Formal and Real Organ Allocation

Yingni Guo Xuyang Xia*

June 18, 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. We further 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 allocation, strategic intermediaries, market design, market power, justified envy, Dutch auction, first-price 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 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 first 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

*We thank Kuichen Wang for excellent research assistance. Guo: Department of Economics, Duke University; email: yingni.guo@duke.edu. Xia: Department of Economics, Duke University; email: xuyang.xia@duke.edu.

deceased-donor organ must be listed with and managed by a transplant center, which is typically a hospital division. Organ offers are routed to such centers rather than to 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). Importantly, centers possess private, real-time information about their own patients' clinical conditions, preferences, and logistical constraints that would be practically impossible for a central designer to monitor continuously.

Transplant centers face different incentives than individual patients do. 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; New York Times, 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? Third, what does our model reveal when we 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 it to patient $i + 1$. 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, the highest-priority patient among those with positive values does not receive the organ. Having positive value and higher priority than the recipient, this patient has *justified envy*. 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 this 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 to later patients in the batch without competitive pressure from other centers, 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 offer to that patient 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 with which patient 1 is affiliated 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 center 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 offer 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 to a rival 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-value 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 reduces priority adherence. We show that the interaction of these two forces can cause the priority adherence index to be nonmonotonic 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: a transplant must occur within hours after the organ is recovered. 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 percent to 19 percent. But this expedient comes at a cost: open offers circumvent the priority list, raising concerns about fairness (New York Times, 2025). This motivates interest in alternative market designs that would preserve speed while increasing 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 priority-proposal 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 for OPOs to use as an alternative market design.

Finally, we ask whether the priority-proposal rule, or equivalently the sequential-offer rule, is the best a designer could do. Casting the problem as one of mechanism design without transfers, we show that it is not: because both rules let centers communicate only coarse information, a mechanism that elicits centers' full private information can raise priority adherence and expected allocation value at the same time. We view this as a first step toward characterizing the frontier of implementable outcomes, and we leave a full treatment to future work.

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 have 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 organ transplant market—a framework that is amenable to analyzing alternative designs and how market outcomes respond to policy reforms.

We also contribute to the literature on how transplant centers respond to allocation incentives. 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 choose for them, and show that greater market power can improve welfare by better targeting resources. As we do in our paper, they model centers that decide on patients' behalf and study how market structure shapes equilibrium outcomes. However, while they—and related empirical work (Sweat, 2025)—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 both 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.

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), since it measures the probability that this property holds. Reducing 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, the top trading cycles mechanism 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 the relevant distribution is determined by the market structure and the match value distribution rather than by a uniform prior.

We contribute to the literature on mechanism design without transfers under multidimensional private information. Existing work studies how (i) using allocation probabilities as a numeraire, (ii) linking decisions, or (iii) imposing quota-like delegation rules can elicit information or align incentives without transfers (Jackson and Sonnenschein, 2007; Miralles, 2012; Frankel, 2014; Ball, Jackson and Kattwinkel, 2022). Our setting is distinct: transplant centers act as strategic intermediaries, privately observe match values for their own patients, and make acceptance decisions under a priority-based rule. We focus on the resulting wedge between formal and real allocation, and on how market structure and allocation procedures shape both priority adherence and expected allocation value.

The rest of the paper proceeds as follows. Section 2 introduces the model and the main outcome measures: the priority adherence index and expected allocation value. Section 3 analyzes the monopoly benchmark. Section 4 studies multicenter markets and identifies the forces of batch shortening, pool thinning, and strategic delay. Section 5 introduces the priority-proposal rule, establishes its strategic equivalence to the current sequential-offer rule, and then formulates the broader mechanism-design problem. Section 6 connects our findings to related empirical evidence. Section 7 concludes with the broader interpretation of our analysis, returning to the role of transplant centers as informed intermediaries and underscoring the fact that imperfect priority adherence reflects an allocation-design tradeoff.

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 map $c : \{1, \dots, n\} \to \{1, \dots, k\}$. Given the realization of this map, 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 map c may still be imbalanced due to randomness.[1]

Allocation mechanism. The OPO allocates the organ using a *sequential-offer rule* in which 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 the organ 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.

