the list? We capture this thought experiment via a class of allocation schemes—deemed *waiting list policies*—that depart from first-come-first-served allocation systems by giving the planner a wide range of control over agents' deferral decisions. This control can be any arbitrary function of the object type being offered and/or the agent's current position in the list.

There are two ways to interpret these controls or restrictions on deferral decisions: as *explicit rules* enforced by the planner, or as the planner's indirect *influence* over an agent's ability or incentive to make decisions. The "rules" interpretation is straightforward:,deferrals can be constrained essentially by force, such as in applications where public housing applicants must accept the next offered unit or depart the queue.

The "influence" interpretation can encompass various ways that a planner indirectly alters deferral decisions. Literal examples of influence exist wherever imperfectly informed agents receive advice on their deferral decisions. A primary example is organ transplant patients whose deferral decisions are heavily influenced by their health care providers.[1] A more subtle example of influence would be to obscure the quality of offered objects, reducing the agents' ability to discriminate amongst them when making deferral decisions.

Our approach captures these direct and indirect forms of "influencing" deferral decisions, bypassing any specification of *how* the planner might accomplish such outcomes. For the sake of generality, we set aside the question of what set of waiting list policies might be "implementable" by the planner since our results would apply to any of them.

1.1 Overview of results

We provide welfare results from both an interim perspective (the welfare of agents conditional on their current positions in the list) and an ex ante perspective (the welfare of agents as they arrive to the list). The former perspective (Sections 3–5) is relevant when the planner is primarily concerned only with the welfare of agents already present in a waiting list, e.g. if the political feasibility of a policy change depends mainly on such agents. These interim welfare results are then used to derive welfare results under an ex ante perspective (Section 6), e.g. where the planner is primarily concerned with the long run expected welfare of agents who join the list over time.

In more detail, we start with a benchmark case of risk-neutral, non-discounting agents where a "payoff equivalence" result (Theorem 1) implies the following. Any (marginal) influence over deferral decisions in early positions of the waiting list has no effect on the expected payoffs to agents in later positions of the list. Later agents face the same expected waiting time and the same distribution over object consumption as if the planner had committed not to influence these decisions. Consequently, a "no influence"

[1] As the Washington Post (Kindy et al. (2018)) reports, "Some transplant centers are working hard to persuade patients to accept less-than-perfect organs" due to organ shortages. In Section 7 illustrates a tradeoff between the welfare gains achievable from such persuasion and the ethical principal to advise individual patients in their own best interest.

policy—fully allowing selfish deferral decisions—maximizes expected interim payoffs starting from any position in the list (Theorem 2).

Under more general preferences this payoff equivalence breaks down since restricting deferral decisions (at early positions) alters the *joint* distribution of waiting time and object type consumption (at later positions). For example, deferral restrictions at earlier positions typically reduces the variance of waiting times for later agents (Corollary 2). Though this effect seems to benefit risk-averse agents, there is an additional effect. Deferral restrictions typically reduce the correlation between later agents' waiting time (a cost) and object quality (a benefit), harming risk-averse agents. Under CARA preferences we show that this latter effect dominates via an "aligned interests" result (Theorem 3): agents in consecutive positions of the list have common preference over (marginal) changes to deferral rights policies. It follows that uninfluenced deferral rights enhance welfare: a "no influence" policy strictly maximizes each position's interim expected utility (Theorem 4).

In contrast, when agents discount future payoffs we obtain an "opposed interests" result (Theorem 5): a marginal change in deferral rights between positions k and $k+1$ lowers position k's (expected, discounted) interim payoff if and only if it increases it for *all* positions $k+1$ and later. The main takeaway of this result is that, by committing to influence deferral decisions at some early position k of the waiting list, we can improve welfare for (typically many) agents who occupy or join the list at positions $k+1$ and beyond.

When considering the (ex-ante) welfare of agents arriving over time, observe that deferral restrictions benefit those agents by shortening the list (and hence their average waiting time). In the risk-averse case this creates a tradeoff with the optimality result of Theorem 4. In Section 6 we show that this tradeoff is resolved by (partially) restricting deferral decisions only in later positions of the waiting list, while the "no influence" prescription of Theorem 4 continues to hold in earlier positions.

In Section 7 we apply our results to a stylized "organ spoilage problem:" objects with limited shelf life must be discarded after being offered to and deferred by many agents. In the discounting case, restrictions on early positions' deferral rights improve welfare due to both waste reduction and our opposed interests result (Theorem 5). In the risk-averse case there is a tradeoff between the benefits of waste reduction and a negative welfare effect highlighted by our aligned interests result (Theorem 3). Nevertheless we argue that even in this case, "marginal" restrictions on deferrals typically enhance welfare at later positions in the list.

1.2 Related literature

Our environment is related to others studied in operations and economics. A special case of our model is the *parallel processor problem*, where jobs of unknown sizes must be completed using a set of processors of different speeds. Agrawala et al. (1984) derive the utilization policy minimizing the expected sum of time to completion across all jobs, which is a special case of our Equation 2. Interpreting total time of completion as utilitarian welfare, their result is a corollary of our Theorem 2.[2] In a model more similar to ours, Su and Zenios (2004) also obtain this kind of utilitarian result where wait-listed agents (with unrestricted deferral rights) are offered randomly arriving objects of varying quality.

When accounting for the welfare of agents *yet to arrive* to a waiting list, it is known that equilibrium behavior in waiting lists need not maximize utilitarian welfare. Naor (1969) points out that an agent deciding to join a queue fails to internalize the additional waiting cost he imposes on future arrivals. The same effect occurs in our setting when agents decide to defer objects. Hassin (1985) succinctly solves this problem by switching the to *last*-come-first-served (LCFS) allocation, forcing agents to internalize the probabilistic arrival of future agents when making joining decisions. Su and Zenios (2004) estimate hypothetical welfare gains of using this approach in kidney allocation, but point out (as does Hassin) that LCFS methods can be impractical due to manipulability, inequity, inducing excessive risk, and being politically unacceptable. We do not consider this approach here.

Related work on dynamic matching and assignment has focused on heterogeneous preferences and discrete time environments. Leshno (2019) improves match quality in such a setting via a "buffer policy," encouraging deferrals with probabilistically improved waiting list positions. Where agents' priority orderings vary over object types, Thakral (2016) proposes a multiple-list procedure: An arriving object "proposes" to its highest priority agent, requiring the agent to accept the object or commit to a separate waiting list dedicated to another object type. Combining this with an idea in Abdulkadiroğlu and Sönmez (1999) yields a strategy-proof, priority-respecting mechanism with a desirable efficiency property. Arnosti and Shi (2020) compare multiple waiting list procedures along with the use of lotteries for horizontally differentiated objects.

[2] See also Kumar and Walrand (1985). Coffman et al. (1987) derive the optimal "Rawlsian" policy of minimizing *makespan* (time elapsed until all jobs are complete).

Bloch and Cantala (2017) consider both horizontal (nonpersistent) preferences and vertical preferences. In a constant-sized waiting list (a new agent replaces any departing one), they consider randomizing the order in which agents are sequentially offered objects (with unrestricted deferral rights). One of their results is that agents prefer a fixed priority order mechanism over any randomized mechanism in their class. This somewhat complements our (logically distinct) Theorem 2 stating that, in our model with a fixed priority order, unrestricted deferral rights yield Pareto dominant outcomes.

Distantly related work considers dynamic two-sided matching of randomly arriving agents. In the kidney exchange model of Ünver (2010), agents *and* objects arrive in pairs leading to additional considerations of property rights. Doval (2018) considers stability in two-sided matching when agents who arrive in different time periods may postpone their arrivals. Akbarpour et al. (2019) analyze the limit behavior of a two-sided market in which agents randomly arrive and depart, comparing mechanisms that differ in their degree of buffering. Baccara et al. (2016) also consider buffering randomly arriving agents before creating a matching in a setting where assortative matchings are utilitarian-efficient. Multi-period matching models are studied by Damiano and Lam (2005), Kurino (2019), Pereyra (2013), Kennes et al. (2014), and Kadam and Kotowski (2018).

2 Model

Agents are waiting to consume an object, where objects arrive randomly over time and waiting is costly. The agents are ordered, which is interpreted as their relative positions in a waiting list. Since our results pertain to the welfare of identical agents who occupy or join some specific position of the list, we specify notation for the set of potential *positions* rather than for individual agents. Denote the set of positions by $\mathbb{N} = \{1, 2, \ldots\}$.

There is a finite set of object types, $\mathcal{O} = \{1, 2, \ldots, n\}$, interpreted as quality levels. Objects of type $i \in \mathcal{O}$ arrive according to a Poisson process with arrival rate μ_i, i.e. the time between consecutive arrivals is exponentially distributed with mean $1/\mu_i$. Furthermore these arrivals are independent of the arrival times of objects of any other types.[3] It is convenient to denote

[3]Equivalently, objects are Poisson arrivals with rate $\sum \mu_i$, and their types are independently determined with probabilities proportional to the μ_i's.

the arrival rate for any set of types $\hat{\mathcal{O}} \subseteq \mathcal{O}$ as

$$\mu_{\hat{\mathcal{O}}} \equiv \sum_{i \in \hat{\mathcal{O}}} \mu_i.$$

An object of type $i \in \mathcal{O}$ has value $v_i \in \mathbb{R}$, where $v_1 > v_2 > \cdots > v_n > 0$.[4] Agents in the list incur a constant flow of waiting costs normalized to one unit of cost per unit of time. Therefore a risk-neutral, non-discounting agent (Section 3) who receives a type i object after waiting t units of time receives a total payoff of $v_i - t$. We consider risk-aversion and discounting of such payoffs in Section 4 and Section 5.

2.1 Waiting list policies

We specify how objects are allocated to agents by motivating and then introducing the concept of a *waiting list policy*, which is the object of our analysis. This concept captures settings in which objects are allocated via first-come first-served (FCFS) waiting lists with deferral rights, but the planner has some method by which deferral decisions can be systematically influenced or constrained as a function of an agent's position in the list. To give a few examples of such influence, imagine a planner who is deciding amongst the following approaches.

1. Allow agents to make selfish, unconstrained deferral decisions ("no influence").

2. Discourage deferrals in some way so that agents early in the waiting list make marginally fewer deferrals than they otherwise would.

