Influencing Waiting Lists

James Schummer*

November 20, 2019

Abstract

Stochastically arriving objects of varying quality (e.g. transplant organs, public housing units) often are allocated via waiting lists exhibiting *deferral rights*: agents may decline offered objects while keeping their position in line. We consider the welfare implications of bestowing these rights, concluding that their desirability depends—in opposite ways—on the agents' risk-aversion and impatience. For risk-averse agents deferral rights are typically Pareto improving. For discounting agents, certain restrictions on deferral rights benefit almost all agents who have or will join the list. Applying our results to a stylized model of "organ spoilage" (non-durable objects) demonstrates that policy evaluations cannot be based solely on simple metrics (e.g. organ utilization rates) that do not indicate preference characteristics.

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

*MEDS, Kellogg School of Management, Northwestern University. Email: schummer@kellogg.northwestern.edu. I thank Hyoduk Shin for a discussion that eventually motivated this paper, and Philip Marx and Edwin Munoz-Rodriguez for research assistance. 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.

1 Introduction

Waiting lists are used to allocate stochastically arriving objects such as transplant organs, public housing units, openings in substance abuse treatment programs, vacancies in child care, and more. Agents are prioritized by a pre-specified ordering, e.g. by time of arrival, health condition, neediness, or some more general scoring rule. When an object becomes available it is offered to the highest priority ("earliest position") agent remaining on the waiting list. That agent may not be obligated to accept the offer, however. In many wait list settings agents are granted *deferral rights*, allowing them to decline an offered object while maintaining their position in the list. A deferred object is offered to the next agent in line, who also may have deferral rights. We consider the welfare effects of granting or restricting these rights.

Agents who already occupy early positions of the waiting list clearly benefit from the bestowal of deferral rights; e.g., the first agent in line has access to all arriving objects and decides for herself whether to accept an offer or wait for a better one. Agents in later positions benefit similarly, *conditional* on reaching earlier positions in which they have the chance to exercise the same option. Before that happens, however, deferrals increase such agents' expected waiting time to reach earlier positions where they have acceess to higher quality offers, while decreasing the average wait for lower quality offers. Without further assumptions on preferences there are ambiguous overall welfare implications for agents joining such lists.

Our approach in addressing this topic is a generalized version of the following thought experiment. Imagine asking an agent already occupying some position in a waiting list to selfishly decide whether deferral rights should be granted to all agents or to none. How does the agent's answer depend on her current position in the queue, or on preference characteristics such as risk tolerance or discount rate? More generally, would the agent instead prefer that deferral rights only be *partially* granted? The planner might allow agents to defer only certain types of objects, or only from certain positions in the list. Asking this kind of question to all agents in a waiting list, what general welfare tradeoffs exist as we partially bestow or curtail deferral rights?

To formalize these questions, we specify a general class of allocation schemes—deemed *waiting list policies*—that capture first-come-first-served waiting list systems in which the planner can control the degree to which agents exercise deferral rights. More specifically, our class of schemes allows for the analysis of arbitrary restrictions on agents' deferral/acceptance decisions, made as a function of an agent's current position in the queue and the type of object being offered.

In our framework, restrictions on deferral rights can be interpreted in at least two ways: as an *explicit rule* enforced by the planner, or as a form of *influence* over the agents' deferral decisions. Examples of the former involve planners who can limit deferral decisions through law or force. In some cities, for instance, public housing applicants must accepted an offered unit or leave the queue.

Examples of *influence* exist in settings where planners or third parties advise imperfectly informed agents on their deferral decisions. For instance, transplant patients accept or defer organ offers in consultation with health care providers. Particularly in marginal cases, these providers can substantially influence a patient's willingness to accept an organ. Indeed the Washington Post reports that "Some transplant centers are working hard to *persuade* patients to accept less-than-perfect organs" (emphasis added) (Kindy et al. (2018)).[1]

A more indirect example of influence would be to obscure the qualities of objects that are being offered to the agents. The inability to precisely distinguish between object types removes an agent's ability to make different deferral decisions over those types. All of these examples are captured via our concept of *waiting list policies*.

Planners may differ in the extent to which they can restrict or influence deferral decisions. We set aside the question of which waiting list policies might be "implementable" by any particular planner since our results would apply to any implementable subset of our class of waiting list policies.

We emphasize that we are specifically interested in the filtering effect on welfare described above: when earlier agents defer objects of certain types, this affects the timing with which later agents receive various types of offers, impacting their welfare. Though this effect is our main interest, there are other channels through which deferral options can impact welfare. A significant one occurs when agents have sufficiently *heterogeneous* (horizontal) preferences over objects, where deferrals allow for improved efficiency of matchings through self-sorting. Leshno (2019) demonstrates how to optimally *encourage* deferrals in such settings (with probabilistically higher

[1]Of course doctors are ethically constrained to give advice that benefits their own patients, leading to a conundrum: Could such ethical constraints lead to Pareto dominated outcomes? Our results in Section 6 address this possibility.

priority in the waiting list). When agents have vertical preferences, however, sorting plays no role, leaving the above filtering effect as the primary effect on welfare. Therefore, to deliver our insights as cleanly as possible we consider such a model of *vertically* differentiated objects.

1.1 Overview of Results

In a benchmark case of risk-neutral, non-discounting agents, we first provide a "payoff equivalence" result (Theorem 2) implying the following. If a planner commits to marginally influencing the deferral decisions of agents who reach early positions in the waiting list, then the expected payoffs to agents starting in later positions would be *unaffected*. 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" policy (fully allowing agents to make selfish deferral decisions) Pareto dominates any other policy in the benchmark case.

Under more general conditions on preferences, payoff equivalence breaks down since restrictions on (early positions') deferral rights can alter the joint distribution over (later positions') waiting time and object type consumption. For example, reducing deferrals at earlier positions typically reduces the variance of waiting time for later agents (Corollary 3). Though this effect seems to benefit risk averse agents, there is an additional effect. A restriction on (selfish) deferrals also tends to reduce the correlation between an agent's waiting time (a cost) and object quality (a benefit), harming risk averse agents. We show that this latter effect dominates through an "aligned interests" theorem, roughly stating that agents in different positions of the waiting list have common interest on whether deferral rights should be (marginally) restricted. It then follows that unrestricted deferral rights necessarily enhance welfare in this case: Under CARA preferences, any restriction on selfish deferral decisions leads to Pareto dominated outcomes (Theorem 7).

In contrast, when (risk neutral) agents discount future payoffs we obtain an "opposed interests" theorem that yields the opposite effect. To interpret the result, imagine that a planner makes a "marginal" restriction in the ability of agents to make deferrals at some position k. We show that this would harm an agent starting in position k if and only if it benefits *all* agents starting positions $k+1$ and later. Thus there is a welfare tradeoff between (the sole agent currently occupying) position $k-1$ and the (typically many agents who do or will occupy) all later positions. That is, certain restrictions on deferral rights would be considered desirable under many measures of social welfare, in that they achieve almost-Pareto-dominant distributions on outcomes. Thus our two departures from the benchmark case demonstrate how risk-aversion and impatience determine, in opposite ways, the welfare consequences of influencing or limiting agents' deferral decisions.

To illustrate these ideas further we consider a stylized "organ spoilage problem," where non-durable objects (interpreted as donor organs) are offered to agents in a waiting list, but must be discarded after being deferred sufficiently many times. This spoilage constraint can be written conveniently as a form of influence in our framework, where agents in sufficiently late positions of the queue are required to "defer" objects of any type because they have spoiled by the time they would have been offered. This allows us to easily incorporate the spoilage phenomenon into our model and to assess the welfare implications of allowing or restricting deferral decisions in such settings. Based on our earlier results we argue that (at least marginal) restrictions on deferral rights can be beneficial from a welfare perspective, but that a stronger case can be made when agents significantly discount future payoffs (or are sufficiently risk neutral).

