Formal and Real Organ Allocation

Yingni Guo Xuyang Xia*

February 19, 2026

Abstract

We study priority-based allocation with strategic intermediaries. In the U.S. transplant market, organs are offered sequentially to patients ranked by priority, but transplant centers—not patients—make acceptance decisions. Centers strategically bypass high-priority patients to secure better matches, creating a wedge between formal and real allocation. We introduce the *priority adherence index* to measure this wedge and analyze how market conditions shape it through three forces: pool thinning, batch shortening, and strategic delay. Finally, we show that a simultaneous priority-proposal rule is strategically equivalent to the current sequential-offer rule but avoids costly delays, offering a faster practical alternative.

JEL: D47, D82, I11, D44, L13

Keywords: priority-based systems, strategic intermediaries, market design, market power, justified envy, Dutch auction

1 Introduction

High-stakes resources are often allocated through priority-based systems in settings such as public housing and school seats. In deceased-donor organ allocation, for example, an

*Guo: Department of Economics, Duke University; email: yingni.guo@duke.edu. Xia: Department of Economics, Duke University; email: xuyang.xia@duke.edu.

organ is offered sequentially to patients in priority order. In theory, the mechanics are straightforward: the resource is offered down a ranked list until a candidate accepts. When a candidate is reached, they evaluate the match and accept if they find it desirable. The resource would therefore go to the highest-priority candidate willing to accept it.

In practice, however, a key institutional feature complicates this picture: intermediaries assess matches and decide on behalf of candidates. For instance, every patient needing a deceased-donor organ must be listed with and managed by a transplant center, typically a hospital division. Organ offers are routed to centers rather than patients, and centers can evaluate organ-patient match quality. It is both permitted and common for centers to reject offers without informing patients (Husain et al., 2025).

Transplant centers face different incentives than individual patients. Each center manages multiple patients, so rejecting an offer for one preserves the option to accept it for another of its patients. Centers are also monitored on center-level performance measures, such as transplant outcomes (CMS, 2007; Ng, 2025). These incentives can motivate centers to bypass patients who would individually accept, in order to transplant a patient with better expected outcomes. This ability to bypass is especially powerful when a center has multiple consecutive patients on the ranked list—a patient batch—since the center can select from the batch without risking losing the organ to another center. Recent evidence documents aggressive bypassing of higher-priority patients, raising public concerns about whether the system respects patients' priority rankings (King et al., 2023; NYT, 2025).

This paper analyzes priority-based systems with strategic intermediaries. Focusing on the U.S. organ transplant market, we investigate three questions. First, how does strategic intermediation affect adherence to priority rankings and allocation outcomes? Second, how do market concentration and competition among transplant centers shape these effects? Finally, we use our model to compare alternative allocation rules and evaluate policy reforms.

We model a market with k transplant centers; center j serves a share s_j of patients. The distribution of market shares captures market structure. When an organ becomes available, it is offered sequentially to n patients in priority order, from patient 1 (highest priority) to patient n. Each patient is affiliated with one center, with affiliation to center j occurring independently with probability s_j. Thus, in expectation, larger centers have more patients on the list than smaller centers. Once patients' affiliations are realized, each center observes the match values and priority positions of its own patients, but not other patients' match values or affiliations. When the offer reaches patient i, their center either accepts on the patient's behalf or passes to the next patient. If a center accepts the organ for one of its patients, its payoff equals that patient's match value; otherwise, its payoff is zero.

As a starting point, suppose there is no intermediation: patients observe their match values and decide whether to accept. A patient accepts if and only if their match value is positive, so the organ goes to the highest-priority patient among those with positive values. We call this outcome *formal allocation*, as it reflects the conventional expectation of a priority system. By contrast, we call the equilibrium outcome under intermediation *real allocation*. When real allocation differs from formal allocation, *justified envy* arises. We thus introduce the *priority adherence index* as the probability that the highest-priority patient among those with positive values obtains the organ, conditional on the organ being allocated to some patient. Equivalently, the index is the probability that no justified envy arises.

Consider first a monopoly center. Without competition, the center cannot lose the organ to a rival and can choose any of its n patients. It therefore accepts for the patient with the highest match value, provided that value is positive. This strategy maximizes realized match value but drives the priority adherence index far below one; indeed, internal reallocation makes patient priorities irrelevant. When match values are always positive, the index falls to $1/n$, the probability that the highest-priority patient happens to have the highest match value.

With competition, centers face a risk absent under monopoly: losing the organ to a rival.

Whether the risk is present depends on the sequence of patient affiliations. Within a patient batch—a run of consecutive patients affiliated with the same center—the center can pass the offer freely, reaching for its highest-valued patient within that batch without risk of loss. However, when the next patient belongs to a rival center, passing the organ becomes risky because the rival may accept it. This risk raises the cost of bypassing high-priority patients and induces closer adherence to priority. When match values are always positive, centers avoid this risk by never passing the offer to rivals. As a result, the center that owns patient 1 accepts the organ for its highest-valued patient in the initial batch. As the number of centers increases, consecutive patients become less likely to share the same affiliation, shortening the expected length of each batch. This *batch-shortening* effect increases the probability that the highest-priority patient obtains the organ. Competition thus increases priority adherence; as the number of centers grows large, the priority adherence index approaches one and the market outcome coincides with the formal allocation.

When match values can be negative, a new strategic margin emerges: a center may pass the organ to a rival even when the current patient yields a positive payoff. The reason is that a rival center rejects with positive probability—for example, when all of its patients have negative match values—so passing does not necessarily forfeit the organ; it may return after circulating through rivals. This can create a priority-value tradeoff for the center currently holding the offer: accept now for a sure but low positive match, or delay in the hope that the organ returns for a lower-priority but higher-value patient. We call the latter choice *strategic delay*. It arises precisely when organs are sufficiently selective—i.e., when negative match values occur often enough—so that rivals reject frequently enough.

An increase in the probability of negative match values—making the organ more selective—affects priority adherence through two opposing forces. The first is a *pool thinning* effect: as the pool of willing acceptors shrinks, the highest-priority patient among them becomes more likely to obtain the organ, thereby improving priority adherence. The second is the *strategic* *delay* discussed above, which encourages centers to bypass high-priority patients with low positive values to pursue higher match values deeper in the priority list. This behavior depresses priority adherence. We show that the interaction of these two forces can cause the priority adherence index to be non-monotonic in the probability of negative match values.

Recent policy changes have expanded the geographic scope of competition for organs, but they also amplify a binding operational constraint: on average, a transplant must occur within roughly 20 hours after recovery. Under this time pressure, organ procurement organizations (OPOs) have increasingly relied on *open offers*, in which an organ is routed to a selected center and that center can allocate it to any of its patients. The use of open offers has risen sharply, from about 2% to 19%. But this expedient comes at a cost: open offers circumvent the priority list, raising concerns about fairness (NYT, 2025). This motivates interest in alternative market designs that preserve speed while improving priority adherence. One proposal is a *priority-proposal rule*, under which each center is invited to propose at most one patient and the organ is awarded to the highest-priority patient among those proposed (Henson et al., 2026). Despite the fundamental shift from a dynamic, sequential procedure to a simultaneous proposal process, we show that this rule is strategically equivalent to the current *sequential-offer rule*. However, by avoiding the delays inherent in contacting centers one by one, it can deliver the same incentives with less wasted time, and thus merits serious consideration.

1.1 Related literature

Our work builds on the conceptual foundation of Aghion and Tirole (1997), who distinguish between formal authority (the contractual right to decide) and real authority (effective control driven by information). We extend this duality to a market design setting. We define formal allocation as the benchmark outcome if candidates observed match values and decided for themselves. In contrast, real allocation is the equilibrium outcome induced by strategic intermediaries who hold effective control. By introducing a tractable model of market structure, we explore how intermediary concentration and competition shape the wedge between formal and real allocation.

Agarwal and Budish (2021, p. 69) argue that "a holistic study of markets requires analyzing both market design and market power," and observe that "with the notable exception of auction markets and a handful of examples discussed in this chapter, these two issues have largely been studied independently." We answer this call with a framework integrating market design and market power in the transplant market, which is amenable to analyzing alternative designs and how market outcomes respond to policy reforms.

We contribute to the literature on transplant centers' strategic behavior. Munoz-Rodriguez and Schummer (2025) study an allocation policy that prioritizes patients who take certain treatments. They compare settings in which patients choose their own treatment versus settings in which centers do, and show that greater market power can improve welfare by better targeting resources. Like our paper, they model centers deciding on patients' behalf and study how market structure shapes equilibrium outcomes. However, whereas they analyze the tradeoff between better resource targeting and distorting treatment choice, we focus on acceptance decisions and the tradeoff between priority adherence and allocation value. Chan and Roth (2024) observe that centers cherry-pick the safest transplants because they are penalized for poor outcomes, while OPOs avoid recovering marginal organs because they are penalized for discards. Using a two-player laboratory experiment between one OPO and one center, they show that holistic regulation—rewarding both parties for health outcomes across the entire patient pool—increases organ recovery and appropriate transplants relative to fragmented regulation. Although we do not explicitly model OPOs, their increasing reliance on open offers to avoid discards motivates our advocacy for the priority-proposal rule over the current sequential-offer rule.