3. Offer certain object types directly to agents beginning in some position k of the list (i.e. require certain deferrals at positions $\ell < k$).

Each of these is an example of a waiting list policy as we define below. Our full class of policies allows for generalizations of these alternatives, arbitrary combinations of them, etc. Our objective is to make welfare comparisons across all such policies.

To further illustrate, consider the first alternative above. For any (highest-priority) agent occupying position 1 of the waiting list, optimal

[4]The assumption of strict inequalities simplifies various statements and proofs. The common value assumption avoids the sorting effect in Leshno (2019); see Section 1.

selfish behavior is to defer objects that are not of sufficiently high value, i.e. to accept only object types that belong to some set, say $W^*(1) \subseteq \mathcal{O}$, whose values exceed some threshold.[5] Any agent in position 2 can foresee this behavior (both by the preceding agent and by himself were he to reach position 1). Due to the expected waiting time needed to reach position 1, however, optimal behavior for a position-2 agent is to accept some set of object types $W^*(2) \supseteq W^*(1)$ by using a (weakly) lower quality threshold than in position 1. Similarly, optimal behavior for agents in subsequent positions is described by some collection of sets $W^*(3) \subseteq W^*(4) \subseteq \cdots$ that determine the set of accepted object types at those respective positions.

This function $W^*(k)$ is one example of a waiting list policy as defined below. It describes outcomes under a particular set of assumptions: selfish, informed agents make ("equilibrium") decisions in a kind of waiting list game where they fully observe their positions and the object types being offered. Under those assumptions W^* is the probabilistic social choice function that describes realized consumption (waiting time and object type) for any realized sequence of object types and arrival times. Our main definition generalizes this concept to capture arbitrary deferral decisions rather than just selfish ones.

For instance, suppose the planner can influence the deferral decisions of agents who sit in position 1 of the waiting list, so that they accept any object whose type belongs to some set, $W(1) \subseteq \mathcal{O}$, and defer any other. If $W(1) \supset W^*(1)$ this represents some *restriction* in deferrals (e.g. the second of the three alternatives described above). Likewise $W(1) \subset W^*(1)$ represents an *increase* in deferrals; agents in position 1 are influenced to defer certain object types that they would otherwise accept. In full generality $W(1)$ may satisfy neither inclusion relation. An arbitrary set $W(1) \neq \emptyset$ represents some systematic way in which a planner has committed to influencing the deferral decisions of whichever agent currently occupies position 1 of the waiting list.

The idea extends to positions $k \geq 2$, though it remains convenient to express these sets cumulatively. A set $W(2) \supseteq W(1)$ represents types that are cumulatively accepted by agents within the first two positions, whereas $\mathcal{O} \setminus W(2)$ are deferred. Given $W(1)$, types in $W(2) \setminus W(1)$ are the ones ultimately accepted by agents in position 2. As above, one can imagine that $W(2)$ is determined by the optimal decisions of agents in position 2

[5] This set is easily derived from the model's primitives, as in Subsection 3.2; see also Su and Zenios (2004).

(given $W(1)$ as a constraint); more generally $W(2) \supseteq W(1)$ can be arbitrary, determined by the planner's influence over position 2's deferral decisions. A similar interpretation applies to $W(k) \supseteq W(k-1)$ for each position k.

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

The non-emptiness assumption is innocuous since if $W(1) = \emptyset$, our payoff expressions would apply after relabeling the positions starting with the first non-empty $W(k)$. The restriction to deterministic policies also is not crucial. Payoffs under *randomized* policies can be expressed by re-specifying the set of object types.[6]

2.1.1 Examples

Two natural classes of waiting list policies further illustrate the concept and play a role in later discussion. First, when deferrals are not permitted an agent in position 1 accepts whatever object arrives next. Such a scenario is described by a policy where $W(k) \equiv \mathcal{O}$. A more general class of policies, playing a technical role below, represents a planner who first discards some (possibly empty) set of object types, and disallows deferrals of the remaining types. Under a *no-deferrals* policy, object types $\hat{\mathcal{O}}$ go to the agent occupying position 1 and the rest are discarded.

Definition 2. Policy W is a **no-deferrals policy** if for some $\hat{\mathcal{O}} \subseteq \mathcal{O}$ we have $W(k) \equiv \hat{\mathcal{O}}$.

Second, when agents are unconstrained in their (selfish) deferral decisions, we described the resulting outcomes via a policy W^* discussed in Subsection 2.1. Each set $W^*(k)$ contains some set of object types i whose values v_i exceed some "sufficiently good" threshold from the perspective of position k. More generally imagine that agents defer objects based on such a threshold but that the planner can influence (only) this threshold itself. In this case the planner's set of implementable policies is contained in the following class.

Definition 3. Policy W is a **threshold policy** when, for any position k, $i \in W(k)$, and $j \in \mathcal{O}$, if $v_j > v_i$ (i.e. $j < i$) then $j \in W(k)$.

[6]E.g., a policy where position 1 agents defer type i objects with probability 50% can be considered by splitting $i \in \mathcal{O}$ into types $\{i', i''\}$, each with values v_i and arrival rates $\mu_i/2$. Resplitting allows position 2 to randomize, etc., and our payoff expressions apply.

Many of our results have their most natural interpretation when a planner has "marginal" influence over threshold policies. That is, we imagine a planner who can encourage or discourage deferrals when agents are making "close" decisions. Nevertheless our results apply to arbitrary policies beyond these two classes; in fact the organ spoilage problem in Section 7 is modeled as a constraint on policies that fits neither of these classes.

2.1.2 Comments

Since we model *influence* through the abstraction of Definition 1, a few clarifying comments are in order. First, we do not explicitly model *how* the planner influences deferral decisions, instead leaving this as a black box. We do this in the interest of generality, since any feasibility or implementability constraints (including ethical ones) on the set of achievable policies can be imposed after the fact.

Despite our generality in this sense, our approach also has limitations. One is that we allow a planner to influence agents' decisions only *anonymously*, i.e. solely as a function of an agent's position in the queue, and not the agent's identity or history. In one sense this good: we rule out scenarios in which a planner gives different treatment to agents with identical priorities. On the other hand we set aside the planner's ability to punish or reward agents as a function of their previous deferral decisions.[7] Our policies also ignore other "state of the world" information such as the current queue length, thus ruling out *last*-come-first-served methods. Such methods can overcome certain inefficiencies but suffer from practical difficulties as discussed in Subsection 1.2.

3 Risk-neutral, patient agents

3.1 Equivalence results

We start by showing that the expected payoff starting from some position k of the waiting list is invariant to certain restrictions on deferrals at positions $\ell < k$. Specifically, this expected payoff under a policy W depends only

[7] In this vein, Chen et al. (2018) and Arnosti and Shi (2020) consider waiting lists in which agents lose their positions after making a certain number of deferrals. Abstractly, Leshno (2019) and Thakral (2016) utilize histories by placing deferring agents into secondary waiting lists.

on the set $W(k)$. This result implies that a "no-influence" policy Pareto dominates all other policies from the perspective of non-discounting, risk-neutral agents already present in the list. Furthermore the set of achievable payoffs at position k can be deduced by considering only the simple class of no-deferrals policies.

To establish intuition, contrast two policies W and W' satisfying

$$\begin{aligned} &W(1) \subsetneq W(2) = \mathcal{O} && \text{(position 1 defers some objects)} \\ &W'(1) = W'(2) = \mathcal{O} && \text{(position 1 defers no objects)} \end{aligned}$$

Under W, object types $\mathcal{O} \setminus W(1)$ are deferred by agents in position 1 and accepted at position 2. Under W', all object types are accepted at position 1. Under either policy there is positive probability that an agent in position 2 ultimately receives some object type in the set $W(1)$. Our first observation is that this probability is the same under both policies.

Under W this event happens when the first object to arrive has a type in $W(1)$. This arrival moves the agent in position 2 to position 1, guaranteeing she eventually receives some type in $W(1)$. Under W' this agent is guaranteed to move into position 1, hence the event happens when the second object to arrive has a type in $W(1)$. Thus in both cases the probability she consumes a type in $W(1)$ is $\mu_{W(1)}/\mu_{\mathcal{O}}$.

This idea extends to individual object types and to later positions. An agent in position k consumes type $i \in W(k) \setminus W(k-1)$ only when the next arrival from the set $W(k)$ is i, which has probability $\mu_i/\mu_{W(k)}$. Since this is true for all types and positions, the probabilities of consuming different object types remain proportional to their arrival rates. Hence if $W(k) = W'(k)$, policies W and W' induce equivalent distributions over object types to an agent starting in position k. Omitted proofs are in the Appendix.

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

A similar reasoning applies to expected waiting times. Under W' in the above example, an agent starting in position 2 consumes the second object to arrive regardless of type, for an expected waiting time of $2/\mu_{\mathcal{O}}$. Under W, the arrival of the first object (expected wait of $1/\mu_{\mathcal{O}}$) leads to two cases. If its type is in $W(1)$ (probability $\mu_{W(1)}/\mu_{\mathcal{O}}$), the agent waits for a second arrival from $W(1)$, adding $1/\mu_{W(1)}$ of expected waiting time. Otherwise the agent departs immediately. Total expected waiting time under W is thus $2/\mu_{\mathcal{O}}$.

Proposition 2 generalizes this observation and more broadly can be seen as a consequence of Little's Law (Little (1961)).[8] This fundamental queueing result states that the average time agents spend in a "system" (the first k positions of the waiting list) is the average number of agents in that system divided by the rate at which they are served (by receiving an object).

Henceforth we let random variable $\boldsymbol{t_k^W}$ denote the waiting time for an agent starting from position k under policy W.

Proposition 2 (Expected-waiting-time equivalence)**.** *For any policy W, an agent in position $k \in \mathbb{N}$ has an expected waiting time of $E(t_k^W) = k/\mu_{W(k)}$.*

Combining these propositions, an agent's expected payoff starting from position k (henceforth denoted $\boldsymbol{\Pi(k;W)}$) depends only on the set $W(k)$.

Theorem 1 (Expected-payoff equivalence)**.** *For any policy W, the expected payoff to an agent starting in position $k \in \mathbb{N}$ is*

$$\Pi(k;W) = \frac{\sum_{i \in W(k)} \mu_i v_i - k}{\mu_{W(k)}} \tag{1}$$

As a consequence, the expected payoffs achievable from some position k can be determined by considering only the class of no-deferrals policies.