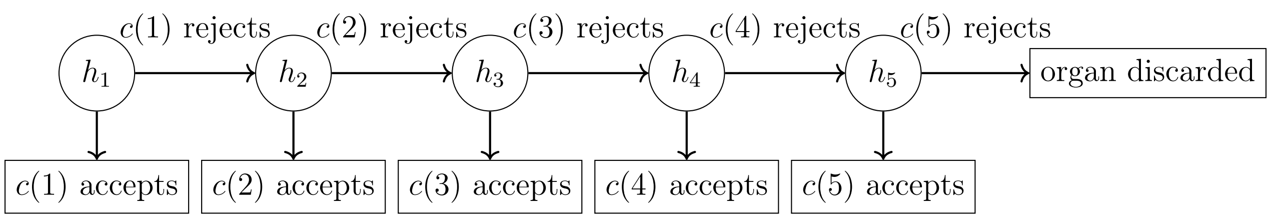

Figure 1: Sequential-offer rule for $n = 5$ patients.

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

1. Nature draws a value vector $x = (x_1, \dots, x_n)$ and an affiliation map $c : \{1, \dots, n\} \to \{1, \dots, k\}$.

2. Each center j privately observes its *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 whether 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 map and its realization, whenever the distinction is clear from context.

Example 1. Consider $n = 5$ patients and $k = 2$ centers, with the realized affiliation map 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 the affiliation map 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 both of the following hold:

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 $S \subseteq \Theta_{-j}$,

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

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 a rival center rejects unexpectedly (an event occurring only when a universal organ is involved), 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 hypothetical benchmark in which each patient i observes their match value x_i and decides whether to accept the organ. In reality, as in our model, patients cannot evaluate match quality, and centers decide on their behalf.

Patients with weakly negative values reject the offer, while those with positive values accept it.[2] 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 lighten notation, we sometimes omit the dependence of i_{pos} on x when no confusion arises, and similarly 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 who has 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 $x = (x_1, \ldots, x_n)$.

[2] We assume that patient i's payoff from receiving the organ equals x_i. This need not be the case; as long as patient i's payoff from receiving the organ has the same sign as x_i, the same decision rule holds.

The *priority adherence index* is a probabilistic generalization of *no justified envy* from the school-choice literature (Abdulkadiroglu and Sönmez, 2003). In that setting, an assignment is said to have no justified envy if no student prefers another student's seat to their own while having higher priority for that seat. 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).

We define the *expected allocation value* as the expected value of the recipient, conditional on the organ being allocated. Under the patient-decision benchmark, the priority adherence index equals one. The expected allocation value is the conditional mean of x_i given $x_i > 0$:

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

2.2 Independence of Match Values and Priority Ranking

In our model, we assume that the match value x_i is independent of the patient's priority ranking i. In practice, the relationship between ranking and value is weak and varies by organ type.

First, transplant outcomes depend heavily on organ-recipient match qualities (e.g., donor-recipient ages and weights), operation complexities, and other factors that are not captured by the priority ranking. For instance, in some scenarios, a low-quality organ may have better outcomes on a less urgent, younger patient with better resilience, yet in other cases, a very urgent patient may benefit more from the transplant even though the organ quality is not ideal. These practical, idiosyncratic, and often stochastic factors in real-life implementations generate considerable noise that separates the priority ranking and the expected outcomes.

There are also mechanical reasons created by the allocation rule that weaken the relationship between priority and outcome. For non-kidney organs, higher priority is generally associated with lower allocation value, though substantial noise remains. Candidates for non-kidney organs are primarily ranked by medical urgency (i.e., how soon they would die without a transplant) and by geographical distance from the donor. For example, patients waiting for a deceased-donor liver transplant are ranked by *urgency score (MELD) bin* $\times$ *distance bin*:[3] patients with MELD$\geqslant$ 37 within 150 nautical miles (NM) of the donor organ are ranked before patients with MELD$\geqslant$ 37 within 150–250 NM, who are then ranked before

[3]MELD stands for Model for End-Stage Liver Disease. The score ranges from 6 to 40; a higher MELD score indicates greater medical urgency, i.e., a higher probability of dying within 90 days without a transplant.

patients with MELD$\geqslant$ 37 within 250–500 NM, followed by patients with 37 > MELD $\geqslant$ 33 within 150 NM, and so on. Within each group, patients are further ranked by blood type, MELD, wait time, and other criteria (Organ Procurement and Transplantation Network, 2026). While greater urgency is associated with a higher post-transplant mortality rate (Kwong et al., 2025), the association between distance and urgency depends on organ availability and the spatial distribution of compatible patients at a given moment, both of which introduce considerable noise.