Finally, we introduce the *priority adherence index* to quantify how often real allocation coincides with formal allocation. The index is closely related to the property of "no justified envy" in the school-choice literature (Abdulkadiroglu and Sönmez, 2003), as it measures the probability that this property holds. Less justified envy is a fairness criterion in this literature, motivating the search for efficient mechanisms that reduce justified envy (e.g., Abdulkadiroglu et al., 2020; Abdulkadiroglu and Grigoryan, 2021). Abdulkadiroglu et al. (2020) show that, in many-to-one matching, top trading cycles admits less justified envy than serial dictatorship *in an average sense* when priorities are drawn uniformly at random. Our priority adherence index also admits an average interpretation, but with the relevant distribution determined by the market structure and the match value distribution rather than by a uniform prior.

2 Model

Setup. Consider a setting with $k \geqslant 1$ transplant centers and $n \geqslant 2$ patients competing for a single organ. Let x_i denote the match value between patient i and the organ, drawn independently from a distribution on $[-1, 1]$ with PDF f and CDF F. Let $x = (x_1, \ldots, x_n)$ denote the vector of match values. We classify the organ as *universal* if $F(0) = 0$, meaning all match values are strictly positive. Conversely, it is *selective* if $F(0) \in (0, 1)$, implying some match values may be negative. Throughout, when we refer to a positive match value, we mean strictly greater than zero.

Let $s = (s_1, \ldots, s_k)$ denote the market shares of the centers, where $s_j > 0$ and $\sum_{j=1}^{k} s_j = 1$. Each patient is independently affiliated with center j with probability s_j, which induces a random affiliation mapping $c : \{1, \ldots, n\} \to \{1, \ldots, k\}$. Under this affiliation, center $c(i)$ decides whether to accept the organ on behalf of patient i. For parts of the analysis, we assume uniform market shares, meaning $s_1 = \cdots = s_k = 1/k$, although the realized affiliation c

may still be imbalanced due to randomness.[1]

Allocation mechanism. The OPO allocates the organ using a *sequential-offer rule* where lower-indexed patients have higher priority. The organ is first offered to center $c(1)$ for patient 1. If rejected, the offer proceeds to center $c(2)$ for patient 2, and so on, until a center accepts or all offers are declined. If patient i receives the organ, their center $c(i)$ receives a payoff of x_i; all other centers receive zero. Figure 1 illustrates this process for $n = 5$ patients. At each decision node h_i (where the organ is offered to center $c(i)$ for patient i), center $c(i)$ either accepts or rejects.

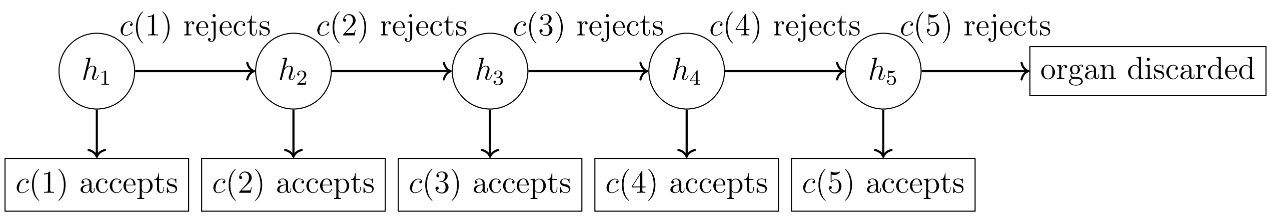

Figure 1: Sequential-offer rule for $n = 5$ patients.

Timing and information. The timing of the game and the information available to each center unfold as follows:

1. Nature draws a value vector $x = (x_1, \ldots, x_n)$ and an affiliation $c : \{1, \ldots, n\} \to \{1, \ldots, k\}$ mapping each patient to a center.

2. Each center j privately observes its own *type* $\theta_j = \{(i, x_i) : c(i) = j\}$, which includes the values and priority positions of its affiliated patients. It does not observe the values or affiliations of other patients.

3. The OPO implements the sequential-offer rule: when the organ is offered to center $c(i)$ for patient i, center $c(i)$ decides to accept or reject based on its type $\theta_{c(i)}$, knowing only that all previous offers were rejected.

[1] We use c to refer to both the random affiliation mapping and its realization, where the distinction is clear from context.

Example 1. Consider $n = 5$ patients and $k = 2$ centers, with the realized affiliation and value vector shown in Table 1. Center 1's type is $\{(1, -0.1), (2, 0.4), (5, 0.6)\}$ (blue), while center 2's type is $\{(3, -0.2), (4, 0.7)\}$ (red). For clarity, Table 2 restates the information available to center 1. With two centers, center 1 can deduce that patients 3 and 4 must belong to center 2. If there were three or more centers, center 1 would be uncertain about their affiliations. Thus, with two centers, it is as if affiliation c were publicly known.

Patient i	1	2	3	4	5
Center $c(i)$	1	1	2	2	1
Value x_i	-0.1	0.4	-0.2	0.7	0.6

Table 1: Affiliation c and value vector x

Patient i	1	2	3	4	5
Center $c(i)$	1	1	xx	xx	1
Value x_i	-0.1	0.4	xx	xx	0.6

Table 2: Information available to center 1

□

Strategies and equilibrium. Center j's type space Θ_j consists of all possible subsets of patients that could be affiliated with it, along with their possible value realizations:

$$\Theta_j := \{\theta_j = \{(i, x_i) : i \in I_j\} \mid I_j \subseteq \{1, \ldots, n\},\ x_i \in [-1, 1] \text{ for all } i \in I_j\}.$$

Given its type $\theta_j = \{(i, x_i) : i \in I_j\}$, center j's decision histories are: $H_j(\theta_j) := \{h_i : i \in I_j\}$, where h_i is the history at which patient i is offered the organ. A strategy σ_j for center j is a function:

$$\sigma_j(\cdot \mid \theta_j) : H_j(\theta_j) \to \Delta(\{\text{accept}, \text{reject}\}),$$

which, for each $\theta_j \in \Theta_j$, specifies the probability of accepting the organ at each decision history $h \in H_j(\theta_j)$. A belief function μ_j for center j is a function:

$$\mu_j(\cdot \mid \theta_j) : H_j(\theta_j) \to \Delta(\Theta_{-j}),$$

which, for each $\theta_j \in \Theta_j$, assigns a belief over the types of the other centers at each history $h \in H_j(\theta_j)$.

A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_k)$ and belief system $\mu = (\mu_1, \ldots, \mu_k)$ constitute a *Perfect Bayesian Equilibrium (PBE)* if:

1. For every center j, type $\theta_j \in \Theta_j$, and history $h \in H_j(\theta_j)$, the continuation strategy $\sigma_j(\cdot \mid \theta_j)|_h$ maximizes center j's expected continuation payoff, given belief $\mu_j(h \mid \theta_j)$ and the strategies σ_{-j} of the other centers.

2. Beliefs $\mu_j(\cdot \mid \theta_j)$ are updated via Bayes' rule wherever possible, given the strategy profile σ and the observed history h. That is, for every j, $\theta_j \in \Theta_j$, $h \in H_j(\theta_j)$, and every measurable set $S \subseteq \Theta_{-j}$,

$$\mu_j(h \mid \theta_j)(S) \;=\; \Pr\big(\theta_{-j} \in S \mid h, \theta_j, \sigma\big),$$

   whenever the conditioning event has positive probability under σ.

We adopt PBE as the solution concept. All model primitives $\{k, n, f, s\}$ are assumed to be common knowledge.

Remark 1 (Mild Refinement of PBE). PBE places no restrictions on beliefs at off-path histories. In our model, only surprising rejections require specified beliefs, since surprising acceptances, although also off-path, end the game. We assume that a surprising rejection by center j does not alter its own belief about other centers' types. This reflects the "no-signaling-what-you-don't-know" condition (Fudenberg and Tirole, 1991), which requires that a player's deviation does not reveal information they do not possess. Additionally, if an opponent center rejects unexpectedly (an event occurring only under a universal organ), we assume center j does not assign positive probability to values it had previously ruled out.

2.1 Patient-Decision Benchmark

We first consider a benchmark in which each patient i observes their own match value x_i and decides whether to accept the organ, whereas in our model centers decide on behalf of their patients.

Patients with weakly negative values reject the offer, while those with positive values accept it. Let $i_{\text{pos}}(x) := \min\{i \in \{1, \ldots, n\} : x_i > 0\}$ denote the first patient with a positive value, with $i_{\text{pos}}(x) = +\infty$ if no such patient exists. The organ is therefore allocated to patient $i_{\text{pos}}(x)$ if $i_{\text{pos}}(x) \leqslant n$, and discarded otherwise. (To simplify notation, we sometimes omit the dependence of i_{pos} on x when no confusion arises, and do the same for other terms.)