Corollary 1. *For any policy W, the expected payoff to an agent in position k is the same as the expected payoff under the no-deferrals policy W' defined as $W'(\ell) = W(k)$ for all $\ell \in \mathbb{N}$.*

3.2 Dominance of the no-influence policy

The main implication of Theorem 1 is that, once we fix a set of types $W(k)$ to be consumed by agents through position k, position k's payoff is unaffected by varying the sets $W(\ell) \subseteq W(k)$ for earlier positions $\ell < k$. Consequently we can show that, in the baseline case of risk-neutral, non-discounting agents, a policy which allows agents to make uninfluenced deferral decisions in their own self interest is a Pareto dominant policy for agents already present in the list, in that it simultaneously maximizes $\Pi(k;\cdot)$ for each k.[9]

Specifically, the "no-influence" policy defined below is not only dominant in this sense but also corresponds to equilibrium behavior in a naturally

[8] Little's Law is applied in a related way by Bloch and Cantala (2017).

[9] We account for the payoffs of future agent arrivals in Section 6.

defined game where agents in our model make rational, uninfluenced acceptance/deferral decisions. Formally we construct a policy that maximizes a single position's payoff (Lemma 1), then observe that this construction can be accomplished simultaneously across all positions yielding the dominance result (Theorem 2). We conclude with a straightforward argument relating this construction to selfish, unrestricted deferral decisions.

To construct a policy that maximizes $\Pi(k, \cdot)$ for some k, it is sufficient to specify a set $W_0^*(k)$ that maximizes the right-hand side of Equation 1. Not surprisingly, such a set is a "threshold set" of the form $W_0^*(k) = \{1, 2, \ldots, i\}$, where i is the highest index (i.e. lowest quality type) for which v_i exceeds the expected payoff with respect to $W_0^*(k)$.

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

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

Furthermore $i^*(k)$ is a weakly increasing function of k.

The monotonicity of the thresholds $i^*(k)$ reflects the fact that agents in later positions are willing to accept lower valued objects in order to avoid greater expected wait times. It also means that it is feasible to *simultaneously* provide each position k with its favorite policy.

Definition 4. The **no-influence policy**, W_0^*, is the threshold policy where, for each k, $W_0^*(k) = \{1, 2, \ldots, i^*(k)\}$ as defined in Lemma 1.

In a related model, Su and Zenios (2004) show that such a policy is utilitarian-optimal, maximizing total expected payoffs to agents currently in the list.[10] Lemma 1 leads to the stronger conclusion that W_0^* dominates any other policy under any Pareto-consistent welfare objective applied to agents currently in the list.

Theorem 2 (Dominance)**.** *The no-influence policy maximizes each position's payoff: For any policy W and position k, $\Pi(k; W_0^*) \geq \Pi(k; W)$.*

Though W_0^* maximizes welfare for agents already present in the list, it is known that the expected welfare of agents yet to arrive to the list can be improved by constraining deferrals, as we address in Section 6.

[10]Agrawala et al. (1984) prove an analogous result in a special case of our model.

Finally we observe that W_0^* represents equilibrium behavior in a "waiting list game" where the agents in our model make rational, unconstrained deferral decisions.[11] As is clear from Equation 2, $W_0^*(1)$ describes an optimal stopping policy for any agent in position 1: the agent should accept any object whose value would exceed the continuation value of waiting for a better one. An agent in position 2 also should accept any object whose value would exceed the continuation value of waiting for a better one. Clearly such objects form a threshold set. Furthermore the threshold set $W_0^*(2)$ yields the maximum feasible payoff to position 2 among all threshold sets (Theorem 2). Thus, even taking $W_0^*(1) \subseteq W_0^*(2)$ as given, an agent in position 2 obtains the highest feasible payoff by accepting only object types in $W_0^*(2)$; this set describes optimal (equilibrium) behavior at position 2. The same argument applies iteratively to all later positions.

3.3 Waiting time distributions

By Proposition 2, the expected waiting time $E(t_k^W)$ depends only on $W(k)$. Nevertheless the distribution of t_k^W depends on the sets $W(\ell)$ for $\ell < k$. The distribution of waiting times can be relevant for reasons outside the scope of our model. In organ waiting lists, for example, waiting time variability can hinder doctors' choices of treatment (Bandi et al. (2018)). Third parties may desire lower variability in waiting costs (e.g. insurance companies reimbursing dialysis costs).

Though the distribution of t_k^W is difficult to describe in general, in this section we derive its variance for an arbitrary policy W. We then show that certain restrictions on deferrals reduce variance in waiting times while keeping *average* waiting times constant (Proposition 2). First we observe that the distribution of t_k^W can be described in the special case of no-deferrals policies.

Lemma 2 (Waiting time distributions for no-deferrals)**.** *Consider a no-deferrals policy W, where $W(k) = \hat{O} \subseteq \mathcal{O}$ 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. An agent in position k departs after k arrivals from $\hat{O}$, the wait for each arrival being exponentially distributed with parameter $\mu_{\hat{O}}$. The sum

[11] We omit formalities since a related idea appears in Su and Zenios (2004); see also Agrawala et al. (1984).

of k (i.i.d.) exponential distributions yields an Erlang distribution, and the mean and variance calculations follow directly. □

The derivation of $\mathrm{Var}(t_k^W)$ is easily explained in the case of two object types $\mathcal{O} = \{1, 2\}$ for a policy with $W(1) = \{1\}$ and $W(2) = \{1, 2\}$. Decompose position 2's waiting time into $t_2^W = t' + t''$ as follows. First an agent in position 2 waits t' units of time for an arrival of a type from $W(2) = \{1, 2\}$, so $t' \sim \mathrm{Exp}(\mu_1 + \mu_2)$. Conditional on that object being type $1 \in W(1)$, the agent waits $t'' \sim \mathrm{Exp}(\mu_1)$ additional units of time for a *second* arrival of a type 1 object. Conditional on it being type $2 \in W(2) \setminus W(1)$, the agent departs immediately (waiting $t'' = 0$ additional units of time). The (unconditional) variance of t'' thus can be shown to be

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

Since t' and t'' are independent we have

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

Note that this variance exceeds that of a policy W' satisfying $W'(1) = W'(2) = \{1, 2\}$, which is $\mathrm{Var}(t_2^{W'}) = 2/(\mu_1 + \mu_2)^2$ by Lemma 2. Even though W and W' yield the same expected waiting time from position 2, the one exhibiting deferrals (W) has higher variance. This idea extends to the general case. Proposition 3 is proven recursively in k, replacing the above object types $\{1, 2\}$ with sets $W(k-1)$ and $W(k) \setminus W(k-1)$. There, t' is the wait for an arrival from $W(k)$ and t'' is the entire (continuation) waiting time.[12]

Proposition 3 (Waiting time variance)**.** *For any policy W, the waiting time from position k has a variance of*

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

[12]In the special case $k = 2$, t'' has a hyper-exponential distribution with known expression for variance. Since this does not extend to $k > 2$, we use our recursive approach.

The expression shows that waiting time variance decreases as deferral rights are restricted. Precisely, $\mathrm{Var}(t_k^W)$ decreases if we expand any set $W(\ell)$ with $\ell < k$. In words, if agents in position $\ell < k$ no longer defer an object type $i \in W(k)$ that otherwise would be accepted in some position $\ell' \leq k$, then $\mathrm{Var}(t_k^W)$ necessarily decreases. This establishes a tradeoff between the dominance of "no-influence" (Theorem 2) and any concern a planner or third party may have for waiting time predictability.

Corollary 2 (Restricting earlier deferrals reduces waiting time variance)**.** *Fix position $k \in \mathbb{N}$ and policies W and W' satisfying (i) $W'(k) = W(k)$ and (ii) for all $\ell < k$, $W(\ell) \subseteq W'(\ell)$. Then $Var(t_k^{W'}) \leq Var(t_k^W)$. Hence subject to any constraint that $W(k) = \hat{\mathcal{O}}$ for some $\hat{\mathcal{O}} \subseteq \mathcal{O}$, $\mathrm{Var}(t_k^W)$ is minimized by the no-deferrals policy $W(\ell) \equiv \hat{\mathcal{O}}$.*

Finally we note that, at position k, expanding the set $W(k)$ itself has ambiguous effects on $\mathrm{Var}(t_k^W)$. Intuitively, adding a type $i \notin W(k)$ to $W(k)$ lowers variance through an increased object arrival rate; but could increase variance by increasing the chance that an agent in position k makes a quick exit from the list. Simple examples show that either effect can dominate.

4 Risk-averse agents

Deferral rights impact the welfare of risk-averse agents by altering both the distribution of waiting time (Subsection 3.3) and its correlation with object consumption. Under the assumption of constant absolute risk-aversion we extend the idea behind Theorem 2: no-influence policies simultaneously maximize each position's expected utility. More strongly, expected utility maximization for any single position k *requires* the use of a no-influence policy at all earlier positions $\ell < k$. This follows from an "aligned interests" result that is interesting in its own right: agents in any two consecutive positions must agree on whether any "marginal" policy change is desirable.

Intuitively, risk-averse agents in later positions of the list might seem to benefit from deferral restrictions at earlier positions since waiting time variability decreases (Corollary 2) while its expectation remains constant (Proposition 2). However this change also alters the relationship between an agent's two payoff components, waiting time and object value. Typically, payoff variability is reduced by increasing the correlation between these two (cost and benefit) components. Fixing a set $W(k)$, a risk-averse agent in position k might prefer that $W(1) \subseteq W(k)$ contain only the highest-value types in $W(k)$ rather than only the lowest-value types, since the former increases this correlation. While no-influence policies obviously correlate higher value types with longer waiting times to some degree, we show there is a sense in which they do so Pareto-optimally.