In contrast, for kidneys, the availability of dialysis allows the ranking to prioritize candidates with better expected transplant outcomes and longer waiting time. The former directly links high priority to high allocation value, whereas the latter provides dynamic incentives for high-priority patients to wait for a better organ (Agarwal et al., 2018)—in other words, linking a high ranking to a high value of the outside option. In our static setting, x_i can be interpreted as the difference between the transplant payoff and the value of the outside option. Hence, dynamic incentives in kidney allocation weaken the association between priority ranking and allocation value.

We therefore regard the assumption that x_i is independent of patient rank as a parsimonious baseline specification that captures this empirical orthogonality. This benchmark is far from a knife-edge case; our core qualitative insights remain robust to systematic correlations between ranking and value. By shutting down this channel initially, our model isolates the pure strategic effects of intermediation, leaving the formal parameterization of rank-dependent distributions to future extensions.

3 Monopoly Center ($k = 1$)

We now turn to our main analysis, beginning with the polar case of a monopoly center that observes all match values and makes decisions for all patients. This setting highlights how information and discretion shape allocation in the absence of competition.[4]

Let $i_{\max}(x) \in \arg\max_i x_i$ denote the patient with the highest value. If $x_{i_{\max}} > 0$, 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 $x_{i_{\max}}$ is given by $[F(x)]^n$. Therefore, the expected allocation value is:

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

[4]The monopoly setting is not hypothetical; in some isolated regions, a single center effectively dominates the top of the priority list (King et al., 2023).

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 patient $i_{\max}$. 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 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 that exactly $m \in \{1, \ldots, n\}$ 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:

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

This calculation shows that the index depends on distribution F only through $F(0)$. In the monopoly setting, 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.

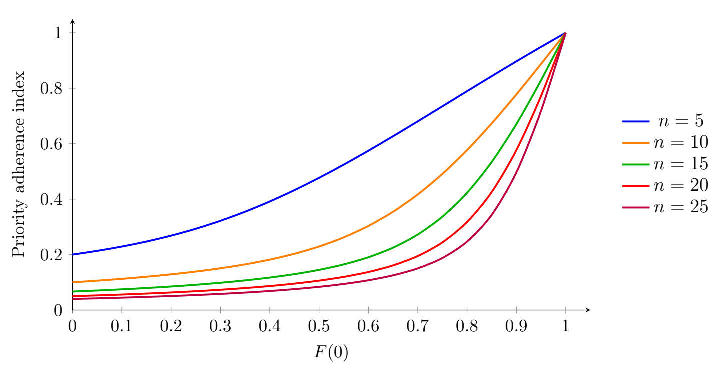

Figure 2: Effect of varying the number of patients n or $F(0)$ on the index.

Figure 2 illustrates the fact 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.

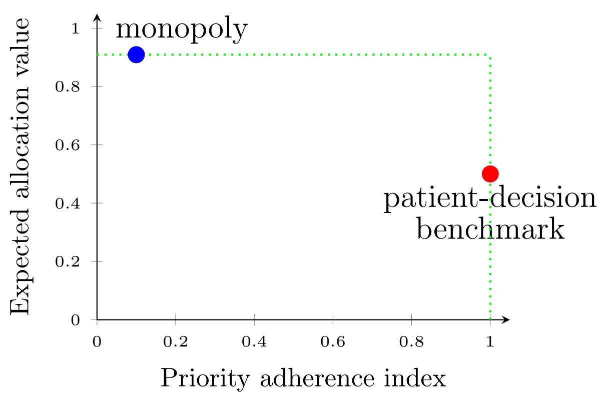

Figure 3: Monopoly versus patient-decision benchmark.

4 Multiple Centers ($k \geqslant 2$)

We now turn to the multicenter setting. Each patient is independently affiliated with a center according to the centers' 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 map is trivial, here c shapes how the organ is allocated.

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

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

Figure 4 illustrates the i-batch under a given affiliation map 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.

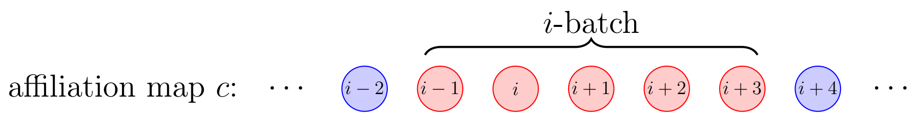

Figure 4: The i-batch under a given affiliation map c; colors denote patients' centers.