1.2 Related Literature

Our environment has some relation to others in the operations and economics literatures. One special case of our model is the *parallel processor problem*, where a set of jobs of unknown sizes must be completed by a set of processors with differing speeds. Agrawala et al. (1984) derive the utilization policy that minimizes the expected sum of waiting and processing times across all jobs, which is a special case of Equation 2. Their optimality result can be reinterpreted as a *utilitarian welfare* result, and thus is also a corollary of our Theorem 3.[2] Su and Zenios (2004) also obtain this kind of utilitarian result in a model similar to ours: wait-listed patients are offered randomly arriving transplant kidneys of varying quality and have unrestricted deferral rights. Though this also would follow from our Theorem 3, Su and Zenios's model also allows for agent arrivals and deaths.

When accounting for the potential welfare of agents *yet to arrive* to a waiting list, it is known that equilibrium behavior in queues need not max-

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

imize total (utilitarian) welfare. Naor (1969) shows that an agent deciding to join a queue fails to internalize the additional (expected) waiting cost he imposes on future arrivals. Restated in a waiting list setting, an agent deciding to defer an offer does not internalize the additional waiting cost imposed on future arrivals. Hassin (1985) solves this problem by switching the waiting list's priority structure to *last*-come-first-served (LCFS), forcing agents to internalize the probabilistic arrival of future agents when making deferral decisions. Su and Zenios (2004) exploit this idea to estimate hypothetical welfare gains from using LCFS for kidney allocation. At the same time, they point out (as does Hassin) that LCFS methods are not practical: they are manipulable in various ways, increase risk to the agents, are inequitable, and are unlikely to be politically acceptable. For these reasons we do not consider such an approach here.

An emerging economics literature on dynamic matching and assignment focuses on agents with heterogeneous (horizontal) preferences over the objects in discrete time settings. As mentioned above, Leshno (2019) reduces coordination inefficiency in such settings via a "buffer policy" that, essentially, rewards agents who defer objects by probabilistically improving their waiting list positions. In a related model where object types prioritize agents differently, Thakral (2016) proposes a multiple-list procedure an arriving object "proposes" to its highest priority agent, requiring that agent either to accept the object immediately or to commit to a separate waiting list dedicated to a specific object type. Combining this with an idea of Abdulkadiroğlu and Sönmez (1999) yields a strategy-proof mechanism that has a desirable efficiency property and respects the objects' priority orders. Arnosti and Shi (2017) compare multiple waiting list procedures (with and without deferral rights) and the use of lotteries to allocate horizontally differentiated objects.

Bloch and Cantala (2017) consider both (nonpersistent) horizontal preferences as above, and vertical preferences as we do. They model a constant-sized waiting list (new agents replace departing ones), and focus attention on allocation schemes which depart from first-come-first-served in that objects are offered to agents by some randomized reordering of their queue positions (and where agents always have full deferral rights). Interestingly, despite the differences in our models, they show in the presence of full deferral rights, agents prefer the mechanism resulting from the first-come-first-served (degenerate) lottery over any other randomized-priorities mechanism in their class. Given the differences in our classes of allocation methods, their result is logically distinct from our baseline result (Theorem 3) stating that "no influence" is Pareto dominant once we take priority orders as fixed.

Finally, more distantly related work considers dynamic two-sided matching by incorporating the assignment of randomly arriving agents. In the kidney exchange model of Ünver (2010), agents *and* objects arrive in pairs, where the agents have implicit property rights over their initial endowments. Doval (2018) considers stability in 2-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, but where assortative matchings are utilitarian-efficient.[3]

2 Model

There is a set of agents, each waiting to consume a single object, where objects arrive randomly over time and waiting is costly. Furthermore the agents are ordered, which is to be interpreted as specifying their relative positions in a waiting list. Since our results pertain to the welfare of agents conditional on their current position in the waiting list, we specify notation for the set of waiting list *positions* rather than for the agents themselves. Without loss of generality, denote the set of positions by $\mathbb{N} = \{1, 2, \ldots\}$, since our results apply in an obvious way to any finite waiting list.

There is a finite set of object types, $\mathcal{O} = \{1, 2, \ldots, n\}$, which can be interpreted as quality levels. Objects of type $i \in \mathcal{O}$ arrive according to a Poisson process with arrival rate μ_i: The time between consecutive arrivals of i-type objects is exponentially distributed with mean $1/\mu_i$. Furthermore these arrivals are independent of the arrival times of objects of any other types.[4] It is convenient to denote the arrival rate for any *set* of types $\hat{\mathcal{O}} \subseteq \mathcal{O}$ as

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

[3]Somewhat less related are dynamic models of multi-period matching by a fixed set of agents, e.g. Damiano and Lam (2005), Kurino (2019), Pereyra (2013), Kennes et al. (2014), and Kadam and Kotowski (2018).

[4]Equivalently, the arrival time between any two consecutive objects of any type is exponentially distributed with mean $1/\sum_{i \in \mathcal{O}} \mu_i$, and independently a separate random process determines each object's type using probabilities proportional to the μ_i's.

An object of type $i \in \mathcal{O}$ has value $v_i \in \mathbb{R}$, where $v_1 > v_2 > \cdots > v_n > 0$.[5] While waiting to receive an object, an agent incurs a constant flow of waiting costs normalized to one unit of cost per unit of time, i.e. an agent who waits $t \in \mathbb{R}$ units of time to receive an object incurs t units of waiting cost.

We begin with the benchmark case of risk-neutral, non-discounting agents in Section 3, where an agent who receives an i-type object after waiting t units of time receives a total payoff of $v_i - t$. We consider the implications of risk-aversion and discounting in Section 4 and Section 5.

2.1 Waiting List Policies

We specify how objects are allocated to agents by introducing the concept of a *waiting list policy*, which is the object of our analysis. This concept, which may not appear standard, is meant to capture 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 alterations to the way in which a first-come-first-served waiting list operates.

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 earlier positions).

Each of these alternatives can be formalized as some waiting list policy as we define below. In fact the full class of waiting list policies allows for more general versions of these alternatives, arbitrary combinations of them, etc. Our objective is to make welfare comparisons across all such policies.

To further illustrate the concept, consider how objects would be allocated under the first alternative above, when a planner fully allows deferrals without influence. For any agent occupying position 1 of the waiting list (i.e. the

[5] As explained in Section 1, we focus on vertically differentiated objects in order to avoid the miscoordination effect addressed by Leshno (2019). Hence each agent perceives the same value of a given object type.

highest-priority agent of those remaining), optimal (selfish) behavior in our setting 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.[6] An agent in position 2 foresees this behavior, both by the preceeding agent and himself were he to reach position 1. Due to the expected waiting time needed to reach position 1, however, optimal behavior for this agent would be to accept some set of object types $W^*(2) \supseteq W^*(1)$ by using a 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.

The construction of W^* implicitly describes equilibrium behavior of some game in which rational agents can observe their positions in the queue and the types of objects being offered. Alternately, W^* can be interpreted as a kind of probabilistic social choice function, in that it describes each agent's realized consumption (waiting time and object type) as a function of any realized sequence of object arrival times and types. Namely, W^* dictates that whenever an object of type $i \in W^*(k) \setminus W^*(k-1)$ arrives, it is consumed at that time by the agent currently in position k (or discarded if no such agent exists).

Definition 1 generalizes this latter interpretation, allowing for arbitrary specifications of the way objects are assigned (i.e. arbitrary ways in which a planner constrains deferral decisions as a function of their positions in the queue). For instance, suppose a planner has the ability to influence the deferral decisions of agents who reach position 1 of the waiting list so that they accept any object whose type belongs to some set $W(1)$ and defer the rest. When $W(1) \supset W^*(1)$ this represents some *restriction* in deferrals (e.g. the second of the three alternatives described above); agents in position 1 defer fewer object types than they otherwise would. 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 agents who reach position 1 of the waiting list.

This idea extends to positions $k = 2, 3, \ldots$ of the waiting list. A set

[6] This set is easily derived from the model's primitives; e.g. see Subsection 3.2. Related derivations appear in Su and Zenios (2004), among others.

$W(2) \supseteq W(1)$ represents those object types that are accepted by agents in one of the first two positions, while $\mathcal{O} \setminus W(2)$ are those that are deferred.[7] As a special case one can imagine informed, rational agents optimizing $W(2)$, taking some $W(1)$ is given. More generally, as above, we consider any arbitrary set $W(2) \supseteq W(1)$ that, again, represents a planner who can influence agents in a way that leads to the deferral of types $\mathcal{O} \setminus W(2)$.