Formally we consider agents with constant absolute risk-aversion: an agent's utility from a payoff of $v - t$ is $u(v-t) = -e^{-\alpha(v-t)}$, fixing the risk parameter $\alpha > 0$. We assume that $\alpha < \mu_i$ for each $i \in \mathcal{O}$, which is innocuous under realistic assumptions on parameters.[13]

We use the following two facts. First, if a random variable t (e.g. waiting time) is exponentially distributed with parameter $\mu > \alpha$, 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, then

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

For the remainder of this section we let $\boldsymbol{U_k^W}$ denote the **expected utility** for an agent starting in position k under policy W (suppressing the dependence on α). It is simple to derive U_1^W since an agent in position 1 receives some object type from $W(1)$ after waiting an exponentially distributed amount of time. From the two facts above,

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

It is more tedious to describe U_k^W for $k > 1$ since the distribution of waiting time is tied to the possible sequences of object type arrivals. Nevertheless Equations (4) and (5) lead to a recursion relation (Equation 6) that in turn yields Equation 7.

[13]This assumption bounds expected (dis)utility. Our results hold whenever $\mu_{W(1)} > \alpha$.

Proposition 4. *Suppose agents have CARA utility with parameter α. For any policy W and position $k \in \mathbb{N}$ we have*

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

The expected utility under W starting from position k is

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

Proposition 4 allows us to see how a "marginal" influence in deferral decisions impacts U_k^W as follows. Fix policy W and position k, and consider an object type $j \in W(k) \setminus W(k-1)$, i.e. j is deferred by agents in positions 1 through $k-1$ and accepted by agents in position k. Holding everything else constant, imagine that we change the policy so that agents in position $k-1$ now accept type j objects rather than defer them. How does this change affect the expected utility of an agent in position k? Equation 7 indicates two opposite effects.

- Since j now belongs to $W(k-1)$, the summation term corresponding to $j \in W(k)$ is now multiplied by a product that contains an additional "penalty term" $\mu_{W(k-1)}/(\mu_{W(k-1)} - \alpha) > 1$, decreasing U_k^W. The interpretation is that *an agent in position k is worse off due to a relatively longer expected wait for an object of type j.*

- Adding j to $W(k-1)$ decreases this "penalty term" $\mu_{W(k-1)}/(\mu_{W(k-1)} - \alpha)$, which is applied to each of the summation terms for each other type $i \in W(k-1) \setminus \{j\}$. The interpretation is that *an agent in position k is better off due to a relatively shorter expected wait for each such type i.*

Which effect dominates depends on the relative magnitude of v_j. When v_j is sufficiently high, the second (positive) effect outweighs the first; an agent in position k benefits from the restriction placed on position $k-1$'s deferral rights. This captures the idea discussed above that *risk-averse agents want to correlate high object values with long waiting times.* Since placing type j into the set $W(k-1)$ changes neither the distribution of objects received (Proposition 1) nor the average waiting time from position k (Proposition 2), this effect entirely drives the conclusion. When v_j is sufficiently low the reverse arguments apply.

It is more straightforward to see how this policy change affects the expected utility of an agent in position $k-1$. Requiring such an agent now to accept type j objects is beneficial only when v_j is sufficiently high. That is, we reach the same qualitative conclusion as we did for position k; their interests are somewhat aligned.

It turns out that this "alignment of interests" is sharp. If we change a policy W by allocating type j objects to position $k-1$ rather than position k, this change benefits position k if and only if it harms position $k-1$.

Theorem 3 (Aligned interests)**.** *Suppose agents have CARA utility with parameter α. Fix a policy W, a position $k \geq 2$, and an object type $j \in W(k) \setminus W(k-1)$. Let W' be obtained from W by allocating j to position $k-1$ instead of to k, that is $W'(k-1) = W(k-1) \cup \{j\}$, and $W'(\ell) = W(\ell)$ for all $\ell \neq k-1$, Then $U_k^W \geq U_k^{W'}$ if and only if $U_{k-1}^W \geq U_{k-1}^{W'}$.*

Two agents occupying positions $k-1$ and k would agree on whether type j objects should be accepted by whomever occupies position k or occupies position $k-1$. This leads us to generalize Theorem 2 to the case of CARA utility. Namely, a policy in which agents make unrestricted, selfish deferral decisions simultaneously maximizes each position's expected utility.

Definition 5. Given $\alpha > 0$, define the **no-influence policy**, W_α^*, as follows. For $k = 1$, let $W_\alpha^*(1)$ be the set of types that maximizes the RHS of Equation 7.[14] For $k = 2$, subject to the constraint that $W(1) = W_\alpha^*(1)$, choose $W_\alpha^*(2)$ to maximize the RHS of Equation 7. Continuing for $k > 2$, taking sets $W_\alpha^*(1), \ldots, W_\alpha^*(k-1)$ as fixed, let $W_\alpha^*(k)$ maximize Equation 7.

Theorem 4 (Dominance)**.** *Suppose agents have CARA utility functions with parameter α. The no-influence policy W_α^* maximizes each position's expected utility: For any policy W and position k, $U_k^{W_\alpha^*} \geq U_k^W$.*

5 Discounted payoffs

Discounting reverses the conclusions obtained under risk-aversion. Mirroring Theorem 3, our *opposed interests* result states that agents in the two consecutive positions impacted by a "marginal" policy change necessarily disagree on the desirability of the change. Since what is good for the earlier agent is

[14] Ties can be broken arbitrarily.

bad for the later one, there is no dominance result in the discounting case analogous to Theorem 4. More significantly, the welfare of (many) agents starting in later positions of the list can be improved by influencing deferral decisions made in (few) earlier positions.

The underlying idea is that, unlike in the risk-averse case, discounting agents prefer to correlate *short* waiting times with high object values. More generally, discounting agents can gain by receiving offers more quickly (particularly high-value ones) even at the cost of delaying other potential offers. As an analogy, compare (i) receiving two gifts tomorrow to (ii) receiving one gift today and another in two days. The discounting agent prefers (ii) *unless today's gift is of sufficiently low relative value*. Similarly, an agent in our setting benefits when an earlier position defers "today's gift" unless its value is too low. So for example, influencing an earlier position to accept marginally lower-valued objects than they otherwise would can benefit agents starting in all later positions. We show that this benefit accrues precisely when the earlier position is *harmed* by this influence.

Formally we consider (risk-neutral) agents who continuously discount the future at rate r: the present value of a payoff x received t units of time in the future is $x \cdot e^{-rt}$. If an agent incurs (unit flow) waiting costs for $t \sim \mathrm{Exp}(\mu)$ units of time, then the expected present value of this total cost is

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

If an agent is to receive an object of value v at $t \sim \mathrm{Exp}(\mu)$ units of time in the future, then the expected present value of this object is

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

We let $\boldsymbol{EPV_k^W}$ denote the **expected present value** (**EPV**) of an agent's payoff starting in position k under policy W (suppressing dependence on r). Under policy W, EPV_1^W is the EPV of the next arrival from $W(1)$ minus the EPV of its associated waiting cost.

$$\begin{aligned} EPV_1^W &= \left(\sum_{i \in W(1)} \frac{\mu_i}{\mu_{W(1)}} \cdot \frac{v_i \mu_{W(1)}}{r+\mu_{W(1)}} \right) - \frac{1}{r+\mu_{W(1)}} \\ &= \frac{1}{r+\mu_{W(1)}} \left[\left(\sum_{i \in W(1)} \mu_i v_i \right) - 1 \right] \end{aligned} \tag{10}$$

The simplest way to think of EPV_2^W is to imagine that an agent in position 2 waits for the next arrival from $i \in W(2)$ and either consumes it (if $i \in W(2) \setminus W(1)$), or otherwise immediately receives a pseudo-object worth $v = EPV_1^W$. That is, analogously to Equation 10 we have

$$EPV_2^W = \frac{1}{r + \mu_{W(2)}} \left[\left(\mu_{W(1)} \cdot EPV_1^W + \sum_{i \in W(2) \setminus W(1)} \mu_i v_i \right) - 1 \right]$$

Repeating the argument yields the following recursive relationship.

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

This yields the following generalization of Equation 1.

Proposition 5 (Discounted payoffs). *Suppose agents discount payoffs at rate r, and fix a policy W. The expected present value of the payoff to an agent starting in position k is*

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

Equation 12 decomposes EPV_k^W into k components, each corresponding to the possible positions $\ell = 1, \ldots, k$ at which the agent ultimately accepts an object. Each ℓth component has both

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

Both parts are multiplied by a more subtle product term, $\prod_{m=\ell}^{k} \mu_{W(m)}/(r + \mu_{W(m)})$, whose interpretation differs across these two parts. Applied to the value part, this product is the amount by which the future object value is discounted to the present (as in 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 (as in Equation 8), while its numerator corresponds to the probability of reaching a position $\ell \leq k$.

We use Equation 12 to prove the main result of this section.

Theorem 5 (Opposed interests). *Suppose agents discount payoffs at rate r. Let policies W, W', position $k \geq 2$ and object type $j \in W(k) \setminus W(k-1)$ be defined as in Theorem 3, i.e. W' is obtained from W by allocating j to position $k-1$ instead of to k. Then*

$$EPV_k^W > EPV_k^{W'} \iff EPV_{k-1}^W < EPV_{k-1}^{W'} \tag{13}$$

The result illustrates a welfare tradeoff between earlier and later positions. A "marginal" restriction in deferrals—in the sense of limiting position $k-1$'s ability to defer a type otherwise accepted at position k—benefits position k precisely when it harms position $k-1$. Due to the recursive nature of payoffs in Equation 11, this benefit to position k extends to *all* positions $\ell \geq k$. In other words, Equation 13 can be strengthened to state

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

Thus while positions k and $k+1$ have "opposed interests" over W vs. W', *all* positions $\ell \geq k+1$ have common preferences over the two policies.

Finally Theorem 5 shows that there is no dominance result in the sense of Theorem 2: any marginal policy change of the kind described in the theorem must benefit exactly one of the two affected positions. On the other hand, for a (e.g. utilitarian) planner wishing to favor a large number of agents joining later positions of the list, the result provides an argument for (at least marginally) suppressing deferral rights in earlier positions. The following two sections further highlight this idea.

6 Agent arrivals and long-run welfare

Theorems 2 and 4 show that, in the benchmark and risk-averse cases, no-influence policies are Pareto-dominant from the perspective of agents already present in the waiting list. However such policies need not maximize the expected welfare of agents yet to join the list, since restrictions on deferral rights can reduce their expected waiting time upon arrival by shortening queue lengths. From the ex ante perspective of agents joining a waiting list, we consider the tradeoff between these two ideas in the non-discounting cases. Most notably we show that the desirability of "no-influence" is strictly preserved across *earlier* positions of the list, even when arriving agents' welfare can be improved by suppressing deferrals at the *latest* positions of the list.