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 affiliated center optimizes locally: if it accepts for a 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 accepting for a patient in the current batch to passing the offer to the next center. Because the next center will accept for a patient in its own batch, an offer passed to that center never returns. Together, these two features ensure that 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:

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

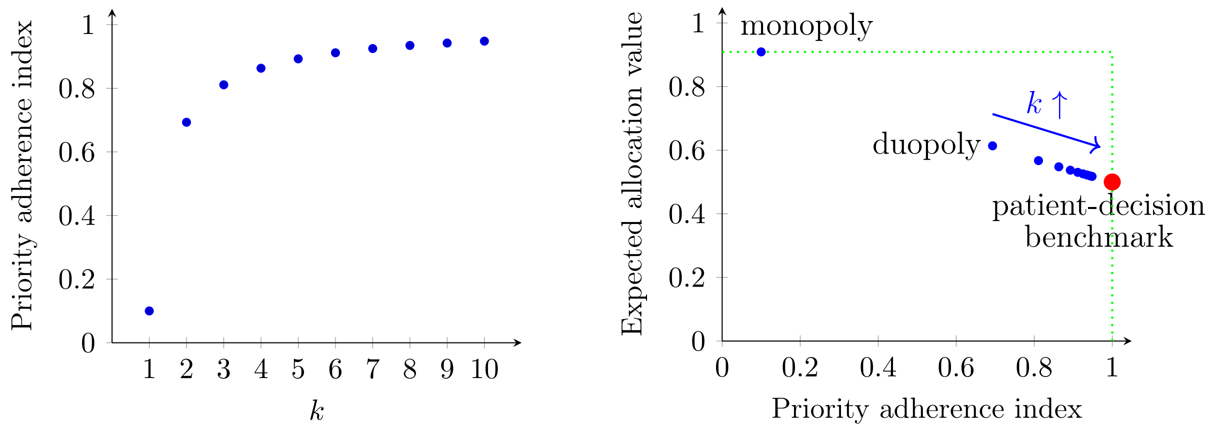

Figure 5: Effect of increasing the number of centers k: $n = 10,\ x_i \sim \text{Unif}[0, 1]$.

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 adherence 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 steeply while the expected allocation value drops substantially. This sharp shift reflects a structural change: introducing a second center creates a positive chance that the first batch ends before all patients have been reached. Beyond that, both measures become less sensitive to which center interrupts the first batch, 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, which accounts 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 summarized 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 with $F(0)$.

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.

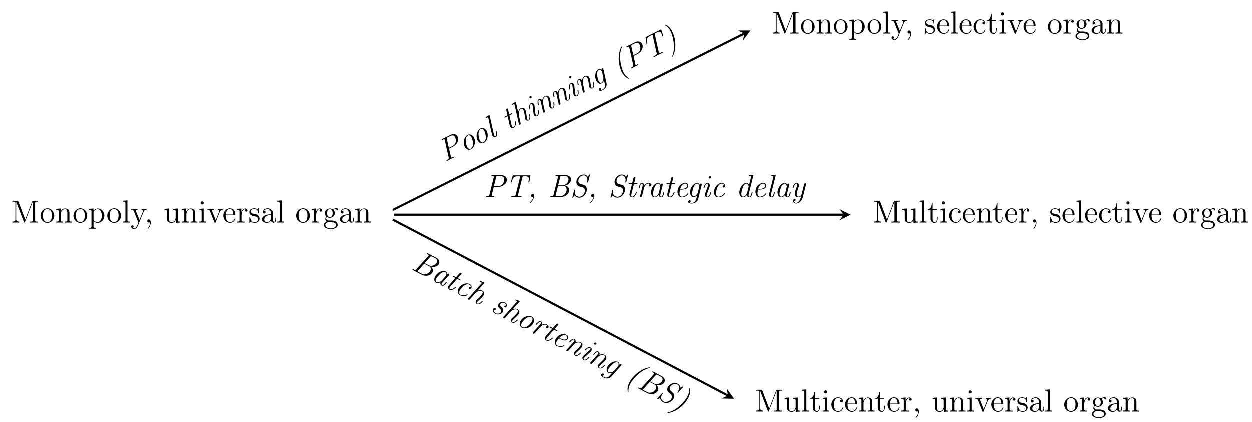

Figure 6: Structural forces shaping priority adherence.