Similarly, conditional on the sets $W(1) \subseteq W(2) \subseteq \cdots \subseteq W(k-1)$, a set $W(k) \supseteq W(k-1)$ represents the extent to which a planner influences deferral decisions by agents in some position k. The agents in position k are *waiting* for these object types in the sense that they defer any type in $\mathcal{O} \setminus W(k)$.

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

Non-emptiness of $W(1)$ guarantees that each agent eventually receives an object. When $W(1) = \emptyset$ our results can be stated by relabeling the positions, beginning with the first non-empty $W(k)$.

2.1.1 Examples

We introduce two classes of waiting list policies that further illustrate the concept, but also play a role in our later discussion: those in which deferral rights are offered fully or not at all. In the latter one, if deferrals are not permitted (or highly discouraged), then an agent in position 1 accepts whatever object arrives next. This scenario is described by a policy where $W(k) \equiv \mathcal{O}$. More generally, a class of policies that plays a technical role for us involves a planner who offers no deferral rights, but also withholds a (possibly empty) set of object types entirely. A *no-deferrals* policy is one that assigns object types $\hat{\mathcal{O}}$ to the agent occupying position 1 and discards the rest.

Definition 2. Waiting list policy W is a **no-deferrals policy** when, for some $\emptyset \neq \hat{\mathcal{O}} \subseteq \mathcal{O}$, we have $W(k) \equiv \hat{\mathcal{O}}$.

At the other extreme, when agents are unconstrained in their (rational, selfish) deferral decisions, we described the resulting outcomes in terms of a policy denoted W^*, above. 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

[7] To be clear, $W(2) \setminus W(1)$ is the set of types that are offered to and accepted by agents in position 2. Our results are more cleanly stated in terms of the cumulative sets $W(k)$.

on such a threshold but that the planner can influence (only) this threshold itself. In this case the planner's implementable policies are contained in the following class.

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

We interpret some of our results specifically in the case that a planner has only "marginal" influence in the agents' deferral decisions. Such a planner makes welfare comparisons only across the subclass of threshold policies. Nevertheless many of our results apply to arbitrary policies beyond the two classes defined above. Furthermore we illustrate how the organ spoilage problem mentioned in the introduction can be modeled simply as a constraint on the class of waiting list policies, in Section 6.

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* planner influences agents' decisions, instead treating it as a black box. We ignore this issue in the interest of generality, since any feasibility or implementability constraints (including ethical ones) can be imposed after the fact; they play no role in our welfare comparisons.

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 of the agent's identity or history. In one sense this good: we rule out scenarios in which a planner would treat two agents differently even when their priorities (waiting list positions) are the same. On the other hand, we are setting aside the planner's ability to punish or reward agents as a function of their previous deferral decisions.[8] We also ignore policies that use information about the current queue length, ruling out *last*-come-first-served methods. Though LCFS can help overcome inefficiencies amongst dynamically arriving agents (e.g. Naor (1969)), most real world examples of waiting lists use priority-based allocation more analogous to the FCFS systems we examine, organ transplant waiting lists being the prime example.

[8] In this vein, Chen et al. (2018) consider waiting lists in which agents lose their positions after making k deferrals, evaluating k's effect on various metrics. Abstractly, Leshno (2019) and Thakral (2016) utilize histories by placing deferring agents into secondary waiting lists.

3 Risk-neutral, Patient Agents

3.1 Equivalence results

We start by showing that the expected payoff to an agent in some position k of the waiting list is invariant to certain restrictions on deferrals made in earlier positions. More precisely, the expected (continuation) payoff from the kth position of the waiting list under an arbitrary policy W depends only on the set of types $W(k)$. In turn this implies that, in our baseline case of non-discounting, risk-neutral agents, a "no influence" waiting list policy is Pareto dominant within the class of *all* waiting list policies. In addition, the set of all achievable expected payoffs for position k can be evaluated by considering only the simple class of no-deferrals policies.

To establish some intuition, let us contrast two waiting list policies, W and W', where

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

There are no deferrals under W': each arriving object is accepted by the agent currently in position 1. Under W, some object types are deferred by position 1 agents, all of which are accepted by position 2 agents. Under either policy, an agent sitting in position 2 has some chance of consuming an object type in the subset $W(1)$. Our first observation is that the probability of this event is identical across the two policies.

Indeed under W, this happens precisely when the first object to arrive has some type in $W(1)$. This arrival, which happens with probability $\mu_{W(1)}/\mu_{\mathcal{O}}$, causes the agent to move into position 1, after which she is guaranteed to receive (eventually) some object type in $W(1)$. Under W', in contrast, the agent is guaranteed to reach position 1, and then to consume whichever object type arrives next. The probability that its type belongs to $W(1)$ is again $\mu_{W(1)}/\mu_{\mathcal{O}}$.

This intuition extends to individual object types and arbitrary positions. Conditional on an agent occupying a position k where she is "eligible" to receive some object type $i \in W(k) \setminus W(k-1)$, her probability of receiving it is $\mu_i/\mu_{W(k)}$. Since this is true for all types and positions, the probabilities of consuming different object types must remain proportional to their arrival rates. Hence whenever $W(k) = W'(k)$, policies W and W' induce equivalent distributions over object types to the agent in position k. All omitted proofs are in the Appendix.

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

Continuing with our example, consider the *expected waiting time* for an agent in position 2. Under W' the agent consumes the second object to arrive regardless of type, averaging $2/\mu_{\mathcal{O}}$ units of waiting time. Under W, the arrival of the first object (expected waiting time of $1/\mu_{\mathcal{O}}$) results in two possible cases. With probability $\mu_{W(1)}/\mu_{\mathcal{O}}$ its type is in $W(1)$, and the agent waits for a second arrival from $W(1)$ (adding $1/\mu_{W(1)}$ of expected waiting time). Otherwise the agent departs immediately. Thus the total expected waiting time is again $2/\mu_{\mathcal{O}}$.

This argument generalizes to arbitrary positions, showing that whenever $W(k) = W'(k)$, policies W and W' yield the same expected waiting time to an agent in position k. With Theorem 1 this argument leads to the expected-payoff equivalence result of Theorem 2: the expected payoff to an agent starting from position k is a function only of the set $W(k)$.

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

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

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

We prove Theorem 2 directly in the Appendix, which yields the following "waiting time equivalence" result as a corollary. To formalize it, we henceforth let $\boldsymbol{t_k^W}$ denote the random variable describing the waiting time for an agent starting in position k to receive an object under policy W.

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

As another consequence of Theorem 2, the expected payoffs that are achievable from any given position k can be determined by restricting attention to the class of no-deferrals policies.

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

3.2 Pareto dominance of the No-influence Policy

The main implication of Theorem 2 is that, once we fix a set of object types $W(k)$ that an agent in position k is "waiting for," the agent neither benefits nor loses when the planner restricts deferral rights at earlier positions of the list. Consequently, in the baseline case of risk-neutral, non-discounting agents, one waiting list policy Pareto dominates all the others, namely the one in which the planner allows agents to selfishly make deferral decisions in their own self interest. This "no-influence" policy corresponds to what would be equilibrium outcomes in a natural, dynamic "waiting list game" where agents can make unconstrained, uninfluenced deferral decisions.

To formalize this without introducing the machinery of equilibrium solution concepts, we establish which policies W would maximize the expected payoff to an agent starting in some position k. To do this it is necessary and sufficient to specify a set $W(k)$ that maximizes the right-hand side of Equation 1. Not surprisingly, such a set $W(k)$ is a "threshold set" of the form $W^*(k) = \{1, 2, \ldots, i\}$,[9] where the threshold i is the highest index (i.e. lowest quality) whose value exceeds the expected (continuation) payoff with respect to $W^*(k)$.[10]

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 set $W^*(k)$, yielding a well defined policy W^*.

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

[9] Recall that objects are indexed from best to worst.