To do this concisely we focus on the following objective: *maximize the (long run) expected welfare of agents who randomly join a waiting list over time*. Consider extending our model so that agents randomly arrive (at Poisson rate $\lambda > 0$) and expire (at rate $\gamma > 0$, departing with a "null object" with value zero). The resulting length of a waiting list under policy W would then be a random variable K whose (steady state) distribution depends on W. The expected payoff of joining the list would be obtained by combining this distribution with our earlier payoff expressions, e.g. Equation 7. Unfortunately this distribution is easily specified only in special cases.[15]

At the same time, the uncertainty in queue length K is not what drives our main point. Furthermore we can approximate this length via the (fluid) limit of our model: when agents and (sufficiently scarce) objects arrive continuously, the length of the list under policy W is the constant $K > 0$ at which the total arrival rate (λ) equals the total departure rate ($\mu_{W(K)}$ depart with objects, γK expire). That is, $\lambda = \mu_{W(K)} + \gamma K$. This motivates the following natural approximation to our objective above: *choose W to maximize the welfare of agents who join the list at "position" $K = (\lambda - \mu_{W(K)})/\gamma$.*

In the risk-neutral, non-discounting case this is formalized as

$$\max_{W} \Pi(K; W) \quad \text{where } K = (\lambda - \mu_{W(K)})/\gamma \tag{15}$$

That is, choose W to maximize the payoff of joining "position" K, where K is the (fluid limit) queue length endogenously determined by W. This exercise is easily solved by substituting for K in Equation 1.[16]

Proposition 6. *Equation 15 is solved by the* ***no-influence until the end policy,*** $\tilde{\boldsymbol{W}}$*, constructed as follows.*

1. *Let $\hat{\mathcal{O}} = W_0^*(\lambda/\gamma)$ be the set of types that would be accepted through position λ/γ under the no-influence policy W_0^*.*
2. *Determine position $K \equiv (\lambda - \mu_{\hat{\mathcal{O}}})/\gamma$.*
3. *Set $\tilde{W}(K) = \hat{\mathcal{O}}$, i.e. deferrals of types $i \in \hat{\mathcal{O}}$ are disallowed at position K.*
4. *For all $\ell < K$ set $\tilde{W}(\ell) = W_0^*(\ell)$, i.e. use no-influence in "early" positions.*

---

[15] E.g. when W is a no-deferrals policy and $\gamma = 0$, K has a geometric distribution. Su and Zenios (2004) consider some other cases via Bellman equations.

[16] Equation 1 is robust to K being non-integer.

As expected, $\tilde{W}$ restricts deferrals at position K ($W_0^*(K) \subseteq \tilde{W}(K)$ since $K < \lambda/\gamma$) so that arriving agents enjoy a shorter average wait, but for objects of lower average quality. Steps 1–3 of the construction of $\tilde{W}$ are sufficient to solve Equation 15. The "no-influence" specification of Step 4 is not necessary under risk-neutrality (by Theorem 1) but guarantees that interim expected payoffs at positions $\ell < K$ also are maximized.

Under risk-aversion, however, this kind of no-influence specification is required to maximize the expected-utility analog of Equation 15. Proposition 7 states that, even if we restrict deferrals at position k ($W(k) \supseteq W_\alpha^*(k)$), expected utility at that position is maximized by allowing uninfluenced deferral decisions at all earlier positions $\ell < k$. The proof is similar to that of Theorem 4 and is available upon request.

Proposition 7. *Suppose agents have CARA utility functions with $\alpha > 0$. For any position k and set $\hat{\mathcal{O}} \supseteq W_\alpha^*(k)$, if policy W' solves*

$$\max_W U_k^W \quad s.t.\ W(k) = \hat{\mathcal{O}}$$

then for all $\ell < k$ we have $U_\ell^{W'} = U_\ell^{W_\alpha^*}$. That is, up to indifference, W' applies "no-influence" to positions $\ell < k$.

Hence the solution to the risk-averse analog of Equation 15 requires a "no-influence until the end" policy, $\tilde{W}_\alpha$, analogous to the one in Proposition 6. It (i) restricts deferrals at some position $K = (\lambda - \mu_{\tilde{W}_\alpha(K)})/\gamma$ so that $\tilde{W}_\alpha \supseteq W_\alpha^*(K)$, and (ii) applies the no-influence policy of Definition 5 (up to indifference) at positions $\ell < K$. In summary, restrictions in deferral rights can improve the welfare of arriving, risk-averse agents by reducing the length of their wait, but ideally these restrictions should take place only at relatively later positions of the waiting list.

In contrast we observe that when payoffs are discounted, arriving agents benefit when these deferral restrictions occur *earlier* in the list. Requiring earlier positions to accept (marginally) lower quality objects than they otherwise would can simultaneously shorten the length of the list while increasing (later positions') discounted payoffs via Theorem 5. Since this idea resurfaces in the following section so we omit further formalities in the discounting case.

7 Application: Organ spoilage

We consider some practical implications of our results within a stylized model of "organ spoilage." Transplant organ waiting lists prioritize patients using various characteristics (e.g. health status, geographic location, join date, physical characteristics). The prioritized agents sequentially accept or defer arriving organs with the advice of health care providers.[17] However the organs' limited shelf life can lead to waste: Lower quality organs (acceptable to low-priority patients) can spoil in the time it takes to process their offer to (and deferral by) high-priority patients (Sack (2012)). In this sense, unrestricted deferral rights create inefficiency.

Abstracting away from the finer details of such environments, our results shed light on the tradeoff between this inefficiency and the potential welfare benefits of deferrals. Start from a scenario with unrestricted deferral rights (a "no-influence" policy), but where some desirable objects spoil due to an excessive number of deferrals. Consider a planner who can influence agents to (marginally) lower their threshold of acceptance, reducing the rate of deferrals (and of spoilage). This affects the welfare of agents in (or joining) later positions of the list in two ways. The positive *resource effect* is the increased consumption rate of desirable objects and the accompanying reduction in average waiting times. The *preference effect*, resulting from a change in the relative timing of various quality offers, depends on the position being considered, the positions being influenced, and the agents' preference characteristics as we have shown.

In the discounting case the preference effect here tends to be positive for agents in later positions. As in our *opposed interests* result, marginal deferral restrictions that harm earlier positions benefit later positions. Since the resource effect is also positive, the kind of influence described above tends to be beneficial to agents in (or joining) later positions.

In the risk-averse case, our *aligned interests* result suggests the opposite: marginal policy changes that harm earlier positions also harm later ones. Whether this outweighs the positive resource effect is ambiguous. We explicitly consider this tradeoff by modeling the above spoilage phenomenon as a simple "spoilage constraint" on policies.

[17] A growing empirical literature studies these decisions; see Arikan et al. (2017), Ata et al. (2017), Agarwal et al. (2019).

7.1 A spoilage constraint

To illustrate the main idea it is sufficient to consider two object types, $\mathcal{O} = \{1, 2\}$, with $v_1 > v_2$ and arbitrarily arrival rates $\mu_1, \mu_2 > 0$. In the context of transplant organs, type 1 represents *high quality* organs that would be rarely (or never) deferred, and hence rarely spoil. Type 2 represents organs of *acceptable quality*, deferred by agents in early positions but acceptable to agents in much later ones. Inefficiency arises when type 2 objects spoil before being offered to those later positions.

We model spoilage by assuming that objects can be offered only to the first κ positions of the waiting list for some fixed κ. To motivate this, imagine it takes a small amount of processing time (ϵ) for an agent to receive and defer an offer. If an object's shelf-life is some value $\kappa\epsilon$, the object spoils after κ deferrals. Since organ shelf-lives are small relative to other parameters,[18] our approximation captures the notion of limited numbers of offers even while ignoring the actual (relatively negligible) transaction times.

Fixing κ, we say that W satisfies the *spoilage constraint* when

$$\forall \ell \in \mathbb{N}, \quad \ell > \kappa \implies W(\ell) = W(\kappa). \tag{16}$$

That is, agents in positions $\ell > \kappa$ must "defer" any type $i \notin W(\kappa)$ by virtue of the fact that the object has spoiled before they can receive an offer.

For spoilage to be a cause of inefficiency, it must be that (i) agents in earlier positions do in fact defer lower-quality objects, and (ii) these objects are considered acceptable by agents in later positions. We formalize these conditions without (yet) specifying preference characteristics; they are considered separately further below. For now, let W^* refer to the "no-influence" policies we have discussed earlier *in the absence of a spoilage constraint*. That is, in the non-discounting cases W^* refers to the policies in Definitions 4 or 5, and in the discounting case W^* refers to the obvious analog[19] of Definition 5.

Under the assumptions of our model we clearly have $1 \in W^*(1)$ (agents in first position accept the best object type). The condition that *spoilage creates inefficiency* is the following assumption that type 2 objects *would be* accepted by agents in sufficiently later positions, had they not spoiled.

Assumption 1 (spoilage creates inefficiency)**.** There exists $\hat{k} > \kappa$ such that $2 \in W^*(\hat{k}) \setminus W^*(\hat{k} - 1)$.

[18] E.g. kidneys last roughly a day whereas patients can spend years on the list.
[19] That is, each $W^*(k)$ maximizes Equation 12 given $W^*(1), \ldots, W^*(k-1)$.

In other words, agents in position $\hat{k}$ (but no earlier) would prefer to accept type 2 objects when offered. However these offers are prevented from happening due to spoilage, i.e. W^* fails Equation 16 because $\hat{k} > \kappa$.

7.2 Restricting deferrals

We consider mitigating the spoilage problem by influencing agents at some (sufficiently early) position j to accept rather than defer type 2 objects. For any $j \in \mathbb{N}$ define policy W^{I_j} as follows.

$$W^{I_j}(k) = \begin{cases} \{1\}, & k \leq j-1 \\ \{1,2\}, & k \geq j. \end{cases}$$

Policy W^{I_j} satisfies the spoilage constraint (16) only when $j \leq \kappa$, but the cases $j > \kappa$ also play a role in some arguments below.