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 a payoff of 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. The corresponding equilibrium characterization for the general case with $k \geqslant 2$ centers is presented in Appendix D.

Theorem 4.2 (Equilibrium for Selective Organs and Two Centers)**.** *Assume $F(0) \in (0,1)$ and $k = 2$. Fix an affiliation map c. An equilibrium $\{\alpha_i, \delta_i\}_{i=1}^n$ satisfies both (4) and (5) below:*

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

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

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 rival 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 rival 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 general case $k \geqslant 2$ admits an analogous characterization (see Theorem D.1 in Appendix D), with the key difference being that reach probabilities are averaged over all possible ways the remaining patients are distributed across rival centers.

The next example solves (4)–(5) to compute $\{\alpha_i, \delta_i\}_{i=1}^n$ for a given affiliation map 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 map 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 rival 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 of Table 4.

The right table in Table 4 shows one realization of match 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 the question as to whether this force rivals pool thinning 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 rival 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 map $c = (c(1), c(2), c(3))$ 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 managing 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 managing 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 percent chance of falling into the lower cluster (near 1/2) and a 50 percent 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 managing patients 1 and 3 bypasses its first positive-valued patient. Accordingly, Figure 7 shows a discontinuous drop in the index at $F(0) = 1/2$. □

5 Alternative Allocation Rules

5.1 The Priority-Proposal Rule

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 five hours after organ recovery, they invited favored hospitals to each propose one patient; the organ was then allocated to the highest-priority patient among those proposed (New York Times, 2025).

Motivated by such practices, we formalize the *priority-proposal rule*: each center may

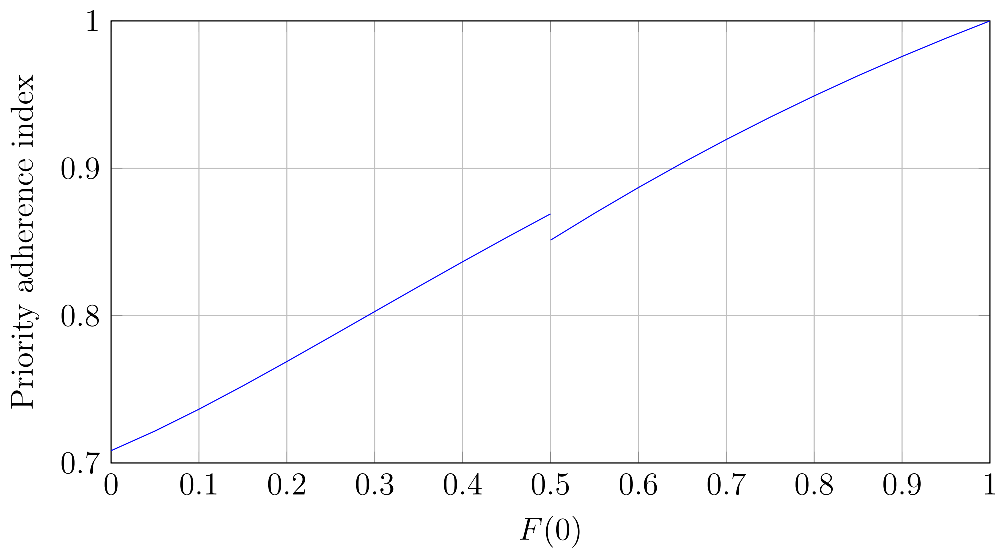

Figure 7: Nonmonotonicity of priority adherence in $F(0)$. Pool thinning tends to raise the index as the set of positive-valued patients shrinks, but strategic delay can create nonmonotone patterns.

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 the *Timing and information* paragraph of 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.

5.1.1 An Equivalence Result 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. Proposing a patient under the priority-proposal rule is equivalent to specifying that patient as the first patient the center would accept under the sequential-offer rule. Hence, for every strategy in the priority-proposal rule, there is an equivalent strategy in the sequential-offer rule, and vice versa.

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.