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

Theorem 3 (Pareto dominance)**.** *The no-influence policy Pareto dominates any other policy: For any policy W and position k, $\Pi(k; W^*) \geq \Pi(k; W)$.*

The theorem immediately follows from Lemma 1. In a related model, Su and Zenios (2004) provide a result saying essentially that W^* is a utilitarian-optimal policy, i.e. maximizing the sum of expected payoffs to agents currently in the list.[11] Theorem 3 strengthens this observation; since W^* is Pareto dominant it would be optimal under any Pareto consistent welfare objective.

Finally we reiterate that W^* represents the outcomes that would be obtained when agents can make rational (uninfluenced) deferral decisions without restriction, i.e. equilibrium behavior.[12] Clearly, once an agent reaches position 1 in the queue it is optimal to accept an object if and only if it belongs to $W^*(1)$; the behavior of agents in later positions is irrelevant. Given this behavior by agents in position 1, an agent in position 2 can achieve his favorite policy by committing to accept only object types in $W^*(2)$. Repeating this argument implies that W^* describes the deferral decisions that would be made in a subgame-perfect equilibrium of the naturally defined "waiting list game" played by uninfluenced agents.

3.3 Waiting time distributions

Corollary 1 says that the expected waiting time from position k, $E(t_k^W)$, only depends on $W(k)$. Though $W(k) = W'(k)$ implies $E(t_k^W) = E(t_k^{W'})$, the distributions of t_k^W and $t_k^{W'}$ may differ. Information about the distribution of t_k^W may be relevant not only to agents in the waiting list but also to third parties. In organ waiting lists, for example, less variability in waiting time can aid in the doctor's choice of patient treatment (Bandi et al. (2018)). A third party that subsidizes waiting costs (e.g. an insurance company reimbursing dialysis costs) might prefer to see lower variability in waiting times (costs).

While the distribution of t_k^W can be difficult to describe in general, we derive the variance of waiting time for an arbitrary policy W. This allows us to show that, in general, certain kinds of restrictions on agents' deferral

[11] Agrawala et al. (1984) also provide what can interpreted as a utilitarian result in a special case of our model.

[12] We omit formalizing a solution concept since these arguments are straightforward. In addition, an analogous is provided in the related model of Su and Zenios (2004) under the restriction to threshold strategies, though in our model this restriction is unnecessary.

decisions can reduce variance in waiting times, even without affecting *average* waiting times are unaffected.

The simplest case is the class of no-deferrals policies, for which waiting time distributions *can* be pinned down.

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 follows an Erlang distribution with mean $E(t_k^W) = k/\mu_{\hat{O}}$ and variance $\mathrm{Var}(t_k^W) = k/\mu_{\hat{O}}^2$.*

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

We provide the variance of t_k^W under arbitrary policies in Theorem 4. To understand the derivation, consider the case of two object types $\mathcal{O} = \{1, 2\}$, and a policy W with $W(1) = \{1\}$ and $W(2) = \{1, 2\}$. Position 2's waiting time decomposes into $t_2^W = t' + t''$ as follows. First an agent in position 2 waits t' units of time for the first arrival of an object from $W(2) = \{1, 2\}$, where t' is exponentially distributed with parameter $\mu_1 + \mu_2$. If that first object is of type $2 \in W(2)$, the agent departs immediately with the object, waiting $t'' = 0$ additional units of time. In the remaining case that it is of type $1 \in W(2)$, the agent must wait for a second arrival of a type 1 object, which requires an additional wait of t'', which is exponentially distributed with parameter μ_1. The overall variance of t'' thus can be shown to be[13]

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

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

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

[13] In this case of position $k = 2$, t'' happens to follow a hyper-exponential distribution which has a known expression for variance. This is not true for $k > 2$ so the variance of t'' is derived recursively in our proof of Theorem 4.

Note that this exceeds the variance we would see under the no-deferrals policy W', where $W'(1) = W'(2) = \{1, 2\}$, which is $\mathrm{Var}(t_2^{W'}) = 2/(\mu_1 + \mu_2)^2$ by Lemma 2. Thus W and W' yield different variances in waiting time even though both policies provide the same expected waiting time (and expected payoff) to an agent in position 2.

The general result for arbitrary policies W and positions k is given in Theorem 4. Its proof recursively applies the arguments above, replacing object types $\{1, 2\}$ with the sets of objects $W(k-1)$ and $W(k) \setminus W(k-1)$. There, t' represents the wait for the first object from $W(k)$ and t'' represents the entire (continuation) waiting time.

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

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

The theorem shows that waiting time variance decreases as deferral rights are restricted. To be more precise, $\mathrm{Var}(t_k^W)$ decreases if, for any position $\ell < k$, we expand the set $W(\ell)$.[14] That is, if agents in some position $\ell < k$ are influenced to *no longer* defer some object type $i \in W(k)$ that otherwise would be accepted no later than position k, this necessarily reduces $\mathrm{Var}(t_k^W)$. One can show that $\mathrm{Var}(t_k^W)$ is minimized by some no-deferrals policy.

Corollary 3 (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 a constraint that $W(k) = \hat{\mathcal{O}}$ for some $\hat{\mathcal{O}} \subseteq \mathcal{O}$, the policy that minimizes* $\mathrm{Var}(t_k^W)$ *is the no-deferrals policy $W(\ell) \equiv \hat{\mathcal{O}}$.*

This result demonstrates a tradeoff between Pareto optimality and waiting time predictability. Pareto dominance is obtained with the the no-influence policy W^* (Theorem 3). By restricting deferrals at earlier positions, however, we can lower the variance of waiting times for agents starting in strictly later positions of the queue. Such restrictions can be made without affecting the expected payoffs of later positions, but clearly hurt agents starting in earlier positions.

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

Finally we note that expanding the set $W(k)$ itself has ambiguous effects on $\mathrm{Var}(t_k^W)$. Intuitively by adding a type $i \notin W(k)$ to $W(k)$, an increased consumption rate tends to lower variance. At the same time, it increases the chance that an agent in position k makes a quick exit from the waiting list, which could sometimes increase variance. Either effect could dominate.

4 Risk-averse agents

Deferral rights affect the welfare of risk-averse agents by altering not only the (marginal) distribution of agents' waiting times (as in Subsection 3.3), but also the (joint) distribution of agents' waiting time and object consumption. Nevertheless, under the assumption of constant absolute risk aversion we obtain a sharp result that strengthens the qualitative idea of Theorem 3: the "no-influence" policy is Pareto dominant.

4.1 Intuition

By Corollary 3, the variability of waiting time for an agent in position k can be reduced by restricting deferrals in earlier positions (i.e. enlarging the sets $W(\ell)$ for $\ell < k$). This might seem to benefit a risk averse agent since it does not change the agent's expected waiting time (Corollary 1).

However there is the additional fact that this change may alter the relationship between the agent's waiting time (a cost) and the value of the object consumed by the agent (a benefit). Generally speaking, variability in payoffs can be reduced by increasing the correlation between these two (cost and benefit) payoff components. If we fix a set $W(k)$, we know that this fixes the distribution of position k's object type consumption (Theorem 1). Considering the set $W(1) \subseteq W(k)$, a risk-averse agent in position k might prefer that $W(1)$ contain only the highest-value types in $W(k)$ rather than only the lowest-value types, since this gives the agent a shorter expected waiting time in the cases where the agent receives lower valued objects.

This positive correlation between waiting time and object quality occurs endogenously when deferral decisions are uninfluenced: Agents in early positions decide to wait for objects of sufficiently high quality; later agents lower their standards due to the additional waiting time. This intuition underlies our generalization of Theorem 3 to the case of constant relative risk aversion, showing that a "no-influence" policy analogous to Definition 4 (now strictly)

Pareto dominates any other waiting list policy. This result follows from an "aligned interests" theorem that is interesting in its own right, showing that agents in any two consecutive positions would agree on the desirability of certain "marginal" reductions in deferral rights.

4.2 Results

We consider agents with constant absolute risk aversion: an agent's utility from a payoff of $v - t$ is $u(v-t) = -e^{-\alpha(v-t)}$, with risk parameter $\alpha > 0$. We assume that $\alpha < \mu_i$ for each $i \in \mathcal{O}$, which is innocuous under realistic assumptions on parameters.[15]

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 $x_1, \ldots, x_k$ (e.g. payoff components) 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. Using the two facts above,

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

It is more tedious to describe U_k^W for $k > 1$ since the distribution of waiting time depends on the sequence of object types that arrive. Nevertheless Equation 4 and Equation 5 lead to a recursion relation.