If objects were immune to spoilage, unrestricted deferral decisions would lead to the outcomes described by policy W^*. In the presence of spoilage, however, unrestricted deferral decisions would cause type 2 objects to go to waste under Assumption 1. Let W^S be the policy that describes such outcomes, i.e.

$$\forall k \in \mathbb{N}, \quad W^S(k) = \{1\}.$$

We make welfare comparisons between W^S and W^{I_j}; in other words, would an agent prefer allowing spoilage to occur as a result of unrestricted deferral decisions, or to prevent it by eliminating deferrals at position $j \leq \kappa$? By construction, an agent *already occupying* position $j \leq \kappa$ prefers policy W^S to W^{I_j}. The more important comparison is for the (typically large number of) agents who occupy or join later positions $k \geq \hat{k}$ of the list. We first address the simpler discounting case, which includes the risk-neutral, non-discounting benchmark as a special case.

7.2.1 Risk-neutral agents

Following our earlier intuition, discounting agents in later positions of the list benefit when acceptable objects are no longer deferred at some earlier position j. Here such agents benefit from (i) increased utilization of type 2 objects, (ii) decreased waiting times, and (iii) the opposed interests effect of Theorem 5. We further show that these benefits are decreasing in j: earlier restrictions are better.

Proposition 8 (later positions prefer earlier deferral restrictions)**.** *Under Assumption 1, consider risk-neutral agents who discount with rate $r \geq 0$, and let j, k satisfy $j \leq \kappa < \hat{k} \leq k$. The expected, discounted payoff to position k is strictly greater under W^{I_j} than under W^S. If $r > 0$ then for any $\ell < j$ this payoff is strictly greater under W^{I_ℓ} than under W^{I_j}.*

The argument is simplest in the non-discounting case ($r = 0$). Since $W^{I_j}(k) = W^*(k)$ by construction, an agent in position k is indifferent between those policies (Theorem 1); that is under $W^{I_j}(k)$, the agent receives an expected payoff *as if deferral rights were unrestricted and objects did not spoil*. By Assumption 1 this is strictly preferred to W^S.

Discounting ($r > 0$) amplifies this argument due to Theorem 5. Since agents in an early position $j \leq \kappa$ prefer to defer type 2 objects, agents in later position $k \geq \hat{k}$ benefit from eliminating such deferrals. Repeating the argument as j decreases implies that earlier restrictions are increasingly beneficial.

7.2.2 Risk-averse agents

As above, restricting an early position j's ($k \leq \kappa$) ability to defer type 2 objects impacts risk-averse agents in later positions in multiple ways. They benefit from (i) increased utilization of type 2 objects and (ii) decreased waiting times, but suffer (iii) the aligned interests effect of Theorem 3.

It follows from Proposition 7 that effect (iii) is minimized by setting $j = \kappa$. That is, among all policies W^{I_j} with $1 \leq j \leq \kappa$, W^{I_κ} maximizes the expected utility to positions $k \geq \kappa$: later restrictions are better.[20]

It is possible, however, that an agent in some later position $k \geq \kappa$ would prefer not to reduce spoilage in this way at all. That is, even under W^{I_κ}, the negative effect (iii) may dominate the positive ones. To avoid tedious expressions, we demonstrate this in the simplest case where $\kappa = 1$ (objects spoil after one deferral) and Assumption 1 holds at $\hat{k} = 2$ (position 2 would have accepted a type 2 object had it not spoiled). Equation 17 characterizes the parameters under which an agent in some position $k \geq 2$ would prefer policy W^{I_1} (preventing spoilage by restricting deferrals) to W^S (allowing deferrals, resulting in spoilage). We interpret it below.

[20] By Proposition 7, subject to the constraint $W(\kappa) = \{1, 2\}$, position κ's expected utility is maximized when there is "no-influence" to positions 1 through $(\kappa - 1)$. This conclusion recursively extends to any position $k \geq \kappa$ by Equation 6.

Proposition 9 (later positions may or may not prefer earlier deferral restrictions)**.** *Suppose that $\kappa = 1$, and that Assumption 1 holds at $\hat{k} = 2$. Consider (non-discounting) risk-averse agents with parameter $0 < \alpha < \mu_1$. For any $k \geq 2$ we have $U_k^{WI_1} > U_k^{WS}$ (position k benefits from restricting position 1's deferrals) if and only if*

$$v_1 - v_2 < L_{\alpha,k} \equiv \frac{1}{\alpha} \ln \left(\frac{\mu_1 + \mu_2}{\mu_2} \left[\frac{\mu_1(\mu_1 + \mu_2 - \alpha)}{(\mu_1 + \mu_2)(\mu_1 - \alpha)} \right]^k - \frac{\mu_1}{\mu_2} \right) \tag{17}$$

Furthermore $L_{\alpha,k} > 0$ is increasing (without bound) in k.

The result tells us two things. First, position k benefits from the deferral restriction whenever v_2 is sufficiently close to v_1 (since $L_{\alpha,k} > 0$). The interpretation is that later position in the list must benefit from a "sufficiently marginal" restriction in deferral rights. Second, regardless of object values, position k benefits whenever k is sufficiently large (since $L_{\alpha,k}$ is unbounded). Hence under risk-aversion, deferral restrictions that reduce spoilage of desirable objects must benefit all agents in sufficiently late positions of the list.

On the other hand when inequality (17) fails, a reduction position j's in deferral rights *increases* the utilization of desirable (type 2) objects, decreases waiting times, and yet *reduces* expected welfare at later positions. This finding has implications in the application of organ waiting lists, where policy changes are evaluated in part using organ utilization rates (Sack (2012)). A policy that *increases* organ utilization rates (via restricted deferrals) nevertheless can *decrease* welfare for risk-averse agents (at any k where (17) fails). Of course this finding does not establish the plausibility of such a scenario, which is left as an empirical question.[21]

8 Conclusion

We have considered the welfare implications of arbitrarily constraining deferral decisions in first-come-first-served waiting list environments, interpreting these constraints as nudges or "influence" over agents' behavior. Our conclusions are driven by the fact that the exercise of deferrals gives later-position agents faster access to low quality offers and slower access to high quality

[21] In addition, this perverse outcome becomes less plausible when a reduction in deferrals (and spoilage) reduces the average waiting time of *arriving* agents as in Section 6.

ones. That is, uninfluenced deferral rights increase correlation in elapsed waiting time and offer quality.

This increase in correlation reduces payoff variability, benefitting risk-averse agents. Under CARA preferences we show that uninfluenced deferral rights maximize welfare starting from any given position of the list. For discounting agents, the gains from earlier access to a large reward exceeds the loss from postponing a small one. Deferral *restrictions* at earlier positions of the list—particularly over lower-quality offers—can benefit discounting agents at all later positions.

These principles extend to settings where deferral restrictions reduce the length of the waiting list (Section 6) or reduce waste when objects (e.g. donor organs) "spoil" after repeated deferrals (Section 7). Even when deferrals are restricted at some position k, risk-averse agents would prefer there to be unrestricted deferrals at all earlier positions, whereas discounting agents do not. Due to the effects of these preference characteristics we conclude that policy evaluations (e.g. for transplant organ waiting lists) should depend on more than simple throughput measurements (e.g. organ utilization rates).

While our model is sufficient to demonstrate these conclusions, its simplicity sets aside other relevant factors. Our formulation rules out history-dependent policies, e.g. those allowing an agent to make at most d deferrals before being expelled. Though such policies do not fit directly within our framework our results suggest welfare implications nonetheless. Starting from $d = \infty$ ("no-influence"), a reduction in d makes strategic agents less choosy—deferral rates are reduced across *all* positions. Under risk-aversion such reductions create negative welfare effects (Theorem 4) that may or may not be offset by any positive effects (e.g. shorter list length as in Section 6). Under discounting, on the other hand, such reductions can benefit agents starting in later positions (e.g. Proposition 8).

We also set aside the potential efficiency gains achieved by *encouraging* deferrals when agents have heterogeneous preferences. Leshno (2019) rewards deferrals of "mismatched" objects by reducing the agent's expected waiting time for a better match. To the extent that this approach gives agents faster access to better offers and slower access to worse ones, our results suggests that discounting would enhance the benefits of this approach while risk-aversion might mitigate them. Of course the tradeoffs between these various effects are dependent on the underlying primitives. Formal analysis is left to future work on specific applications where such primitives can be specified precisely.

Acknowledgements

I thank Hyoduk Shin for a discussion that motivated this paper, and Philip Marx and Edwin Munoz-Rodriguez for research assistance. The excellent comments of two referees led to considerable improvement in the paper. Helpful comments were provided by Nick Arnosti, Anna Bogomolnaia, Ozan Candogan, Eun Jeong Heo, Edwin Munoz-Rodriguez, Tayfun Sönmez, Kathy Spier, and seminar participants at Duke (Fuqua), N.C. State, Paris School of Economics, Rochester, and Vanderbilt.

References

Abdulkadiroğlu, Atila, and Tayfun Sönmez, 1999. "House Allocation With Existing Tenants." *Journal of Economic Theory*, 88(2):233–260.

Agarwal, N., I. Ashlagi, M.A. Rees, P.J. Somaini, and D.C. Waldinger, 2019. "An empirical framework for sequential assignment: The allocation of deceased donor kidneys." NBER working paper.

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, 2019. "Thickness and Information in Dynamic Matching Markets." *Journal of Political Economy*, forthcoming.

Arikan, Mazhar, Baris Ata, John Friedewald and Rodney Parker, 2017. "Enhancing Kidney Supply Through Geographic Sharing in the United States." *Production and Operations Management*, published online 2017.

Arnosti, Nick, and Peng Shi, 2020. "Design of Lotteries and Wait-Lists for Affordable Housing Allocation." *Management Science*, 66:2291–2799.

Ata, Baris, John Freidewald and Cem Randa, 2019. "Organ Transplantation," in *Handbook of Healthcare Analytics*, T. Dai and S. Tayur (Eds.), John Wiley & Sons., In Press.

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv, 2019. "Optimal Dynamic Matching." Mimeo, Princeton University.

Bandi, Chaithanya, Phebe Vayanos and Nikolaos Trichakis, 2018. "Robust Multiclass Queuing Theory for Wait Time Estimation in Resource Allocation Systems." *Management Science*, 65:152–187.