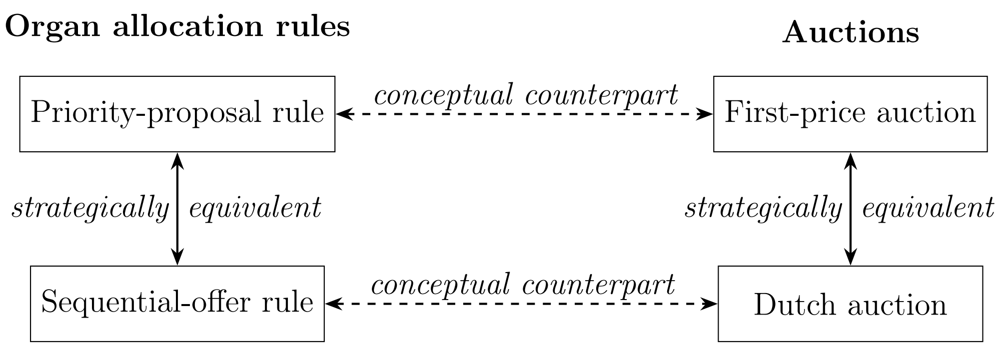

Figure 8: Strategic and conceptual relationships among four mechanisms.

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 a 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 of its patients and their positions on the priority list. Although each center j is restricted to proposing only one of its own patients, we can equivalently enlarge the action space to allow proposals of any patient $i \in \{1, \dots, n\}$. The payoff from winning with proposal i is x_i if i is affiliated with center 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 type θ_j. These features make the problem less tractable than standard auctions, which is why we have developed tools tailored to priority-based systems rather than relying on familiar auction methods.

5.1.2 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 and vice versa. Second, because the two rules are outcome-equivalent, choosing between them hinges on practical considerations outside our model.

One major challenge in organ allocation is time: transplants must occur within hours after an organ is recovered.[5] This constraint makes the sequential-offer rule costly: waiting for a response before sending the next offer can be time-consuming. In response, many OPOs have already adopted practices to accelerate placement, such as allowing centers to make tentative decisions for lower-ranked patients (Guan et al., 2025) or using open offers to place organs out of sequence (Henson et al., 2026). Recent rule changes have exacerbated this time pressure by broadening organ sharing and requiring OPOs to coordinate with more centers nationwide, further straining the sequential-offer rule. OPOs have cited these changes to justify relying more on open offers (New York Times, 2025).

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, it also introduces an effort cost: a center may invest effort in identifying a candidate, only to lose the organ to another center that proposes a higher-priority patient. Even so, the more pressing concern is that out-of-sequence allocations have risen from 2 percent to 19 percent in recent years, an increase often justified by time sensitivity. This trend has contributed to growing public distrust in the transplant system (New York Times, 2025). Given these time and fairness concerns, we recommend 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.[6] 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 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 (Mankowski et al., 2019; Committee on Organ Procurement and Transplantation Policy, 2000). 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 analogy to a Dutch auction.

5.2 Mechanism Design Approach

There are many other possible mechanisms for organ allocation to consider. We have argued that the priority-proposal rule is analogous to a first-price auction. One natural analog of a second-price auction works as follows: (i) each center may propose one of its own patients;

[5] Preservation times vary by organ type: kidneys can typically be transplanted within 36 hours, the liver, pancreas, and intestinal organs within 12 hours, and hearts and lungs within 6 hours (Health Resources and Services Administration, 2025).

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

(ii) the organ goes to the center whose proposed patient has the highest priority among all proposed patients; (iii) that center may then allocate the organ to any of its patients, provided the chosen patient's priority exceeds the second-highest priority among those proposed. Just as bidding one's true value is weakly dominant in a second-price auction, each center's weakly dominant strategy here is to propose its highest-priority patient with a positive match value.

This mechanism, however, raises a practical concern. When a winning center allocates the organ to a patient other than the one it proposed, the proposed patient is documented as viable but passed over in favor of a lower-priority patient. This paper trail may exacerbate ongoing public concerns about the fairness of the allocation system. The priority-proposal rule, by contrast, avoids this problem: the winning center must allocate the organ to the patient it proposed.

The discussion above and Section 5.1.2 together highlight various practical considerations in designing the organ allocation rule. Nevertheless, to understand the theoretical limits of what can be achieved by any mechanism, we now abstract from these concerns and formulate the mechanism design problem. A key question is whether the priority-proposal rule or, equivalently, the sequential-offer rule, lies on the frontier of implementable combinations of the priority adherence index and expected allocation value. We show in Example 4 that it need not: there is room to improve both objectives simultaneously.