[15] This assumption merely bounds expected (dis)utility. More generally our statements hold whenever $\mu_{W(1)} > \alpha$.

Lemma 3. *For any policy W, the expected utility to position $k+1$ satisfies*

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

This recursive expression yields the following result.

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

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

The main interpretation of Theorem 5 is that it tells us how a "marginal" influence in deferral decisions impacts U_k^W. To illustrate, fix a policy W, a position k, and consider an object type $j \in W(k) \setminus W(k-1)$. That is, type j is deferred by agents in positions 1 through $k-1$ but is accepted by agents in position k. Holding everything else constant, imagine that we "marginally influence" policy W so that agents in position $k-1$ accept type j objects rather than defer them. How does such a change affect the expected utility of an agent in position k? There are two opposite effects, highlighted by Equation 7.

- First, 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" $\mu_{W(k-1)}/(\mu_{W(k-1)} - \alpha) > 1$, which decreases U_k^W. The interpretation for this is that *an agent in position k is worse off due to a relatively longer expected wait for an object of type j.*
- Second, 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 for this 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 on position $k-1$'s deferral rights over j-type objects. 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 (Theorem 1) nor the average waiting time from position k (Corollary 1), this correlation effect is what drives the conclusion. When v_j is sufficiently low the reverse arguments apply.

It is more straightforward to determine how this change in W affects the expected utility of an agent in position $k-1$. Requiring such an agent to accept, rather than defer, type j objects is beneficial only when v_j is sufficiently high. This is, we reach the same qualitative conclusion as we did for position k; interests are aligned.

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

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

Two agents occupying positions $k-1$ and k would agree on whether type j objects should be assigned to position k ($j \in W(k) \setminus W(k-1)$) or to position $k-1$ ($j \in W(k-1) \setminus W(k-2)$). Thus it is intuitive that we can generalize the Pareto dominance result of Theorem 3 to the case of CARA utility. Namely, a Pareto dominant policy results from a setting in which agents can make unrestricted deferral decisions in their own self interest.

Definition 5. For $\alpha > 0$, 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.[16] 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 7 (Pareto dominance)**.** *Suppose agents have CARA utility functions with parameter α. The no-influence policy W_α^* Pareto dominates every other waiting list policy.*

[16] Ties can be broken arbitrarily.

5 Discounted payoffs

Discounting reverses the effects obtained under risk aversion. For risk neutral agents who continuously discount future payoffs we obtain an *opposed interests* result mirroring Theorem 6. Its interpretation is that certain "marginal" restrictions on deferral rights would benefit agents in later positions precisely when they harm agents in earlier positions. This immediately implies that there is no Pareto dominant policy. More substantially, it shows that certain policy changes that restrict deferral rights can benefit a large number of agents in the list at the expense of the few agents who start in earlier positions.

The idea underlying the result is that a discounting agent—who faces a lottery over object types and waiting times—prefers to negatively correlate waiting time with object quality. As is intuitive, a *high value now/low value later* lottery is preferred to a *low value now/high value later* lottery, due to the greater impact of discounting on higher valued payoffs. Thus, reversing the idea described in Subsection 4.1, consider a fixed position k and set $W(k)$. An agent in position k would tend to prefer that $W(1) \subset W(k)$ contain only the lower valued types in $W(k)$ rather than only the higher valued types. Clearly, however, an agent already sitting in position 1 would have the opposite preference.

We obtain a sharp *opposed interests* result reversing the conclusion of Theorem 6. Namely, imagine the kind of "marginal" policy change between positions k and $k-1$ considered in that result. *An agent in position k (and beyond) strictly benefits from this change if and only if an agent in position $k-1$ is strictly harmed.*

5.1 Results

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

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

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

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

With these two facts it is straightforward to derive the **expected present value** (EPV) of an agent's payoff starting from position $k = 1$ under policy W. It is the EPV of the next object to arrive from $W(1)$, minus the EPV of the waiting cost for such an arrival.

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

The simplest way to think of EPV_2 is to consider an agent starting in position 2 acquiring a payoff as follows. First the agent waits for an arrival of some type $i \in W(2)$. That object is allocated either to the agent in position 1 (if $i \in W(1)$) or to the agent in position 2. In the former case, we can think of the agent in position 2 also being immediately allocated a pseudo-object with value $v = EPV_1$.

That is, consider a list of pseudo-object values v' such that $v'_i = EPV_1$ for all $i \in W(1)$, and $v'_i = v_i$ otherwise. Now EPV_2 satisfies the analog of Equation 10 with respect to v' and $W(2)$, which one can then rewrite in terms of v as follows.

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

By the same reasoning, this recursion relation generalizes to the following.

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

This leads to the following generalization of Equation 1.

Theorem 8 (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 offered 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 term 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 any such position $\ell \leq k$.

We use this payoff expression to prove our "opposed interests" result.

Theorem 9 (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 6, 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 that in the case of discounted payoffs, there is a welfare tradeoff between earlier and later positions when considering the restriction of deferral rights. If a "marginal" restriction (in the sense of limiting only the right to defer type j objects at position $k-1$) would harm an agent at position $k-1$, this would be (ex ante) beneficial to an agent in position k. In fact, due to the recursive nature of payoffs seen in Equation 11, this benefit spills over into all positions $\ell \geq k$. That is, Equation 13 can be strengthened to

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

Thus, to be clear in our wording, while positions k and $k+1$ have "opposed interests" over W vs. W', all agents in positions $\ell \geq k+1$ have common preferences over these two policies.

Finally, Theorem 9 demonstrates that no single policy is Pareto dominant for discounting agents. Indeed, whenever there exists some $j \in W(k) \setminus$

$W(k-1)$, the theorem tells us that EPV_k^W and EPV_{k-1}^W must be affected in opposite directions by placing $j \in W(k-1)$, i.e. W cannot simultaneously maximize both positions' expected, discounted payoffs. When discounting is a significant part of agents' preferences, agents later in the queue can benefit from certain restrictions placed on earlier positions' deferral rights (e.g. when early positions defer objects that are to be accepted in later positions). This observation is relevant in the next section where we give an application of our results.

6 Application: Organ Spoilage

To demonstrate some implications of our results we consider a stylized restriction in our model motivated by the "organ spoilage problem." Patients on waiting lists for transplant organs are prioritized based on various characteristics, e.g. health status, geographic location, join date, etc. Arriving organs are offered to agents sequentially, according to some such priority ordering, with each patient deciding whether to accept or defer it.[17] Due to the fact that organs have limited shelf life, this process can result in waste. Organs of (relatively) lower quality—which might be acceptable to low-priority patients—end up spoiling in the time that it takes to process the offers to (and deferrals by) high-priority patients earlier in the list.[18] In other words, the bestowal of deferral rights creates inefficiency through waste.

Abstracting away from further details of such environments,[19] we consider the desirability of deferral rights in the presence of such inefficiencies from spoilage. To what extent should deferral rights be restricted? How does the answer depend on agents' preference characteristics?

Our earlier results suggest the following intuition. Start from a scenario in which deferral rights are unrestricted, and objects are spoiling due to an excessive number of deferrals. Imagine a planner who can influence[20] agents

[17]Of course acceptance decisions are difficult ones, and hence are influenced through the advice of patients' health care providers. A growing empirical literature studies patient preferences and decisions in such settings, e.g. Arikan et al. (2017), Ata et al. (2017), Agarwal et al. (2019).

[18]See "In Discarding of Kidneys, System Reveals Its Flaws," Sack (2012).

[19]E.g. transplant waiting lists priorities can also vary with the object type (e.g. blood type compatibility), can change over time (due to health changes), etc.

[20]Such influence might occur explicitly (by forbidding or punishing deferrals) or implicitly (via the advice the agents receive from advisors, e.g. health care providers).

early in the waiting list to (marginally) lower their threshold for acceptance. This would effect the agents' welfare in two different ways. First, by reducing the number of deferrals, fewer objects spoil improving efficiency. This *resource effect* clearly benefits the many agents who begin in later positions of the waiting list by reducing their average waiting time for acceptable objects. Second, however, it also changes the timing with which agents receive offers of various qualities. As we have shown, this *preference effect* can be beneficial or harmful depending on the agents' risk aversion and discounting.