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

Bloch, Francis and David Cantala, 2017. "Dynamic Assignment of Objects to Queuing Agents," *American Economic Journal: Microeconomics*, 9(1):88–122.

Chen, Zhou, Qi Qi, Changjun Wang, and Wenwei Wang, 2018. "What Is the Optimal Deferral Number in Waitlist Mechanism." Mimeo, HKUST.

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, 2018. "A Theory of Stability in Dynamic Matching Markets." Mimeo, Cal Tech.

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

Kadam, Sangram V., and Maciej H. Kotowski, 2018. "Multi-period Matching." *International Economic Review*, 59:1927–1947.

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

Kindy, Kimberly, Lenny Bernstein, and Dan Keating. "Lives Lost, Organs Wasted." Washington Post, Dec 20, 2018.

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

Kurino, Morimitsu, 2019. "Credibility, efficiency, and stability: a theory of dynamic matching markets." *The Japanese Economic Review*, forthcoming.

Leshno, Jacob, 2019. "Dynamic Matching in Overloaded Waiting Lists." Mimeo, University of Chicago.

Little, John, 1961. "A Proof for the Queuing Formula: $L = \lambda W$." *Operations Research*, 9(3):383–87.

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

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

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

Su, Xuanming and Stefanos Zenios, 2004. "Patient Choice in Kidney Allocation: The Role of the Queueing Discipline." *Manufacturing & Service Operations Management*, 6(4):280–301.

Thakral, Neil, 2016. "The Public-Housing Allocation Problem: Theory and Evidence from Pittsburgh." Mimeo, Harvard University.

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

9 Appendix: proofs

Proof of Proposition 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 $i \in W(k-1)$ given that that the agent starts in position $k-1$ is $\mu_i/\mu_{W(k-1)}$. Hence the probability of ultimately consuming $i \in W(k-1)$ conditional on starting in position k is

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

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

Proof of Proposition 2. The result is a consequence of Little's Law. Alternatively it follows from combining Proposition 1 with Theorem 1 proven independently below. □

Proof of Theorem 1. 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

$$\Pi(1; W) = \frac{\sum_{i \in W(1)} \mu_i v_i}{\mu_{W(1)}} - \frac{1}{\mu_{W(1)}} = \frac{\sum_{i \in W(1)} \mu_i v_i - 1}{\mu_{W(1)}}$$

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_i. Accounting for these two possibilities, along with the expected waiting time for the arrival from $W(k)$, we have the following.

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

proving the result. □

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 1,

$$\pi(C) = \frac{\sum_{i \in C} \mu_i v_i - k}{\mu_C}$$

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

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

which (weakly) improves on $\pi(C)$ if and only if $v_j \geq \left(\sum_{i \in C} \mu_i v_i - k\right) / \mu_C$. Since object types are in decreasing order of the v_i's, any W^* defined via Equation 2 maximizes $\Pi(k; \cdot)$.[22]

Since the right-hand side of the inequality within Equation 2 is decreasing in k, the type index $i^*(k)$ is increasing in the position index k. □

[22]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.

Proof of Proposition 3. 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,

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

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,

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

By Proposition 2, $E(t_{k-1}^W) = (k-1)/\mu_{W(k-1)}$. Therefore

$$\begin{aligned}
\mathrm{Var}(t_k^W) &= \mathrm{Var}(t') + \mathrm{Var}(t'') \\
&= \frac{1}{\mu_{W(k)}^2} + \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t_{k-1}^W) + \left(\frac{\mu_{W(k-1)}}{\mu_{W(k)}} - \frac{\mu_{W(k-1)}^2}{\mu_{W(k)}^2}\right)\frac{(k-1)^2}{\mu_{W(k-1)}^2} \\
&= \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t_{k-1}^W) + \frac{(k-1)^2}{\mu_{W(k)}\mu_{W(k-1)}} - \frac{(k-1)^2 - 1}{\mu_{W(k)}^2} \qquad (18)
\end{aligned}$$

which we can solve recursively.

For any policy, t_1^W is exponentially distributed with variance of $1/\mu_{W(1)}^2$ 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 18,

$$\begin{aligned}
\mathrm{Var}(t_k^W) &= \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\mathrm{Var}(t_{k-1}^W) + \frac{(k-1)^2}{\mu_{W(k-1)}\mu_{W(k)}} - \frac{k^2-2k}{\mu_{W(k)}^2} \\
&= \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\frac{1}{\mu_{W(k-1)}}\left(\sum_{\ell=1}^{k-2}\frac{2\ell}{\mu_{W(\ell)}} + \frac{2(k-1)-(k-1)^2}{\mu_{W(k-1)}}\right) + \frac{(k-1)^2}{\mu_{W(k-1)}\mu_{W(k)}} - \frac{k^2-2k}{\mu_{W(k)}^2} \\
&= \frac{1}{\mu_{W(k)}}\left(\sum_{\ell=1}^{k-2}\frac{2\ell}{\mu_{W(\ell)}} + \frac{2(k-1)}{\mu_{W(k-1)}} - \frac{k^2-2k}{\mu_{W(k)}}\right) \\
&= \frac{1}{\mu_{W(k)}}\left(\sum_{\ell=1}^{k-1}\frac{2\ell}{\mu_{W(\ell)}} + \frac{2k-k^2}{\mu_{W(k)}}\right) = \frac{1}{\mu_{W(k)}}\left(\sum_{\ell=1}^{k}\frac{2\ell}{\mu_{W(\ell)}} + \frac{-k^2}{\mu_{W(k)}}\right)
\end{aligned}$$

proving the result. $\square$

Proof of Proposition 4. Fix W, α, and a position k. To prove Equation 6, note 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 the waiting time in (i) is independent of the uncertainties in (ii), the total expected utility of (i) and (ii) is a product of two terms (see Equation 5). The first term is given by Equation 4, while the second term (in parentheses) is the expected utility of the payoffs described in (ii).

To prove Equation 7, recall by Proposition 1 an agent in position k 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 ℓ satisfies $i \in W(\ell)\setminus W(\ell-1)$ 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 arrivals from $W(k)$, $W(k-1)$, $\ldots$, $W(\ell)$.

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 3. Observe that $U_{k-1}^W \ge U_{k-1}^{W'}$ if and only if $U_{k-1}^W \ge u(v_i) = -e^{-\alpha v_i}$, i.e. $k-1$ prefers to defer i whenever the utility from v_i 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 $W(k)\setminus W(k-1) = \{i\} \cup (W'(k)\setminus W'(k-$

1)). This cancels some terms in Equation 6, so that

$$U_k^W \geq U_k^{W'} \Leftrightarrow \frac{\mu_{W(k-1)}}{\mu_{W(k)}} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k)}} u(v_i) \geq \frac{\mu_{W'(k-1)}}{\mu_{W'(k)}} U_{k-1}^{W'}$$

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

$$\frac{\mu_{W(k-1)}}{\mu_{W(k)}} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k)}} u(v_i) \geq \frac{\mu_{W(k-1)} + \mu_i}{\mu_{W(k)}} U_{k-1}^{W'}, \text{ or}$$

$$\frac{\mu_{W(k-1)}}{\mu_{W(k-1)} + \mu_i} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k-1)} + \mu_i} u(v_i) \geq U_{k-1}^{W'} \tag{19}$$

Next we express $U_{k-1}^{W'}$ in terms of U_{k-1}^W. The following equation can be derived (tediously) from Equation 7; however it can be understood as follows. After adding i to $W(k-1)$, with probability $\frac{\mu_{W(k-1)}}{\mu_{W(k-1)}+\mu_i}$ the agent receives the payoff he would have received under W, and with the remaining probability he receives $u(v_i)$. In both cases the term $\frac{\mu_{W(k-1)}+\mu_i}{\mu_{W(k-1)}+\mu_i-\alpha}$ represents the waiting cost utility as in Equation 4. However in the former case, U_{k-1}^W 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

$$\begin{aligned} U_{k-1}^{W'} &= \frac{\mu_{W(k-1)}}{\mu_{W(k-1)} + \mu_i} U_{k-1}^W \left[\frac{\mu_{W(k-1)} - \alpha}{\mu_{W(k-1)}} \frac{\mu_{W(k-1)} + \mu_i}{\mu_{W(k-1)} + \mu_i - \alpha} \right] \\ &\quad + \frac{\mu_i}{\mu_{W(k-1)} + \mu_i} u(v_i) \frac{\mu_{W(k-1)} + \mu_i}{\mu_{W(k-1)} + \mu_i - \alpha} \\ &= \frac{\mu_{W(k-1)} - \alpha}{\mu_{W(k-1)} + \mu_i - \alpha} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k-1)} + \mu_i - \alpha} u(v_i) \end{aligned}$$

Now Equation 19 becomes

$$\begin{aligned} &\frac{\mu_{W(k-1)}}{\mu_{W(k-1)} + \mu_i} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k-1)} + \mu_i} u(v_i) \\ &\qquad \geq \frac{\mu_{W(k-1)} - \alpha}{\mu_{W(k-1)} + \mu_i - \alpha} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k-1)} + \mu_i - \alpha} u(v_i) \end{aligned} \tag{20}$$

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

Proof of Theorem 4. We show that for any position k, no policy W' provides higher expected utility than W_α^*. For $k = 1$, this follows immediately from the construction of $W_\alpha^*(1)$.

For $k = 2$, suppose by contradiction that W' maximizes $U_2^{\cdot}$ and provides strictly higher expected utility than W_α^*. If $W'(1) = W_\alpha^*(1)$, then W^* maximizes $U_2^{\cdot}$ by construction of $W_\alpha^*(2)$; therefore we have $W'(1) \neq W_\alpha^*(1)$. Observe that, by construction, $W_\alpha^*(1)$ is a *threshold set*, i.e. of the form $\{1, 2, \ldots, x\}$ containing the "best x" object types. This follows from the observation that $U_1^{\cdot}$ cannot be maximized by a policy in which position 1 defers object type i and accepts some worse type $j > i$. Using similar reasoning, we can also see that $W_\alpha^*(2)$ (which maximizes $U_2^{\cdot}$ subject to fixing $W_\alpha^*(1)$) and even $W'(2)$ (which, with $W'(1)$, maximizes $U_2^{\cdot}$) also must be threshold sets, since it cannot be optimal to defer a better object than an accepted one.