This allocation, which we term the *formal allocation*, reflects the conventional expectation of priority-based rules: the organ goes to the highest-priority patient with a positive value. To measure adherence to this principle, we define the *priority adherence index* as the probability that patient $i_{\text{pos}}(x)$ receives the organ, conditional on the organ being allocated:

$$\text{Index} := \Pr[\, i_{\text{pos}}(x) \text{ receives the organ} \mid \text{the organ is allocated}\,],$$

where the probability is taken over all realizations of the value vector $x = (x_1, \ldots, x_n)$.

The *priority adherence index* formalizes, in probabilistic terms, the same fairness principle captured by the notion of *no justified envy* in the school-choice literature (Abdulkadiroğlu and Sönmez, 2003). In that setting, an assignment is said to have no justified envy if no student prefers another student's seat while having higher priority for that seat. Analogously, in our setting, the patient $i_{\text{pos}}(x)$, the highest-priority patient with a positive value, plays the role of the "justified claimant." The priority adherence index measures how often this claimant receives the organ. An index value of one indicates perfect adherence to priority (no justified envy), whereas values below one quantify the frequency with which priority is violated (i.e., justified envy arises).

Under the patient-decision benchmark, the priority adherence index equals one. The expected allocation value is the conditional mean of positive patient values:

$$\mathbb{E}\big[x_{i_{\text{pos}}} \mid i_{\text{pos}} \leqslant n\big] = \mathbb{E}[x_i \mid x_i > 0] = \frac{\int_0^1 x\, f(x)\, \mathrm{d}x}{1 - F(0)}.$$

3 Monopoly Center ($k = 1$)

We now turn to our main analysis, beginning with the polar case of a monopoly center that observes all patient values and makes decisions for all patients. This setting highlights how information and discretion shape allocation in the absence of competition.[2]

Let $i_{\text{max}}(x) \in \arg\max_i x_i$ denote the patient with the highest value. If $x_{i_{\text{max}}} > 0$, the center rejects the organ for all patients before patient i_{max} and accepts it for i_{max}; otherwise, it rejects the organ for all patients. This strategy yields the center a payoff of $(\max_i x_i)^+$, the positive part of $\max_i x_i$.

The distribution of $x_{i_{\text{max}}}$ is given by $[F(x)]^n$. Therefore, the expected allocation value is:

$$\mathbb{E}[x_{i_{\text{max}}} \,|\, x_{i_{\text{max}}} > 0] = \frac{\int_0^1 x \cdot n[F(x)]^{n-1} f(x)\, \mathrm{d}x}{1 - [F(0)]^n}.$$

This expected allocation value is the highest achievable for our problem, since the organ always goes to the patient with the greatest value above zero. Furthermore, this value is increasing in n.

The monopoly allocation departs from perfect priority adherence because the center may bypass earlier positive-valued patients in order to reach i_{max}. Recall that the priority adherence index is the probability that the first positive-valued patient, i_{pos}, receives the organ, conditional on allocation. In the monopoly setting, since the center allocates the

[2]The monopoly setting is not hypothetical; in some isolated regions, a single center effectively dominates the top of the priority list.

organ to $i_{\max}$ as long as $x_{i_{\max}} > 0$, the index is given by:

$$\text{Index} = \Pr[\, i_{\text{pos}} = i_{\max} \leqslant n \,|\, x_{i_{\max}} > 0 \,] \,.$$

To compute this probability, suppose exactly $m \in \{1, \dots, n\}$ patients have positive values. This event occurs with probability $\binom{n}{m}[1-F(0)]^m [F(0)]^{n-m}$. Among these m positive-valued patients, the probability that the highest-priority one is also the highest-valued patient is $1/m$. Averaging over all possible m, the priority adherence index is therefore:

$$\text{Index} = \frac{\displaystyle\sum_{m=1}^{n} \frac{\binom{n}{m}[1-F(0)]^m [F(0)]^{n-m}}{m}}{1-[F(0)]^n}. \qquad (1)$$

This calculation shows that the index depends on distribution F only through $F(0)$. In monopoly, the index reflects only "positional luck": the probability that the first positive-valued patient happens to be the highest-valued among all positive-valued patients. This is a combinatorial question that depends only on the number of positive-valued patients (determined by $F(0)$), not on how their positive values are distributed. By contrast, we show in Section 4.2 that in multicenter settings the shape of F on $[0, 1]$ beyond $F(0)$ affects the index.

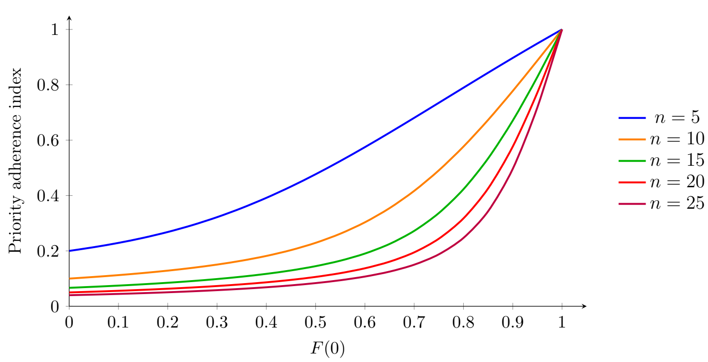

Figure 2: Effect of varying the number of patients n or $F(0)$ on the index

Figure 2 illustrates that the index decreases with n and increases with $F(0)$. As n grows, more patients seek the organ, increasing the likelihood that a later patient has a higher value than the first positive-valued patient, thereby lowering the index. Conversely, since the expected number of patients with positive values is $(1 - F(0))n$, a higher $F(0)$ implies fewer such patients, which raises the index.

Proposition 3.1. *In a monopoly market, the priority adherence index* (1) *is strictly decreasing in n and strictly increasing in $F(0)$, for all $n \geqslant 2$ and $F(0) \in (0, 1)$. Moreover, $\lim_{F(0)\to 0} Index = 1/n$ and $\lim_{F(0)\to 1} Index = 1$.*

The monopoly setting and the patient-decision benchmark lie at opposite ends of the tradeoff between priority adherence and allocation value. Figure 3 illustrates this for $n = 10$ and $x_i \sim \text{Unif}[0, 1]$. The patient-decision benchmark achieves perfect adherence but yields a low expected allocation value, whereas the monopoly setting attains the highest expected allocation value but low adherence to priority.

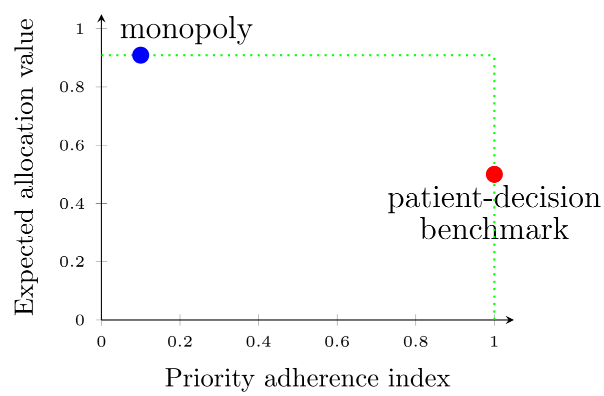

Figure 3: Monopoly versus patient-decision benchmark

4 Multiple Centers ($k \geqslant 2$)

We now turn to the multicenter setting. Patients are affiliated with centers independently according to their market shares. Each center decides whether to accept or reject the organ on behalf of its patients. Table 1 illustrates one such affiliation c: patients 1, 2, and 5 belong to center 1, while patients 3 and 4 belong to center 2. Unlike in the monopoly setting, where the affiliation is trivial, here c shapes how the organ is allocated.

To analyze centers' behavior, we introduce the *i-batch*, denoted B_i. Fix an affiliation c. For each $i \in \{1, \dots, n\}$, B_i is the maximal interval of consecutive patients containing i who are affiliated with the same center:

$$B_i := \Big\{ i' \in \{1, \dots, n\} : c(\ell) = c(i) \text{ for all } \ell \in \{\min(i, i'), \dots, \max(i, i')\} \Big\}.$$

Figure 4 illustrates the i-batch under a given affiliation c, where colors denote patients' centers; in this example, the i-batch consists of patients $i-1$ through $i+3$. We also refer to any such interval of consecutive patients affiliated with the same center as a *batch*. In Table 1, patients 1–2 form the first batch, 3–4 the second, and 5 the third.

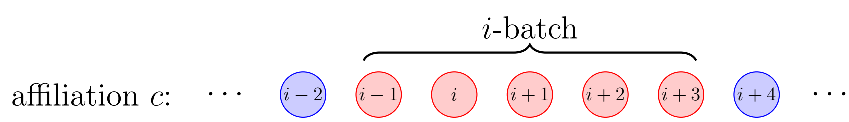

Figure 4: The i-batch under a given affiliation c; colors denote patients' centers.