As we demonstrate below, a marginal restriction in deferral rights tends to benefit discounting agents; i.e. the preference effect above is positive. Intuitively, the *opposed interests* result (Theorem 9) shows that certain restrictions on deferrals that harm earlier positions are (ex ante) beneficial to later positions. Hence, combined with the resource effect, certain restrictions on deferral rights can be unambiguously welfare improving to all current and future agents who join later positions in the queue.

For risk averse agents, in contrast, our *aligned interests* result leads to the opposite preference effect: a policy change that harms agents in earlier positions creates a negative welfare effect on later ones. Whether this negative effect outweighs the resource effect is ambiguous without further formalization.

To formalize these tradeoffs and intuition more concretely, we impose an additional constraint on our waiting list model where an excessive number of offers and deferrals cause an object to spoil. One conclusion to draw from our exercise is that policymakers evaluating organ allocation policies need to consider more than summary statistics that measure organ utilization rates. They also need to evaluate overall welfare effects under different allocation policies by also considering the degree to which patients are risk averse and/or discount future payoffs.

6.1 Waiting lists with a spoilage constraint

We can sufficiently illustrate our points using a model with two object types, $\mathcal{O} = \{1, 2\}$. Recall that $v_1 > v_2$, while arrival rates $\mu_1, \mu_2 > 0$ are arbitrary. In the context of organ waiting lists we interpret type $1 \in \mathcal{O}$ objects as representing a set of *high quality* organ types that would be accepted by agents in early positions of the waiting list, even if they had the right to defer them; such organs never spoil. Type $2 \in \mathcal{O}$ objects represent *marginally undesirable* organs that early position agents would prefer to defer (in favor of waiting for a high quality organ), but that later agents would be willing to accept. Inefficiency arises when these objects spoil before being offered to those later agents willing to accept them.[21]

We model object spoilage by assuming that an arriving object can be offered only to the first κ positions of the waiting list for some κ. To motivate this assumption, imagine that it takes a small amount of processing time, say ϵ, for an agent to receive and defer an offer. After κ such deferrals are made the object spoils, i.e. $\kappa\epsilon$ is the shelf-life of an object. In the context of organs, shelf-life is typically very small relative to other parameters,[22] so as an approximation we are essentially assuming that the planner can process offers instantaneously ($\epsilon \approx 0$), while maintaining the critical point that the number of offers must be limited.

Formally, this assumption implies the following *spoilage constraint* on policies W, where κ is a fixed parameter.

$$\forall \ell \in \mathbb{N}, \quad \ell > \kappa \Longrightarrow W(\ell) = W(\kappa) \tag{15}$$

That is, an agent in any position $\ell > \kappa$ is required to "defer" any object type $i \notin W(\kappa)$ by virtue of the fact that the object has spoiled by the time an offer can be made.

For spoilage to be a possible cause of inefficiency in this model, two facts must hold. The first is that, in the absence of any restrictions on deferral rights, type 2 objects would be deferred by any agents occupying positions 1 through κ. That is, even an agent in position κ would prefer to wait for the κth arrival of a type 1 object rather than to accept a type 2 object immediately. Second, there must be some (minimal) position $\hat{k} > \kappa$ for which the reverse is true: an agent in that position would prefer to immediately accept a type 2 object rather than wait for the $\hat{k}$th arrival of a type 1 object. (Otherwise the spoilage causes no inefficiency.)

We formalize these assumptions extending our earlier language of "no-influence" policies (see Subsection 3.2), where W^* denotes the waiting list policy that would arise when selfish agents can make unconstrained deferral decisions, *setting aside the spoilage constraint*. That is, W^* would represent equilibrium outcomes in our model before we impose the constraint in Equation 15.

[21]More specifically, type 2 represents the highest quality organs among those that end up spoiling, which are the most desirable to save from spoilage.

[22]E.g. kidneys have a shelf-life on the order of one day, whereas the average time patients spend on such lists is measured in years.

Of course the exact specification of W^* depends on the preferences of the agents. When agents do not discount future payoffs, W^* is defined in Definition 4 in the risk neutral case, and in Definition 5 in the risk averse case. Analogously, when (risk neutral) agents discount payoffs, let W^* be such that:

- $W^*(1)$ is the set $W(1)$ that maximizes Equation 12 for $k = 1$;
- given $W(1) = W^*(1)$, $W^*(2)$ is the set $W(2)$ that maximizes Equation 12 for $k = 2$;
- and so on: for each k, $W^*(k)$ maximizes Equation 12 subject to the previous $W^*(j)$'s, $j < k$.

We will separately consider the risk averse and discounting cases below. In all cases, however, we clearly have $1 \in W^*(1)$ (the first agent in line accepts the best object type), and hence $1 \in W^*(k)$ for all k. The assumption that *spoilage creates inefficiency* thus becomes the following.

Assumption 1 (deferrals create wasteful spoilage)**.** There exists $\hat{k}$, with $\kappa < \hat{k} < \infty$, such that $2 \in W^*(\hat{k}) \setminus W^*(\hat{k} - 1)$.

That is, in the absence of spoilage, $\hat{k} > \kappa$ would be the first position at which agents would accept type 2 objects. Since $\hat{k} > \kappa$, spoilage occurs (W^* fails (15)).

6.2 Restricting deferrals

Our goal is to evaluate the potential benefits of (marginally) reducing deferral rates in earlier positions. In the context of our two-type model, this means considering waiting list policies in which type 2 objects are no longer deferred by agents who reach some position $j \leq \kappa$. For any position j, define policy W^{I_j} as follows.

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

A policy W^{I_j} satisfies the spoilage constraint (Equation 15) only when $j \leq \kappa$, but the cases $j > \kappa$ play a role in some arguments below.

In the absence of spoilage, unrestricted deferral decisions (by selfish agents) would result in outcomes described by policy W^* as described above. This policy is defined as follows, for $\hat{k}$ defined as in Assumption 1.

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

In the presence of spoilage, however, unrestricted deferral decisions cause type 2 objects go to waste since they are deferred by positions 1 through κ. We formalize this scenario as policy W^S.

$$W^S(k) \equiv \{1\}.$$

Our main question in this section becomes the following: To what extent can we make a welfare comparison between W^S (unrestricted deferral rights, which allow spoilage to occur) and W^{I_j} (preventing spoilage, but restricting decision rights at some position $j \leq \kappa$)?

Clearly, for an agent *already occupying* position $j \leq \kappa$ in the waiting list, a switch from W^S to W^{I_j} is harmful by definition: since $W^S(j) = W^*(j) = \{1\}$, an agent in position j prefers to wait for a type 1 object rather than to immediately accept a type 2 object ($W^{I_j}(j) = \{1,2\}$).

However our interest is in the (many) agents who currently (or will) occupy the later positions $k \geq \hat{k}$ of the waiting list, i.e. the agents for whom type 2 objects are considered desirable in the sense that $2 \in W^*(k)$. Do these agents gain by a switch from policy W^S to policy W^{I_j}? Following our earlier intuition, the answer depends on the agents' preference characteristics.

6.2.1 Risk neutral agents

The result below shows that when agents are risk-neutral (and whether they discount or not), a restriction in deferral rights reduces spoilage in a way that has overwhelmingly positive welfare effects. A switch to some policy W^{I_j} eliminates the ability to defer type 2 objects at position $j \leq \kappa$. While this harms the few agents who initially start in some earlier positions (j through $\hat{k} - 1$), it unambiguously increases the expected payoffs for all of the agents occupying all positions $\hat{k}$ and beyond.

Proposition 1 (later positions benefit from deferral right restrictions)**.** *Suppose agents are risk-neutral and discount with rate $r \geq 0$, and that Assumption 1 holds. For any j, k satisfying $j \leq \kappa < \hat{k} \leq k$, an agent in position k*

benefits from the restriction of position j's deferral rights (i.e., receives a higher expected, discounted payoff under policy W^{I_j} than under W^S).