Additionally, $W'(1)$ must be a threshold set. If not, it would prescribe position 1 to defer some object type i while planning to accept some worse type $j > i$. In this case position 1 could obtain higher expected utility than from $U_1^{W'}$ either by accepting i or by deferring j, i.e. either by adding $i \in W'(2)$ to $W'(1)$ or by deleting $j \in W'(2)$ from $W'(1)$. By Theorem 3, however, this improvement in position 1's expected utility would also improve the expected utility to position 2, contradicting the fact that W' maximizes $U_2^{\cdot}$.

Since $W_\alpha^*(1)$ and $W'(1)$ are threshold sets, one must be a subset of the other. If $W_\alpha^*(1) \subsetneq W'(1)$ then removing some $i \in W'(1) \setminus W_\alpha^*(1)$ from $W'(1)$ increases $U_1^{\cdot}$. By Theorem 3 this increases $U_2^{\cdot}$, contradicting the fact that W' maximizes $U_2^{\cdot}$.

Therefore $W'(1) \subsetneq W_\alpha^*(1)$. Let $i = \min\{i \in W_\alpha^*(1) \setminus W'(1)\}$ be the highest-valued object type in $W_\alpha^*(1)$ that is not in $W'(1)$. It is immediate from the construction of $W_\alpha^*(1)$ that adding i to $W'(1)$ gives position 1 expected utility higher than $U_1^{W'}$. If $i \in W'(2)$, then by Theorem 3 this change gives position 2 expected utility higher than $U_2^{W'}$, contradicting the fact that W' maximizes $U_2^{\cdot}$.

Therefore $i \notin W'(2)$. Since i is the best object type not in $W'(1)$ and $W'(2)$ is a threshold set, this implies $W'(2) = W'(1)$. In this case, we can see via Equation 6 that the addition of i to $W'(1) = W'(2)$ increases $U_2^{\cdot}$ (which, recall, is negative): the term corresponding to $\mu_{W'(2)}/(\mu_{W'(2)} - \alpha)$ decreases; the term corresponding to $\mu_{W'(1)}/\mu_{W'(2)}$ remains one; $U_1^{W'}$ increases; and the final summation term is null. Hence the overall effect of adding i to $W'(1)$ is to increase position 2's expected utility, contradicting the fact that W' maximizes $U_2^{\cdot}$.

For $k > 2$, the proof continues inductively, with analogous, but tedious, reasoning. The general idea as above is that moving an object type between the set $W(k)$ and one of the earlier sets $W(\ell)$ ($\ell < k$) either benefits both positions or hurts both positions. Since the sets $W_\alpha^*(1), \ldots, W_\alpha^*(k-1)$ already maximize expected utilities to those positions, the expected utility to a later position cannot be improved by a change that would hurt those positions, a là Theorem 3. The related, full proof of Proposition 7 is available upon request. $\square$

Proof of Proposition 5. 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 EPV_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" EPV_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

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

Combining these terms and substituting, we have

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

yielding Equation 12 for $k + 1$. □

Proof of Theorem 5. Let W ,W', $k \geq 1$, and j be such that $j \in W(k+1) \setminus W(k)$, $W'(k) = W(k) \cup \{j\}$, and $\ell \neq k \implies W'(\ell) = W(\ell)$.[23] We prove

$$EPV_{k+1}^{W} < EPV_{k+1}^{W'} \iff EPV_k^{W} > EPV_k^{W'}.$$

[23] The proof is slightly easier to read using indices k and $k + 1$ rather than those in the statement of the theorem.

First it is obvious that

$$v_j < EPV_k^W \iff EPV_k^W > EPV_k^{W'} \tag{21}$$

i.e. under policy W', position k agents would benefit from instead deferring j-type objects when v_j is less than their continuation value under W.

By (12), we can write $EPV_k^W = X/(r + \mu_{W(k)})$ where

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

Since W and W' differ only in that $W'(k) = W(k) \cup \{j\}$, we also have $EPV_k^{W'} = \frac{X + \mu_j v_j}{r + \mu_{W(k)} + \mu_j}$. Therefore by the recursion relation in Equation 11,

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

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

Hence we have $EPV_{k+1}^{W'} > EPV_{k+1}^W$ if any only if

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

With (21) this proves the result. □

Proof of Proposition 6. Under the constraint in Equation 15 the payoff is

$$\Pi\left(\frac{\lambda - \mu_{W(K)}}{\gamma}; W \right) = \frac{\sum_{i \in W(K)} \mu_i v_i - \lambda/\gamma}{\mu_{W(K)}} + \frac{1}{\gamma}$$

so the maximization exercise in (15) is equivalent to

$$\max_{\hat{\mathcal{O}} \subseteq \mathcal{O}} \frac{\sum_{i \in \hat{\mathcal{O}}} \mu_i v_i - \lambda/\gamma}{\mu_{\hat{\mathcal{O}}}} \equiv \max_W \Pi(\lambda/\gamma; W) \tag{22}$$

By Theorem 2 this is solved by setting $\hat{\mathcal{O}} = W^*(\lambda/\gamma)$, the set of types that would be accepted through position λ/γ under a no-influence policy. $\square$

Proof of Proposition 8. The arguments proving the case $r = 0$ are provided in the text. For the case $r > 0$, though preferences over policies can be evaluated explicitly using Proposition 5, we provide purely ordinal arguments for simplicity. Say that "position ℓ prefers policy W to W'" when $EPV_\ell^W > EPV_\ell^{W'}$.

We show that position $\hat{k}$ prefers W^{I_j} to W^*. Since $\hat{k}$ prefers W^* to W^S by Assumption 1, this proves the result for $k = \hat{k}$. Furthermore, due to the recursive nature of payoffs (Equation 11), this proves that any position $k > \hat{k}$ also prefers W^{I_j} to W^S, completing the proof.

Consider policy $W^{I_{\hat{k}-1}}$ which assigns type 2 objects to position $\hat{k}-1$. By Assumption 1, position $\hat{k}-1$ prefers W^* to $W^{I_{\hat{k}-1}}$. By Theorem 5 position $\hat{k}$ has opposed preferences, and prefers $W^{I_{\hat{k}-1}}$ to W^* (since $W^{I_{\hat{k}-1}}$ differs from $W^* = W^{I_{\hat{k}}}$ only in that $2 \in W^{I_{\hat{k}-1}}(\hat{k}-1) \setminus W^*(\hat{k}-1)$).

If $j = \hat{k}-1$ we are done. Otherwise we extend the idea to policy $W^{I_{\hat{k}-2}}$. By the same arguments as above, (i) position $\hat{k}-2$ prefers $W^{I_{\hat{k}-1}}$ to $W^{I_{\hat{k}-2}}$ by Assumption 1, (ii) position $\hat{k}-1$ prefers $W^{I_{\hat{k}-2}}$ to $W^{I_{\hat{k}-1}}$ by Theorem 5.

Finally, due to the recursive structure of payoffs (Equation 11), position $\hat{k}$ thus also prefers $W^{I_{\hat{k}-2}}$ to $W^{I_{\hat{k}-1}}$. (Specifically, this is because, since $W^{I_{\hat{k}-2}}(\hat{k}-1) = W^{I_{\hat{k}-1}}(\hat{k}-1)$, the specification of the set $W(\hat{k}-2)$ affects $EPV_{\hat{k}}$ in the same direction it affects $EPV_{\hat{k}-1}$.)

Repeating this argument, position $\hat{k}$ prefers W^{I_1} to W^{I_2} to…to $W^{I_{\hat{k}-1}}$ to $W^* = W^{I_{\hat{k}}}$ to W^S. Since payoffs are recursive, the same statement applies to any position $k \geq \hat{k}$. $\square$

Proof of Proposition 9. Given our definitions of W^S and W^{I_1}, Equation 7 gives the expected utility to position $k \geq 2$ as

$$U_k^{W^S} = -e^{-\alpha v_1}\left(\frac{\mu_1}{\mu_1 - \alpha}\right)^k$$
$$U_k^{W^{I_1}} = \frac{\mu_1(-e^{-\alpha v_1}) + \mu_2(-e^{-\alpha v_2})}{\mu_1 + \mu_2} \cdot \left(\frac{\mu_1 + \mu_2}{\mu_1 + \mu_2 - \alpha}\right)^k$$

Therefore $U_k^{W^S} \geq U_k^{W^{I_1}}$ if and only if

$$
\begin{aligned}
e^{-\alpha v_1}\left(\frac{\mu_1}{\mu_1-\alpha}\right)^k &\leq \frac{\mu_1(e^{-\alpha v_1})+\mu_2(e^{-\alpha v_2})}{\mu_1+\mu_2}\cdot\left(\frac{\mu_1+\mu_2}{\mu_1+\mu_2-\alpha}\right)^k \\
(\mu_1+\mu_2)\left(\frac{\mu_1(\mu_1+\mu_2-\alpha)}{(\mu_1+\mu_2)(\mu_1-\alpha)}\right)^k &\leq \frac{\mu_1(e^{-\alpha v_1})+\mu_2(e^{-\alpha v_2})}{e^{-\alpha v_1}} \\
(\mu_1+\mu_2)\left(\frac{\mu_1(\mu_1+\mu_2-\alpha)}{(\mu_1+\mu_2)(\mu_1-\alpha)}\right)^k &\leq \mu_1+\mu_2 e^{\alpha(v_1-v_2)} \\
e^{\alpha(v_1-v_2)} &\geq \frac{\mu_1+\mu_2}{\mu_2}\left(\frac{\mu_1(\mu_1+\mu_2-\alpha)}{(\mu_1+\mu_2)(\mu_1-\alpha)}\right)^k-\frac{\mu_1}{\mu_2} \\
v_1-v_2 &\geq \frac{1}{\alpha}\ln\left(\frac{\mu_1+\mu_2}{\mu_2}\left[\frac{\mu_1(\mu_1+\mu_2-\alpha)}{(\mu_1+\mu_2)(\mu_1-\alpha)}\right]^k-\frac{\mu_1}{\mu_2}\right)
\end{aligned}
$$

It is straightforward to verify that the term in square brackets exceeds 1, and hence that the right hand side increases in k. □