4.1 Multicenter with Universal Organs ($F(0) = 0$)

We now characterize the equilibrium for universal organs.

Theorem 4.1 (Equilibrium for Universal Organs)**.** *Assume $F(0) = 0$ and $k \geqslant 2$. There exists a unique equilibrium in which, when patient i is offered the organ, center $c(i)$ accepts if i is the highest-valued among those in the i-batch from i onward; otherwise, it rejects. Consequently, the organ is allocated to the highest-valued patient in the first batch.*

Two features drive this equilibrium. First, within each batch, the owning center optimizes locally: if it accepts for some patient in that batch, it does so for the highest-valued patient, ensuring optimal intra-batch allocation. Second, the center currently offered the organ strictly prefers to accept rather than pass it along. Since the next center will accept for a patient in the next batch upon receiving the offer, a passed organ never returns. Together, these two features ensure the organ goes to the highest-valued patient in the first batch.

This equilibrium yields a closed-form expression for the priority adherence index for any $k \geqslant 2$ and market shares $s = (s_1, \ldots, s_k)$. To compute the index, we sum over centers $j \in \{1, \ldots, k\}$, where the first patient belongs to center j. For each j, we further sum over $m \in \{1, \ldots, n\}$, where m is the size of center j's first batch. In this case, the organ goes to the highest-valued patient among the first m patients. By symmetry, patient 1 is the highest-valued (and thus receives the organ) with probability $1/m$. The resulting index is:

$$\text{Index}(k, n, s) = \sum_{j=1}^{k} \left(\sum_{m=1}^{n-1} \frac{s_j^m (1 - s_j)}{m} + \frac{s_j^n}{n} \right). \tag{2}$$

The expected allocation value is computed analogously, except that when the first batch has size m, the term $1/m$ is replaced by the expected maximum of m value draws:

$$\sum_{j=1}^{k} \left(\sum_{m=1}^{n-1} s_j^m (1 - s_j) \int_0^1 x \cdot m[F(x)]^{m-1} f(x) \, \mathrm{d}x + s_j^n \int_0^1 x \cdot n[F(x)]^{n-1} f(x) \, \mathrm{d}x \right). \tag{3}$$

Figure 5 illustrates the impact of increasing the number of centers under uniform market shares ($s_1 = \cdots = s_k$). The left panel shows that the priority adherence index increases with k, since greater market competition shrinks the expected size of the first batch, bringing the index closer to the patient-decision benchmark. The right panel plots both measures: as k increases, the index rises while the expected allocation value falls, with both converging to their respective patient-decision benchmark values. We formalize these comparative statics in the following proposition.

Proposition 4.1. *Under uniform market shares and universal organs, (i) the priority adher-*

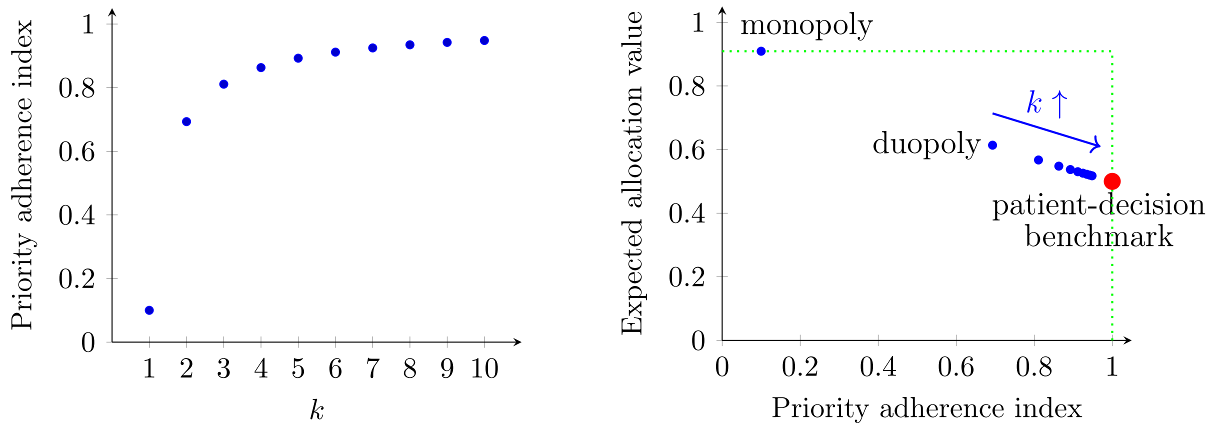

Figure 5: Effect of increasing the number of centers k: $n = 10,\ x_i \sim \text{Unif}[0, 1]$

ence index is strictly increasing in k*; (ii) the expected allocation value is strictly decreasing in* k*; and (iii) as* $k \to \infty$*, both converge to their respective patient-decision benchmark values.*

Figure 5 also shows that both measures change sharply when moving from monopoly ($k = 1$) to duopoly ($k = 2$): the priority adherence index rises dramatically while the expected allocation value drops substantially. This sharp shift reflects a structural change: introducing a second center creates a positive chance that another center cuts the first batch short. Beyond that, both measures become less sensitive to which center does the cutting, so adding more centers yields diminishing marginal effects. Notably, the probability that the first two patients belong to different centers (i.e., $c(1) \neq c(2)$) jumps from 0 to 0.5 when moving from one to two centers, accounting for most of the change in both measures.

4.2 Multicenter with Selective Organs ($F(0) \in (0, 1)$)

Our analysis to this point highlights two structural forces that tend to raise priority adherence, as illustrated in Figure 6. The first is *pool thinning*: as $F(0)$ increases, fewer patients have positive values, so the pool of patients interested in the organ shrinks. With fewer interested patients, the first positive-valued patient faces less competition, thereby increasing priority adherence. This effect is most clearly seen in the monopoly market, where

Proposition 3.1 shows that priority adherence increases in $F(0)$.

The second force is *batch shortening*: as the number of centers k increases, consecutive patients are less likely to belong to the same center, so each batch tends to contain fewer patients. Earlier, we showed that if a center ever accepts for some patient in a batch, it must do so for the highest-valued patient in that batch. Shorter batches make it more likely that the first positive-valued patient in a batch is also the highest-valued patient, thereby increasing priority adherence. This effect is most clearly seen for a universal organ, where the organ goes to the highest-valued patient within the first batch, and Proposition 4.1 shows that priority adherence increases in k.

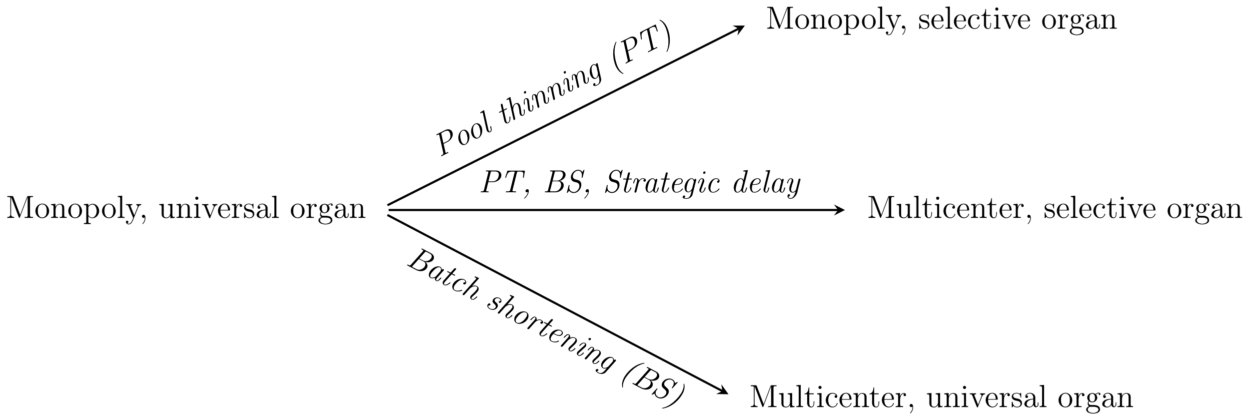

Figure 6: Structural forces shaping priority adherence

When we move to the multicenter setting with selective organs, a third force arises, *strategic delay*: a center may reject the organ for all its patients in an earlier batch even when that batch contains positive-valued patients. Unlike the first two forces, strategic delay reduces priority adherence.

To illustrate, suppose center 1 faces the situation in Table 3. Its two relevant options are: (i) accept immediately for patient 1, securing a payoff of 0.1; or (ii) delay by rejecting for patients 1 and 2 and, if the organ is rejected for patients 3 and 4, accept for patient 5, obtaining 0.9. If the likelihood that the organ is rejected for patients 3 and 4 exceeds 1/9, center 1 optimally delays to accept for patient 5, trading off patient 1's higher priority for patient 5's higher value.

Patient i	1	2	3	4	5
Center $c(i)$	1	1	xx	xx	1
Value x_i	0.1	-0.4	xx	xx	0.9

Table 3: Illustration of a situation in which center 1 may delay

To formalize this priority-value tradeoff, we characterize equilibrium behavior, focusing on the case of two centers. This case captures the essence of the game while yielding the cleanest result.