The argument behind the result is simplest in the non-discounting case ($r = 0$). By construction, for j, k as above, $W^{I_j}(k) = W^*(k)$. Hence by Theorem 2 an agent in position $k \geq \hat{k}$ is indifferent between these policies. By Theorem 3) (with Assumption 1), the agent prefers these policies to W^S. That is, agents in positions beyond $\hat{k}$ receive an expected payoff *as if spoilage were not a constraint*.

Discounting further amplifies this argument. The general idea is that, due to Theorem 9, an agent in position k becomes *strictly* better off when we require an earlier position $j < k$ to accept an object that otherwise would be accepted at position k, as long as an agent in position j would choose to defer it. We refer the reader to the formal proof for details.

If we think of a waiting list where $\hat{k}$ might be very small relative to the total length of the queue, this result provides a strong argument in favor of (at least marginally) restricting deferral rights, or otherwise imposing policies that encourage fewer deferrals. Indeed, when type 2 objects are interpreted as "marginally undesirable" ones from the perspective of early positions in the queue, $\hat{k}$ would tend to be small (close to κ), further justifying this statement.

6.2.2 Risk averse agents

When agents are risk-averse, a restriction in deferral rights may reduce the spoilage of objects, but has a negative welfare effect along the lines of Theorem 6. In particular we show how a switch to policy W^{I_j} can *lower* the expected utility to an agent occupying some position $k \geq \hat{k}$, even though W^{I_j} increases the utilization rate of "desirable" type 2 objects. An agent in position $\hat{k}$ suffers from the fact that these objects are never offered (due to spoilage), but nevertheless prefers this spoilage to occur over a policy that would instead require agents in earlier positions to accept them.

This phenomenon should be of interest to policymakers that measure the success of waiting list policies by emphasizing organ utilization rates. If organ patients are sufficiently risk averse, a policy that increases organ utilization rates can have the perverse effect of harming some of the agents it is intended to help. Fortunately, however, even when such a phenomenon occurs, we argue below that under realistic parameters it is likely to be rare. For agents in (asymptotically) late waiting list positions, the efficiency gains from utilizing organs that would otherwise spoil outweighs any negative welfare effect from risk aversion.

We formally show when this negative welfare effect can dominate. To avoid rather tedious expressions, we do this in the case where $\kappa = 1$ (i.e., after the object is deferred by the agent in position 1, it spoils), and when Assumption 1 holds at $\hat{k} = 2$ (i.e., an agent in position 2 prefers to immediately accept a type 2 object rather than wait for the second arrival of a type 1 object). The following result provides conditions under which an agent in position k prefers to let spoilage occur (W^S) rather than to restrict position 1's deferral rights (W^{I_1}).[23]

Proposition 2 (later positions may and may not benefit from deferral restrictions)**.** *Suppose that $\kappa = 1$, and that (non-discounting) agents are risk averse with parameter $0 < \alpha < \mu_1$. Further suppose that Assumption 1 holds at $\hat{k} = 2$. Then an agent in position $k \geq 2$ benefits from the restriction of position 1's deferral rights (i.e., receives weakly higher expected utility under W^{I_1} than under W^S) if and only if*

$$v_1 - v_2 \leq \frac{1}{\alpha} \ln \left(\frac{\mu_1 + \mu_2}{\mu_2} \left[\frac{\mu_1(\mu_1 + \mu_2 - \alpha)}{(\mu_1 + \mu_2)(\mu_1 - \alpha)} \right]^k - \frac{\mu_1}{\mu_2} \right) \equiv L_{\alpha,k}. \tag{16}$$

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

Thus an agent in position k favors a restriction on deferral rights when v_2 is sufficiently close to v_1, i.e. when objects that spoil are sufficiently valuable relative to the most desirable objects by being worth at least $v_1 - L_{\alpha,k} > 0$. The fact that $L_{\alpha,k}$ increases in k means that agents in lower positions are more prone to benefit from the restriction on deferral rights. In fact there must exist a sufficiently late position k' beyond which all agents would benefit from the removal of position 1's deferral rights over type 2 objects.

Corollary 4. *Under the assumptions of Proposition 2, risk averse agents in sufficiently late positions prefer restricted deferral rights: There exists k' such that an agent in any position $k > k'$ receives higher expected utility under W^{I_1} than under W^S.*

Furthermore if v_2 is sufficiently close to v_1, then an agent in *any* position $\ell \geq 2$ would benefit from this restriction on deferral rights.

[23] The analysis for other parameter values (κ, $\hat{k}$) offers little additional insight, but requires substantially more complicated expressions than those appearing below.

7 Conclusion

We have considered the welfare implications of arbitrarily influencing (or constraining) agents' deferral decisions over vertically differentiated objects, in first-come-first-served waiting lists environments. Allowing deferrals permits agents in earlier positions to be more discriminating amongst their offers. This gives later agents quicker access to lower quality objects, but slower access to higher quality ones. As our results demonstrate, the resulting welfare consequences depend on preference characteristics.

When agents are risk-averse, the ability to correlate higher quality offers with higher waiting costs lowers payoff variability, increasing welfare. Restrictions in deferral rights tend to reduce this correlation, leading to our conclusion that deferral rights are welfare enhancing. On the other hand, agents who discount future payoffs prefer to correlate higher quality offers with shorter waiting times; postponing a small reward is less costly than postponing a large one. In this case the full bestowal of deferral rights is harmful to agents in sufficiently late positions of the waiting list.

We illustrate these effects in a constrained version of our model in which objects (e.g. donor organs) inefficiently "spoil" after excessively many deferrals. Naturally, a reduction in deferral rights reduces this inefficiency. Combined with our results above, this implies that when agents are risk neutral, certain limitations on deferral rights are welfare-improving. When agents are risk averse, however, the efficiency improvement can be offset by the negative welfare consequences of restricting agents' selfish deferral decisions. Though the net effect can be ambiguous in general (Proposition 2), we show that *marginal* restrictions on deferral rights continue benefit even risk-averse agents at sufficiently late positions in the list (Corollary 4).

The elegance of our results comes at the expense of our model's simplicity. For instance, we have focused on settings where agents have common preferences[24] Nevertheless this has been sufficient to make the point that welfare analyses in such settings requires an evaluation of preference characteristics. For example, in the context of organ allocation it would be insufficient to rely on organ utilization rates to compare different allocation policies. As we have shown, even the *direction* in which welfare is effected depends on characteristics such as risk aversion levels and discount rates.

[24] As mentioned in Section 1, when preferences are heterogeneous deferral rights enhance welfare by improving match efficiency. Future work should evaluate the tradeoff between this effect and the ones we consider by evaluating agent preferences in specific applications.

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, 2017. "How (Not) To Allocate Affordable Housing." Mimeo, Columbia University.

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 Systems.” Mimeo, University of Chicago.

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.

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

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.

8 Appendix: proofs

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

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

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 2,

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

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 $\Pi(k;\cdot)$.[25]

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

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

Otherwise $i \in W(k-1)$, so the agent moves into position $k-1$; that is, the total probability of moving into position $k-1$ is $\sum_{j\in W(k-1)}\mu_j/\mu_{W(k)}$. By the induction assumption, the probability of eventually consuming any $i \in W(k-1)$ 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 Theorem 4. The wait time t_k^W is the sum of two independent random variables: the initial wait t' until the arrival of the next object $i \in W(k)$, and the remaining wait t'', which either has the same distribution as t_{k-1}^W (if $i \in W(k-1)$) or is degenerately $t'' = 0$ (if $i \in W(k) \setminus W(k-1)$).

Since t' is exponentially distributed,

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

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

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

[25] 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.

Here,

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

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

$$\begin{aligned}
\text{Var}(t^W_k) &= \text{Var}(t') + \text{Var}(t'') \\
&= \frac{1}{\mu^2_{W(k)}} + \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\text{Var}(t^W_{k-1}) + \left(\frac{\mu_{W(k-1)}}{\mu_{W(k)}} - \frac{\mu^2_{W(k-1)}}{\mu^2_{W(k)}}\right)\frac{(k-1)^2}{\mu^2_{W(k-1)}} \\
&= \frac{\mu_{W(k-1)}}{\mu_{W(k)}}\text{Var}(t^W_{k-1}) + \frac{(k-1)^2}{\mu_{W(k)}\mu_{W(k-1)}} - \frac{(k-1)^2 - 1}{\mu^2_{W(k)}} \qquad (17)
\end{aligned}$$

which we can solve recursively.