One basic reason is that both rules allow centers to communicate only coarse information to the designer. Under the priority-proposal rule, each center proposes a single patient; under the sequential-offer rule, a center simply accepts or rejects when one of its patients is offered the organ. In contrast, a general mechanism elicits full private information from each center, which can enable the designer to simultaneously improve priority adherence and expected allocation value. Our framework provides a foundation for characterizing this frontier, and for exploring practical modifications to the current rule that improve both objectives. We leave both tasks for future research.

5.2.1 Mechanism Setup and Example

The designer's problem arises because they do not observe the match value of each patient. Instead, only patient i's center, denoted by $c(i)$, observes that patient's match value. The affiliation map c, however, is known to the designer. A mechanism is a family of maps $\rho = \{\rho_c\}_c$. For each realized c, the map $\rho_c : [-1, 1]^n \to \Delta^n$ assigns to each profile of reported values, $(\hat{x}_1, \ldots, \hat{x}_n)$, an allocation probability vector:

$$\rho_c(\hat{x}_1, \ldots, \hat{x}_n) = (p_1, \ldots, p_n, p_{n+1}).$$

Here, p_i is the probability that patient $i \in \{1, \ldots, n\}$ receives the organ; p_{n+1} is the probability of discarding the organ; and $\sum_{i=1}^{n+1} p_i = 1$. The game follows the same timing as in the *Timing and information* paragraph of Section 2, except that in step 3, each center j reports the match values of its affiliated patients to the designer, and the map ρ_c is implemented.

The following example uses this setup to show that the priority-proposal rule need not lie on the implementable frontier.

Example 4. Consider a setting with $n = 3$ patients and $k = 2$ centers. The match value x_i is 1 with probability 1/4, 1/3 with probability 1/4, and $-K$ with probability 1/2, where $K > 0$ is large.[7,8] The realized affiliation map is $c = (1, 2, 1)$, so center 1 has patients 1 and 3, and center 2 has patient 2.

We first compute the priority adherence index and expected allocation value under the priority-proposal rule. Center 2 proposes patient 2 if $x_2 > 0$ and proposes no patient if $x_2 < 0$. Center 1 proposes its highest-priority patient with a positive match value, except when $x_1 = 1/3$ and $x_3 = 1$. In that case, center 1 proposes patient 3, since with probability 1/2 center 2 proposes no patient and so patient 3 receives the organ. The event $\{x_1 = 1/3,\, x_3 = 1\}$ occurs with probability 1/16. Since the organ is allocated with total probability $1 - 1/8 = 7/8$, the priority adherence index is

$$\frac{7/8 - 1/16}{7/8} = \frac{13}{14} \approx 0.929.$$

The organ goes to the highest-priority patient with a positive match value except when $x_1 = 1/3$ and $x_3 = 1$; in that case, it goes to patient 2 if $x_2 > 0$, and to patient 3 otherwise. The expected allocation value is

$$\frac{59}{84} \approx 0.702.$$

In the case $x_1 = 1/3$ and $x_3 = 1$, incentive compatibility places a lower bound on the probability assigned to patient 3. To see this, suppose the mechanism assigns probability p_3 to patient 3 and probability $1 - p_3$ to patient 1. If center 1 reports truthfully, its expected payoff is

$$p_3 \cdot 1 + (1 - p_3) \cdot \frac{1}{3}.$$

If center 1 misreports x_1 as negative, patient 3 receives the organ with probability 1/2,

[7]For large K, we can ignore the incentive-compatibility constraints that prevent a center from misreporting a negative match value as positive; such misreporting can be deterred by allocating the organ to any reported positive-valued patient with a probability that converges to zero as $K \to \infty$. These constraints are also nonbinding under the sequential-offer and priority-proposal rules.

[8]While the match values here fall outside the $[-1, 1]$ interval assumed in the general model, this specific support is chosen purely to simplify the exposition of the example. Because the strategic incentives are invariant to affine transformations, all payoffs can be rescaled to fit the original $[-1, 1]$ domain without altering any qualitative results.

yielding an expected payoff of 1/2. Thus, truthful reporting requires

$$p_3 \cdot 1 + (1 - p_3) \cdot \frac{1}{3} \geqslant \frac{1}{2},$$

which gives $p_3 \geqslant 1/4$.