Theorem 4.2 (Equilibrium for Selective Organs and Two Centers)**.** *Assume $F(0) \in (0,1)$ and $k = 2$. Fix an affiliation c. An equilibrium $\{\alpha_i, \delta_i\}_{i=1}^n$ satisfies:*

$$\alpha_i = \Pr\big[x_i > 0 \text{ and } \delta_i x_i \geqslant \delta_{i'} x_{i'} \ \ \forall\, i' \text{ with } c(i') = c(i)\big] = \int_0^1 f(u) \prod_{i' \neq i,\ c(i')=c(i)} F\Big(\tfrac{\delta_i u}{\delta_{i'}}\Big)\, du, \quad (4)$$

$$\delta_i = 1 - \sum_{i' < i,\ c(i') \neq c(i)} \alpha_{i'}, \quad (5)$$

where α_i is the probability that center $c(i)$ makes its first acceptance at patient i, and δ_i is the probability that $c(i)$'s opponent center rejects the organ for all its patients who precede i.

For a realized value vector x, center $c(i)$ accepts for patient i if $x_i > 0$ and $\delta_i x_i \geqslant \delta_{i'} x_{i'}$ for all $i' > i$ with $c(i') = c(i)$, and rejects otherwise.

The theorem can be understood in terms of two jointly determined sequences: the *reach probabilities* $\{\delta_i\}_{i=1}^n$ and the *acceptance probabilities* $\{\alpha_i\}_{i=1}^n$. The organ can reach patient i only if $c(i)$'s opponent center rejects it for all its patients with higher priority than i; δ_i is the probability that this occurs. The expected value to center $c(i)$ from bypassing its higher-priority patients to accept the organ for patient i is $\delta_i x_i$. Each center j thus compares its patients $i \in c^{-1}(j)$ by their *discounted values* $\delta_i x_i$, accepting first for the patient with the highest such value, provided it is positive. Given this equilibrium behavior, α_i is the ex-ante probability that center $c(i)$ makes its first acceptance at patient i. In sum, $\{\delta_i\}_{i=1}^n$ captures how priority affects each patient's reach probability, while $\{\alpha_i\}_{i=1}^n$ characterizes acceptance behavior once this priority effect is internalized.

The next example solves (4)–(5) to compute $\{\alpha_i, \delta_i\}_{i=1}^n$ for a given affiliation c, and then illustrates how realized acceptance decisions follow from comparing discounted values $\delta_i x_i$ across a center's patients.

Example 2. Consider a setting with $n = 5$ patients and $x_i \sim \text{Unif}[-1, 1]$. Consider the affiliation c with $c(1) = c(2) = c(5) = 1$ and $c(3) = c(4) = 2$. Patients 1 and 2 form the first batch and have higher priority than all patients from the opponent center, so their reach probabilities equal 1. Center 2 has only one batch with patients 3 and 4, so it accepts for the patient with the higher value (if positive), yielding $\alpha_3 = \alpha_4 = 3/8$. Consequently, patient 5 has reach probability $1 - \alpha_3 - \alpha_4 = \frac{1}{4}$. Center 1 then compares the discounted values $\{x_1, x_2, (1/4)x_5\}$ and makes its first acceptance at the patient with the highest positive discounted value. The full set of equilibrium probabilities is reported in the left table.

The right table shows one realization of patient values for center 1's patients 1, 2, and 5. In this case, the discounted value $\delta_i x_i$ is highest for patient 5 (equal to 0.225), so center 1 rejects the organ for patients 1 and 2 and accepts it for patient 5.

Patient i	1	2	3	4	5
Center $c(i)$	1	1	2	2	1
Accept prob. α_i	$\frac{275}{768}$	$\frac{275}{768}$	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{61}{384}$
Reach prob. δ_i	1	1	$\frac{109}{384}$	$\frac{109}{384}$	$\frac{1}{4}$

Patient i	1	2	3	4	5
Center $c(i)$	1	1	xx	xx	1
Value x_i	0.1	-0.4	xx	xx	0.9
Disc. value $\delta_i x_i$	0.1	-0.4	xx	xx	0.225

Table 4: Equilibrium acceptance/reach probabilities (left) and strategic delay (right)

□

The preceding example illustrates strategic delay but leaves open whether it rivals pool thinning or batch shortening in shaping priority adherence. We now show that it does. Recall that pool thinning alone causes the priority adherence index to increase in $F(0)$: fewer interested patients means less competition for the first positive-valued patient. The next example shows that strategic delay can reverse this relationship, causing the index to *decrease* in $F(0)$. Intuitively, as $F(0)$ increases, the opponent center is more likely to reject for all its patients, making it more attractive for the current center to bypass earlier, lower-valued patients in favor of later, higher-valued ones.

Example 3. Consider a setting with $n = 3$ patients, $k = 2$ centers, and uniform market shares $s = (1/2, 1/2)$. The affiliation $c = (c(1), c(2), c(3))$ is uniform over $\{1, 2\}^3$: $\{(1, 1, 1), (1, 1, 2), (1, 2, 1), (1, 2, 2), (2, 1, 1), (2, 1, 2), (2, 2, 1), (2, 2, 2)\}$. Strategic delay arises only if two conditions are met: (i) the affiliation is $(1, 2, 1)$ or $(2, 1, 2)$, and (ii) both patients 1 and 3 have positive values. When these conditions hold, the center owning patients 1 and 3 must decide nontrivially whether to reject the organ for patient 1 in order to accept it for patient 3. The center owning patient 2 rejects the organ whenever $x_2 < 0$, which occurs with probability $F(0)$; hence, patient 3's reach probability is $F(0)$. Therefore, conditional on $x_1 > 0$ and $x_3 > 0$, patient 1 is bypassed if and only if $x_1 < F(0)x_3$, an event that occurs with probability $\Pr\big(x_1 < F(0)x_3 \mid x_1 > 0,\ x_3 > 0\big)$.

Consider the density function $f(x)$ for x_i parameterized by $F(0)$:

$$f(x) = \begin{cases} F(0), & x \in [-1, 0], \\ \dfrac{1 - F(0)}{2\varepsilon}, & x \in \left[\frac{1}{2} - \varepsilon, \frac{1}{2}\right] \text{ or } x \in [1 - \varepsilon, 1], \\ 0, & \text{otherwise.} \end{cases}$$

Conditional on being positive, x_i has a 50% chance of falling into the lower cluster (near

1/2) and a 50% chance of falling into the higher cluster (near 1). In the limit $\varepsilon \to 0$, the probability $\Pr\big(x_1 < F(0)x_3 \mid x_1 > 0, x_3 > 0\big)$ is:

$$\begin{cases} 0 & \text{if } F(0) < \frac{1}{2} \\ \frac{3}{16} & \text{if } F(0) = \frac{1}{2} \\ \frac{1}{4} & \text{if } F(0) > \frac{1}{2}. \end{cases}$$

Strategic delay emerges once $F(0)$ reaches 1/2; whenever it occurs, it contributes zero to the priority adherence index because the center owning patients 1 and 3 bypasses its first positive-valued patient. Accordingly, Figure 7 shows a discontinuous drop in the index at $F(0) = 1/2$. □

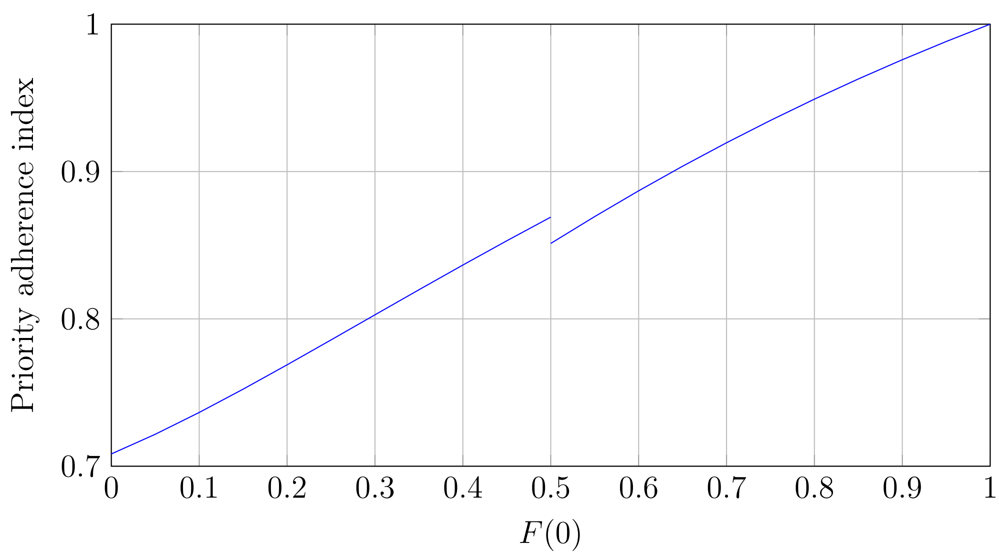

Figure 7: Priority adherence index as a function of $F(0)$

5 Alternative Allocation Rules

5.1 Two Priority-Based Rules

While the sequential-offer rule remains the dominant allocation rule in practice, alternative procedures have begun to appear. For example, staff at LiveOn NY, an OPO in New York, reported that after five hours, they invited favored hospitals to each propose one patient; the organ was then allocated to the highest-priority patient among those proposed.