For any policy, t^W_1 is exponentially distributed with variance of $1/\mu^2_{W(1)}$ which coincides with Equation 3. We show that if Equation 3 holds for some arbitrary $k-1$ then it holds for k. Substituting into Equation 17,

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

proving the result. □

Proof of Lemma 3. 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). □

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

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

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

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

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 $j \leq k$ satisfying $i \in W(j)$, so the last line yields Equation 7. □

Proof of Theorem 6. Observe that $U_{k-1}^W \geq U_{k-1}^{W'}$ if and only if $U_{k-1}^W \geq u(v_i) = -e^{-\alpha v_i}$, 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 $U_k^W \geq U_k^{W'}$ if and only if $U_{k-1}^W \geq u(v_i) = -e^{-\alpha v_i}$.

Observe that $\mu_{W'(k)} = \mu_{W(k)}$ and that $W(k)\backslash W(k-1) = \{i\} \cup (W'(k)\backslash W'(k-1))$. 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

$$\begin{aligned}
&\frac{\mu_{W(k-1)}}{\mu_{W(k)}} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k)}} u(v_i) \geq \frac{\mu_{W(k-1)} + \mu_i}{\mu_{W(k)}} U_{k-1}^{W'}, \text{ or} \\
&\frac{\mu_{W(k-1)}}{\mu_{W(k-1)} + \mu_i} U_{k-1}^W + \frac{\mu_i}{\mu_{W(k-1)} + \mu_i} u(v_i) \geq U_{k-1}^{W'} \qquad (18)
\end{aligned}$$

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 18 becomes

$$\begin{aligned} &\frac{\mu_{W(k-1)}}{\mu_{W(k-1)}+\mu_i} U_{k-1}^{W} + \frac{\mu_i}{\mu_{W(k-1)}+\mu_i} u(v_i) \\ &\qquad \geq \frac{\mu_{W(k-1)}-\alpha}{\mu_{W(k-1)}+\mu_i-\alpha} U_{k-1}^{W} + \frac{\mu_i}{\mu_{W(k-1)}+\mu_i-\alpha} u(v_i) \end{aligned} \tag{19}$$

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

Proof of Theorem 7. We need to show that, for any position k, there is no policy W' that 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^*. By construction of W_α^*, we must have $W'(1) \neq W_\alpha^*(1)$. Observe that, by construction, $W_\alpha^*(1)$ is a *threshold set*, i.e. a set of the form $\{1, 2, \ldots, x\}$ for some 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 some object type i while planning to accept 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, extending this argument even further, $W'(1)$ also 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 $j > i$. Therefore, 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 6, 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 both $W_\alpha^*(1)$ and $W'(1)$ are threshold sets, they must have a weak inclusion relation. Suppose $W_\alpha^*(1) \subsetneq W'(1)$. Then removing some $i \in W'(1) \setminus W_\alpha^*(1)$ from $W'(1)$ increases $U_1^{\cdot}$ which, again by Theorem 6, 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 best (highest-valued) object type in $W_\alpha^*(1)$ that is not in $W'(1)$. It is clear by the construction of $W_\alpha^*(1)$ that adding i to $W'(1)$ would give position 1 expected utility higher than $U_1^{W'}$. If $i \in W'(2)$, this change would also give position 2 expected utility higher than $U_2^{W'}$, by Theorem 6, which contradicts the definition of W'.

Therefore $i \notin W'(2)$, which implies $W'(2) = W'(1)$ since $W'(2)$ is a threshold set. In this case, we can see the implications of adding i to $W'(1) = W'(2)$ on $U_2^{\cdot}$ via Equation 6: the term corresponding to $\mu_{W'(2)}/(\mu_{W'(2)} - \alpha)$ decreases, the term corresponding to $\mu_{W'(1)}/\mu_{W'(2)}$ remains one, and $U_1^{W'}$ (which is negative) increases. Hence the overall effect of adding i to $W'(1)$ is to increase position 2's expected utility, contradicting the fact that W' maximized $U_2^{\cdot}$.

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 6. □

Proof of Theorem 8. 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 9. 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 $\ell \neq k \implies W'(\ell) = W(\ell)$. We prove[26]

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

First it is obvious that

$$v_j < EPV_k^{W} \iff EPV_k^{W} > EPV_k^{W'} \tag{20}$$

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

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

[26] The proof is slightly easier to read using indices k and $k+1$.

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 (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 (20) this proves the result. □

Proof of Proposition 1. We consider the case $r > 0$. The risk neutral case follows easily from the the arguments appearing after the proposition. Though agents' preferences over policies can be evaluated using Theorem 8, the arguments below are purely ordinal, using Theorem 9. We say that an agent in some position ℓ prefers policy W to W' when $EPV_\ell^{\cdot}$ is higher under W.

Consider the (possibly infeasible) policy $W^{I_{\hat{k}-1}}$, where type 2 objects are accepted by agents in position $\hat{k}-1$. Though this policy may violate the spoilage constraint (if $\hat{k}-1 > \kappa$), we can still conclude that, by Assumption 1, an agent in position $\hat{k}-1$ prefers W^S to $W^{I_{\hat{k}-1}}$ (i.e. prefers to defer type 2 objects). Observe that $W^*(\ell)$ and $W^{I_{\hat{k}-1}}(\ell)$ are equivalent for all $\ell \leq \hat{k}-1$. Therefore an agent in position $\hat{k}-1$ prefers W^* to $W^{I_{\hat{k}-1}}$ (i.e. prefers to defer type 2 objects).

In addition, $W^*(\ell)$ is identical to $W^{I_{\hat{k}-1}}(\ell)$ for all ℓ except that type $2 \in W^*(\hat{k}) \setminus W^*(\hat{k}-1)$. Therefore by Theorem 9 position $\hat{k}$ has the opposite preferences over these two policies, preferring $W^{I_{\hat{k}-1}}$ to W^*. In words, the hypothetical scenario ($W^{I_{\hat{k}-1}}$) in which position $\hat{k}-1$ consumes type 2 objects is even better (for position $\hat{k}$) the scenario where agents make unrestricted deferral decisions (W^*).

If $j = \hat{k}-1$ we are done. Otherwise we extend the above argument by considering the (possibly infeasible) policy $W^{I_{\hat{k}-2}}$. Using similar arguments, we conclude that (i) an agent in position $\hat{k}-2$ prefers W^* to $W^{I_{\hat{k}-2}}$ and (ii) is indifferent between W^* to $W^{I_{\hat{k}-1}}$, and therefore by Theorem 9, (iii) an agent in position $\hat{k}-1$ prefers $W^{I_{\hat{k}-2}}$ to $W^{I_{\hat{k}-1}}$.

Now, due to the recursive structure of payoffs (Equation 11), conclusion (iii) implies that an agent in position $\hat{k}$ *must also* prefer $W^{I_{\hat{k}-2}}$ to $W^{I_{\hat{k}-1}}$. Intuitively, the marginal restriction of position $(\hat{k}-1)$'s deferral rights benefits position $\hat{k}$ by the opposed interests theorem. By the same argument, restricting position $(\hat{k}-2)$'s deferral rights benefits position $\hat{k}-1$ (by opposed interests), but this also benefits the subsequent position $\hat{k}$ due to the recursive nature of payoffs.

If $j = \hat{k}-2$ we are done. Otherwise we repeat this argument: an agent in position $\hat{k}$ continues to benefit as we continue to restrict deferral rights in positions $\hat{k}-3$, $\hat{k}-4$, etc., using $W^{I_{\hat{k}-3}}$, $W^{I_{\hat{k}-4}}$, etc., until reaching W^{I_j}. At that point we have shown that position $\hat{k}$ prefers W^{I_j} to W^*. Finally note that agents in any position $\ell > \hat{k}$ also prefer W^{I_j} to W^*, again due to the recursive payoff structure (Equation 11). □

Proof of Proposition 2. 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

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

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