Motivated by such practices, we formalize the *priority-proposal rule*: each center may propose one of its own patients, and the organ is allocated to the proposed patient with the highest priority. The game follows the same timing as in Section 2, except that in step 3 the OPO applies the priority-proposal rule. Conceptually, this rule is analogous to a first-price auction, but with two key differences: (i) the winner is determined by priority rank rather than bid amount, and (ii) there is no payment.

Likewise, the sequential-offer rule is analogous to a Dutch (descending-price) auction. The OPO starts with the highest-priority patient and proceeds sequentially down the list until a center accepts on behalf of a patient. Here the "clock" runs down the priority list rather than the price, and again there is no payment.

5.2 Strategic Equivalence and Connection to Auctions

Mirroring the strategic equivalence between the first-price and Dutch auctions, we establish a similar equivalence between the priority-proposal and sequential-offer rules.

Theorem 5.1. *The priority-proposal rule and the sequential-offer rule are strategically equivalent.*

Figure 8 illustrates the strategic and conceptual relationships between two organ allocation rules (left column) and their auction counterparts (right column). Horizontal arrows represent conceptual analogies, while vertical arrows indicate strategic equivalence. Our focus is on the allocation rules, with the auctions serving as familiar benchmarks.

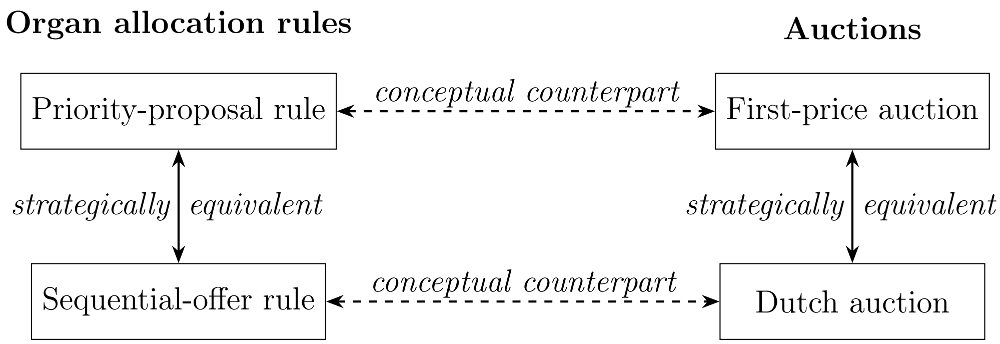

Figure 8: Strategic and conceptual relationships among four mechanisms

To better understand the connection between the priority-proposal rule and the first-price auction, we compare their type and payoff structures. The same comparison applies between the sequential-offer rule and the Dutch auction.

In the standard first-price auction, each bidder j has a one-dimensional private value v_j and, if it wins with bid b, earns payoff $(v_j - b)$, which is linear in both value and bid.

By contrast, under the priority-proposal rule each center j has a multi-dimensional type θ_j, consisting of the values and priority positions of all its patients. Although each center j is restricted to proposing one of its own patients, we can equivalently enlarge the action space to allow proposing any patient $i \in \{1, \ldots, n\}$, and define the payoff from winning with proposal i to be x_i if $i \in c^{-1}(j)$ and $-\infty$ otherwise. Under this representation, the payoff is neither monotone in the "bid" (the priority rank of the proposed patient) nor monotone in any one-dimensional summary of θ_j. These features make the problem substantially less tractable than standard auctions, explaining why we developed analytical tools tailored to priority-based systems rather than relying on familiar auction techniques.

5.3 Practical Considerations and Policy Implications

Theorem 5.1 yields two key implications. First, every equilibrium characterization for the sequential-offer rule carries over to the priority-proposal rule. Second, because the two rules are outcome-equivalent, choosing between them hinges largely on practical considerations outside our model.

One major challenge in organ allocation is time: on average, transplants must occur within 20 hours after an organ is recovered. This constraint makes the sequential-offer rule costly, as it requires the OPO to contact centers sequentially, a time-consuming process. In response, some OPOs have sent open offers to selected centers within hours of organ recovery without clear guidelines, raising fairness concerns.[3]

In contrast, the priority-proposal rule invites all centers to submit proposals simultaneously, saving the time otherwise spent contacting centers sequentially after rejections.

While the priority-proposal rule may save time in allocating the organ, it may also introduce a tradeoff: Centers with lower-priority patients may invest effort in identifying a candidate, only to see the organ go to a higher-priority center. Nonetheless, with the rate of out-of-sequence allocations rising from 2% to 19% in recent years, often justified by time sensitivity, and with growing public distrust in the transplant system, we suggest that the priority-proposal rule be seriously considered as a practical alternative.

Remark 2 (Why the Dutch-Auction Analogy Does Not Imply Speed). In some markets for perishable goods, a Dutch auction can clear sales very quickly.[4] At the Aalsmeer flower auction, for example, the price ticks down automatically; the first bidder to press the button wins, and the sale can conclude in seconds. By contrast, this automated progression is absent

[3]In interviews, the heads of Mid-America Transplant (St. Louis) and LiveOn NY (New York City) defended their policies of initiating open offers eight and five hours, respectively, after organ recovery. They explained that recent rule changes, requiring offers to be made to patients nationwide, had imposed additional time constraints.

[4]We thank Curtis Taylor for raising this point.

under the sequential-offer rule for deceased-donor organ allocation. Transplant centers can take up to one hour to decide whether to accept an offer; if they decline (or the clock runs out), the OPO must actively advance the offer to the next patient (Committee on Organ Procurement and Transplantation Policy, 2000; Mankowski et al., 2019). The OPO therefore contacts centers one by one and waits for each response. This difference—automatic progression versus response-dependent advancement—explains why the sequential-offer rule is slow despite its conceptual parallel to a Dutch auction.

6 Conclusion

We develop a canonical model of organ allocation that explicitly incorporates the strategic behavior of transplant centers as "big players" in the system. Within this framework, our key innovation is the introduction of a *priority adherence index*, which quantifies the fundamental tradeoff between priority adherence and allocation value across different market structures and allocation rules. Our analysis reveals three structural forces, *pool thinning*, *batch shortening*, and *strategic delay*, that systematically determine when and why priority is violated.

As the *New York Times* recently emphasized, "For decades, fairness has been the guiding principle of the American organ transplant system ... meant to ensure that donated organs are offered to the patients who need them most, in careful order of priority." Our framework provides the first formal lens for understanding when and why this principle is upheld or violated in practice. Looking forward, important extensions include incorporating OPO incentives (following Chan and Roth, 2024) and applying our framework to evaluate ongoing policy innovations in the transplant system.

A Proofs for Section 3

Proof of Proposition 3.1. We first derive a different expression for the priority adherence index by summing over all $i \in \{1, \ldots, n\}$ such that i is both the first positive-valued patient and the highest-valued patient, and dividing this sum by the probability that the organ is allocated. We also explicitly write the index as a function of $F(0)$ and n:

$$\begin{aligned}
\text{Index}(F(0), n) &= \Pr[\, i_{\text{pos}} = i_{\max} \leqslant n \,|\, x_{i_{\max}} > 0\,] = \frac{\sum_{i=1}^{n} \Pr[\, i_{\text{pos}} = i \text{ and } i_{\max} = i\,]}{\Pr[\, x_{i_{\max}} > 0\,]} \\
&= \frac{\sum_{i=1}^{n} F(0)^{i-1} \int_0^1 F(s)^{n-i} f(s) \, \mathrm{d}s}{1 - F(0)^n} = \frac{\sum_{i=1}^{n} \frac{F(0)^{i-1} - F(0)^n}{n-i+1}}{1 - F(0)^n} \\
&= \sum_{i=1}^{n} \frac{F(0)^{i-1} - F(0)^n}{(n-i+1)(1 - F(0)^n)}.
\end{aligned}$$

Taking the partial derivative of Index(F(0),n) with respect to $F(0)$, we obtain that

$$\begin{aligned}
\frac{\partial \text{Index}(F(0), n)}{\partial F(0)} &= \sum_{i=1}^{n} \frac{(i-1)F(0)^i + F(0)^n \left(F(0)^i(-i+n+1) - F(0)n\right)}{F(0)^2 \left(F(0)^n - 1\right)^2 (-i+n+1)} \\
&= \frac{F(0)^{n-1}}{\left(1 - F(0)^n\right)^2} \sum_{i=1}^{n} \left(F(0)^{i-1} \left(\frac{(i-1)F(0)^{-n}}{-i+n+1} + 1 \right) + \frac{n}{i-n-1} \right) \qquad (6)
\end{aligned}$$

Let

$$g(F(0), n, i) = F(0)^{i-1} \left(\frac{(i-1)F(0)^{-n}}{-i+n+1} + 1 \right) + \frac{n}{i-n-1}.$$

It is readily verified that $\lim_{F(0)\to 1} g(F(0), n, i) = 0$. The partial derivative of $g(F(0), n, i)$ with respect to $F(0)$ is:

$$\frac{\partial g(F(0), n, i)}{\partial F(0)} = (1-i)\left(1 - F(0)^n\right) F(0)^{i-n-2},$$

which equals zero if $i = 1$, and strictly negative for $i \geqslant 2$. Hence, $g(F(0), n, i) = 0$ if $i = 1$, and strictly positive for $i \geqslant 2$. This implies that (6) is strictly positive, so Index$(F(0), n)$

strictly increases in $F(0)$.

We next show that $\text{Index}(F(0), n)$ strictly decreases in n. Let $p = F(0)$ and $j = n-i+1$. When i goes from 1 to n, j goes from n to 1. So:

$$\text{Index}(p,n) = \sum_{i=1}^{n} \frac{p^{i-1} - p^n}{(n-i+1)(1-p^n)} = \sum_{j=1}^{n} \frac{p^{n-j} - p^n}{j(1-p^n)} = \frac{p^n}{1-p^n} \sum_{j=1}^{n} \frac{1-p^j}{jp^j}.$$

Thus,

$$\begin{aligned}
\text{Index}(p,n+1) - \text{Index}(p,n) &= \frac{p^{n+1}}{1-p^{n+1}} \sum_{j=1}^{n+1} \frac{1-p^j}{jp^j} - \frac{p^n}{1-p^n} \sum_{j=1}^{n} \frac{1-p^j}{jp^j} \\
&= \frac{p^{n+1}}{1-p^{n+1}} \frac{1-p^{n+1}}{(n+1)p^{n+1}} + \left(\frac{p^{n+1}}{1-p^{n+1}} - \frac{p^n}{1-p^n} \right) \sum_{j=1}^{n} \frac{1-p^j}{jp^j} \\
&= \frac{1}{n+1} - \frac{(1-p)p^n}{(1-p^{n+1})(1-p^n)} \sum_{j=1}^{n} \frac{1-p^j}{jp^j}.
\end{aligned}$$

Hence, $\text{Index}(p,n+1) - \text{Index}(p,n) < 0$ is equivalent to:

$$H(p) := \sum_{j=1}^{n} \frac{1-p^j}{jp^j} - \frac{(1-p^{n+1})(1-p^n)}{(1+n)(1-p)p^n} > 0.$$

It is readily verified that $\lim_{p \to 1} H(p) = 0$, so in order to show that $H(p) < 0$ for $p \in (0,1)$ we only need to show that $H'(p) < 0$ for $p \in (0,1)$:

$$\begin{aligned}
H'(p) &= \sum_{j=1}^{n} (-p^{-j-1}) + \frac{n(1-p)\left(1-p^{2n+1}\right) - p\left(1-p^n\right)^2}{(n+1)(1-p)^2 p^{n+1}} \\
&= \frac{\left(1-p^{n+1}\right)\left[\left(1+n\left(1-p\right)\right)p^n - 1\right]}{(n+1)(1-p)^2 p^{n+1}} < 0 \\
&\iff 1 - (1+n(1-p))p^n > 0 \\
&\iff \frac{1-p^n}{1-p} > np^n \iff \sum_{i=0}^{n-1} p^i > np^n.
\end{aligned}$$

This last inequality is true since $p^i > p^n$ for any $i \in \{0, 1, ..., n-1\}$, $p \in (0, 1)$, and $n \geqslant 2$.

The limiting results $\lim_{F(0)\to 0} \text{Index}(F(0), n) = 1/n$ and $\lim_{F(0)\to 1} \text{Index}(F(0), n) = 1$ follow directly from the expression of $\text{Index}(F(0), n)$. □

B Proofs for Section 4

Proof of Theorem 4.1. We first show that the characterization indeed constitutes an equilibrium, and then prove its uniqueness.

Suppose the affiliation c induces a total of z batches, with batch sizes $n_1, n_2, \ldots, n_z \geqslant 1$, ordered from earliest to latest, such that $\sum_{j=1}^{z} n_j = n$. We begin with the last batch of patients, managed by center $c(n)$.

1. When the game reaches history h_n, center $c(n)$ strictly prefers to accept rather than reject, since acceptance yields a payoff of $x_n > 0$, while rejection yields zero.

2. If the last batch contains only one patient (i.e., $n_z = 1$), the analysis for the last batch is complete. Otherwise, suppose the last batch contains multiple patients. Consider the game at h_{n-1}. If center $c(n)$ accepts, it receives x_{n-1}; if it rejects, the game proceeds to h_n, in which case it receives x_n. Thus, $c(n)$ prefers to accept if $x_{n-1} \geqslant x_n$, and to reject otherwise.

3. This logic extends recursively. Suppose the game reaches h_i for some i in the last batch. If $c(n)$ accepts, it receives x_i; if it rejects, the game continues and yields a payoff of $\max_{i+1 \leqslant g \leqslant n} x_g$ to center $c(n)$. Therefore, center $c(n)$ prefers to accept if $x_i \geqslant \max_{i+1 \leqslant g \leqslant n} x_g$, and to reject otherwise.

This completes the analysis of the last batch: center $c(n)$ accepts for patient i if x_i is the highest value among those weakly after i within the last batch; otherwise, it rejects.

Next, consider the second-to-last batch. When the game reaches $h_{\sum_{j=1}^{z-1} n_j}$—the last patient in that batch—center $c\left(\sum_{j=1}^{z-1} n_j\right)$ strictly prefers to accept, since $x_{\sum_{j=1}^{z-1} n_j} > 0$, and rejecting passes the organ to a different center, which will accept it for one of its own patients—resulting in a payoff of zero for center $c\left(\sum_{j=1}^{z-1} n_j\right)$. By the same logic as in the last batch, center $c\left(\sum_{j=1}^{z-1} n_j\right)$ accepts for patient i if x_i is the highest value among those weakly after i within the second-to-last batch; otherwise, it rejects.

Repeating this argument for each of the earlier batches completes the argument that the characterization is an equilibrium.

We next prove uniqueness using an induction argument.[5] We show that for any $n = 1, 2, \ldots$, the characterization describes the unique equilibrium. The claim obviously holds for $n = 1$. Suppose it holds for $n = m \geqslant 1$; we now show that it also holds for $n = m + 1$. We divide into two cases depending on whether patients 1 and 2 belong to the same center.

1. Patient 1 belongs to a different center from patient 2. Then $c(1)$ will accept for patient 1. Accepting gives $c(1)$ a payoff of x_1, while rejecting passes the organ to $c(2)$. With m patients remaining, the unique outcome is that the organ never circles back to $c(1)$ by the induction hypothesis.

2. Patient 1 belongs to the same center as patient 2. Suppose the first batch includes patients 1 through $g \geqslant 2$. If $c(1)$ accepts for patient 1, its payoff is x_1. If $c(1)$ rejects for patient 1, then with m patients remaining, the unique outcome is that $c(1)$ receives a payoff of $\max_{2\leqslant i\leqslant g} x_i$. Thus, $c(1)$ prefers to accept for patient 1 if and only if

$$x_1 \geqslant \max_{2\leqslant i\leqslant g} x_i,$$

which yields the same characterization of behavior as in our theorem.

Hence, the characterization is the unique equilibrium for $n = m + 1$.

[5] We thank Attila Ambrus for suggesting this argument.

$\square$

Proof of Proposition 4.1. Substituting $s_j = 1/k$ into (2), we obtain

$$\text{Index}(k, n, s) = k\left(\sum_{m=1}^{n-1} \frac{(k-1)\left(\frac{1}{k}\right)^{m+1}}{m} + \frac{\left(\frac{1}{k}\right)^n}{n}\right) = \sum_{m=1}^{n-1} \frac{(k-1)\,k^{-m}}{m} + \frac{k^{1-n}}{n}.$$

We next show that the right-hand side strictly increases in k for $k \in \mathbb{R}$ with $k \geqslant 1$. Let

$$H(k) = \sum_{m=1}^{n-1} \frac{(k-1)\,k^{-m}}{m} + \frac{k^{1-n}}{n}.$$

Then

$$H'(k) = \sum_{m=1}^{n-1} \frac{k^{-m-1}\big(m(1-k)+k\big)}{m} - \frac{(n-1)k^{-n}}{n}.$$

Separate the $m = 1$ term:

$$H'(k) = \frac{1}{k^2} + \sum_{m=2}^{n-1} \frac{k^{-m-1}\big(m(1-k)+k\big)}{m} - \frac{(n-1)k^{-n}}{n}.$$

Bound the remaining terms:

$$\sum_{m=2}^{n-1} \frac{k^{-m-1}\big(m(1-k)+k\big)}{m} \;\geqslant\; \sum_{m=2}^{n-1} \frac{k^{-m-1}\big(m(1-k)+0\big)}{m} = k^{-n} - \frac{1}{k^2}.$$

Therefore,

$$H'(k) \;\geqslant\; \frac{1}{k^2} + k^{-n} - \frac{1}{k^2} - \frac{(n-1)k^{-n}}{n} = \frac{k^{-n}}{n} \;>\; 0.$$

Hence, under uniform market shares $s_j = 1/k$, the index is strictly increasing in k for $k \geqslant 1$. Finally, the limiting result $\lim_{k\to\infty} \text{Index} = 1$ holds because, as $k \to \infty$, all the terms except $(k-1)/k$ converge to zero.

Substituting $s_j = 1/k$ into (3), and letting $x(m)$ denote the expected maximum of m draws and $V(k)$ the expected allocation value as a function of k, we can write this expected value as:

$$V(k) := k\left(\sum_{m=1}^{n-1}\frac{1}{k^m}\left(1-\frac{1}{k}\right)x(m)+\frac{1}{k^n}x(n)\right) = \sum_{m=1}^{n-1}\left(k^{1-m}-k^{-m}\right)x(m)+k^{1-n}x(n).$$

Then, the derivative $V'(k)$ is given by:

$$V'(k) = \sum_{m=1}^{n-1} k^{-1-m}(k+m-km)x(m)+k^{-n}(1-n)x(n).$$

The inequality $k+m-km \leqslant 0$ is true for all integers $k \geqslant 2$ and $m \geqslant 2$ and strictly so if either $k > 2$ or $m > 2$. Also, $x(m)$ is increasing in m. We thus have the following inequality:

$$\begin{aligned} V'(k) &< \sum_{m=1}^{n-1} k^{-1-m}(k+m-km)x(1)+k^{-n}(1-n)x(1) \\ &= \left(\sum_{m=1}^{n-1} k^{-1-m}(k+m-km)+k^{-n}(1-n)\right)x(1) = 0. \end{aligned}$$

Hence, the expected value $V(k)$ decreases in k. The limiting result $\lim_{k\to\infty} V(k) = x(1)$ holds because, as $k \to \infty$, all the terms except $(1-1/k)x(1)$ converge to zero. □

Proof of Theorem 4.2. By the strategic equivalence between the priority-proposal rule and the sequential-offer rule (established in Theorem 5.1), we first characterize the Bayesian Nash Equilibrium for the priority-proposal rule. Fix affiliation c. Let α_i be the probability that center $c(i)$ proposes patient i. Then, from center $c(i)$'s perspective, the probability that the opponent center proposes a patient before patient i is:

$$\sum_{i'<i, c(i')\neq c(i)} \alpha_{i'}.$$

Define $\delta_i := 1-\sum_{i'<i,c(i')\neq c(i)} \alpha_{i'}$. Hence, center $c(i)$'s expected payoff from proposing patient i is $\delta_i x_i$. Center $c(i)$ proposes patient i if and only if $x_i > 0$ and $\delta_i x_i \geqslant \delta_{i'} x_{i'}$ for all i' with

$c(i') = c(i)$. This pins down α_i to be:

$$\alpha_i = \Pr[\, x_i > 0 \text{ and } \delta_i x_i \geqslant \delta_{i'} x_{i'} \ \ \forall\, i' \text{ with } c(i') = c(i)\,]\,.$$

This completes the characterization of the BNE for the priority-proposal rule.

For the sequential-offer rule, we specify a center's choice at every decision history. When the offer reaches patient i, center $c(i)$ accepts patient i if $x_i > 0$ and

$$x_i \geqslant \left(1 - \frac{\sum_{i<i'<i'',c(i')\neq c(i)} \alpha_{i'}}{\delta_i}\right) x_{i''}, \text{ for all } i'' > i \text{ with } c(i'') = c(i).$$

Substituting $\delta_i = 1 - \sum_{i'<i,c(i')\neq c(i)} \alpha_{i'}$ into this expression yields:

$$\delta_i x_i \geqslant \left(\delta_i - \sum_{i<i'<i'',c(i')\neq c(i)} \alpha_{i'}\right) x_{i''} = \delta_{i''} x_{i''}, \text{ for all } i'' > i \text{ with } c(i'') = c(i).$$

This completes the proof. □

C Proofs for Section 5

Proof of Theorem 5.1. Under the *priority-proposal rule*, center j with type $\theta_j = \{(i, x_i) : i \in I_j\}$ simultaneously proposes a patient from the action space $I_j \cup \{\emptyset\}$, where $\emptyset$ indicates that center j proposes no patient. Given the proposal profile $(b_1, \ldots, b_k)$, the organ is allocated to the proposed patient with the highest priority.

Under the *sequential-offer rule*, a pure strategy for center j specifies whether to accept or reject at each decision history $h \in H_j(\theta_j)$. However, we can partition center j's strategy space into outcome-equivalent classes based on which patient in $I_j \cup \emptyset$ is the first for whom center j accepts. Each such equivalence class can be represented by a single element $p_j \in I_j \cup \{\emptyset\}$, where p_j is the first patient that center j would accept, and $\emptyset$ represents the strategy of rejecting for all patients. Given the first-patient-to-accept profile $(p_1, \ldots, p_k)$, the organ is allocated to the patient with the highest priority among $(p_1, \ldots, p_k)$.

Therefore, both rules have identical reduced-form strategy spaces: each center j effectively chooses an element from $I_j \cup \{\emptyset\}$.

Given any strategy profile $(s_1, \ldots, s_k)$ where $s_j \in I_j \cup \{\emptyset\}$, both rules allocate the organ to the patient with highest priority among $(s_1, \ldots, s_k)$. Since payoffs depend only on which patient receives the organ, both rules induce identical payoff functions over strategy profiles.

Since the two rules have identical strategy spaces and identical payoff functions over strategy profiles, they are strategically equivalent. □

References

Abdulkadiroglu, Atila, Yeon-Koo Che, Parag A Pathak, Alvin E Roth, and Olivier Tercieux (2020), "Efficiency, justified envy, and incentives in priority-based matching." *American Economic Review: Insights*, 2, 425–442.

Abdulkadiroglu, Atila and Aram Grigoryan (2021), "Efficient and envy minimal matching." *Working Paper*.

Abdulkadiroglu, Atila and Tayfun Sönmez (2003), "School choice: A mechanism design approach." *American economic review*, 93, 729–747.

Agarwal, Nikhil and Eric Budish (2021), "Market design." In *Handbook of Industrial Organization*, volume 5, 1–79, Elsevier.

Aghion, Philippe and Jean Tirole (1997), "Formal and real authority in organizations." *Journal of political economy*, 105, 1–29.

Chan, Alex and Alvin E Roth (2024), "Regulation of organ transplantation and procurement: A market-design lab experiment." *Journal of Political Economy*, 132, 3827–3866.

CMS (2007), "Medicare program; hospital conditions of participation: requirements for approval and re-approval of transplant centers to perform organ transplants. final rule." *Federal Register*, 72, 15197–15280.

Committee on Organ Procurement and Transplantation Policy (2000), *Organ procurement and transplantation: assessing current policies and the potential impact of the DHHS final rule*. National Academies Press.

Fudenberg, Drew and Jean Tirole (1991), *Game Theory*. MIT Press, Cambridge, MA.

Henson, Jacqueline B, Xuyang Xia, Ryan McDevitt, Norine W Chan, Andrew S Barbas, Lindsay Y King, Andrew J Muir, Debra L Sudan, Stuart J Knechtle, and Lisa M McElroy

(2026), "Evolution of out-of-sequence liver allocation, 2019-2024." *American Journal of Transplantation.*

Husain, Syed Ali, Jordan A Rubenstein, Seshma Ramsawak, Anne M Huml, Miko E Yu, Lindsey M Maclay, Jesse D Schold, and Sumit Mohan (2025), "Patient and provider attitudes toward patient-facing kidney organ offer reporting." *Kidney International Reports*, 10, 1122–1130.

King, Kristen L, S Ali Husain, Miko Yu, Joel T Adler, Jesse Schold, and Sumit Mohan (2023), "Characterization of transplant center decisions to allocate kidneys to candidates with lower waiting list priority." *JAMA network open*, 6, e2316936–e2316936.

Mankowski, Michal A, Martin Kosztowski, Subramanian Raghavan, Jacqueline M Garonzik-Wang, David Axelrod, Dorry L Segev, and Sommer E Gentry (2019), "Accelerating kidney allocation: Simultaneously expiring offers." *American journal of transplantation*, 19, 3071–3078.

Munoz-Rodriguez, Edwin and James Schummer (2025), "Rationing through classification." *Available at SSRN 5225290.*

Ng, Han (2025), "Federal oversight and strategic choices of kidney transplant centers." *Working Paper.*

NYT (2025), "Organ transplant system 'in chaos' as waiting lists are ignored." *The New York Times.*