The index-maximizing mechanism therefore sets $p_3 = 1/4$ when $x_1 = 1/3$ and $x_3 = 1$, allocating the organ to patient 1 with probability 3/4 and to patient 3 with probability 1/4:

$$\rho_c(1/3, \hat{x}_2, 1) = (3/4, 0, 1/4, 0),$$

where the coordinates are ordered as (p_1, p_2, p_3, p_4), with p_4 denoting the probability of discarding the organ.[9] By the preceding incentive-compatibility calculation, center 1 has no incentive to misreport. The resulting priority adherence index and expected allocation value are

$$\left(\frac{55}{56}, \frac{19}{28}\right) \approx (0.982,\ 0.679).$$

More generally, any $p_3 \in [1/4, 1]$ yields an incentive-compatible mechanism, with priority adherence index and expected allocation value given by:

$$\left(\frac{14 - p_3}{14}, \frac{14 + p_3}{21}\right).$$

We refer to these mechanisms as *variants of the index-maximizing mechanism*, which are depicted by the blue line in Figure 9. As p_3 increases from 1/4 to 1, the combination moves from

$$\left(\frac{55}{56}, \frac{57}{84}\right) \approx (0.982,\ 0.679)$$

to

$$\left(\frac{13}{14}, \frac{5}{7}\right) \approx (0.929,\ 0.714),$$

covering a range of combinations that improve on the priority-proposal rule with respect to both the priority adherence index and expected allocation value.

□

[9] This index-maximizing mechanism resembles the optimal mechanism in the multiproject environment of Guo and Shmaya (2023), where the agent can hide projects. For the agent to be incentivized to propose all projects, he must receive the maximal payoff from proposing each project alone. Similarly here, center 1 can report either patient as having a negative value. To incentivize center 1 to report both patients as positive, it must be given the maximal payoff from proposing each patient as positive alone.

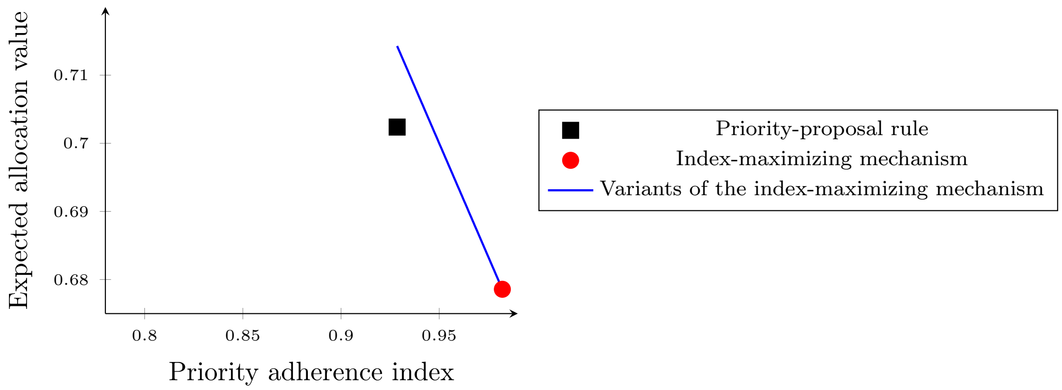

Figure 9: The priority-proposal rule is not on the implementable frontier.

6 Related Empirical Evidence

Our theory is motivated by empirical patterns documented in economic and medical studies, investigative reports, and news articles, and it provides a unified framework for interpreting these patterns.

To start with, evidence has suggested that transplant centers indeed have the incentive and discretion to optimize organ allocation internally. The first type of empirical evidence is through local monopoly centers. King et al. (2023) focus on eleven geographically isolated kidney transplant centers, and find that 68 percent of organs to these centers are not accepted by patients with the highest priority. This pattern cannot be fully explained by organ incompatibility ($x_i < 0$ in our model). Even among the highest-quality kidneys, only 44 percent were accepted by the highest-ranked patients; the authors believe this demonstrates that "transplant centers often overlook higher allocation priority candidates to place organs with the 'right' recipient" (King et al., 2023, p. 9). The authors also note that these local monopoly centers have "full discretion to decline offers for higher-priority candidates and accept them for lower-ranked candidates at their center" (King et al., 2023, p. 1). This underscores a key premise of our model: centers need not mechanically follow the priority ranking. Our Figure 3 speaks to monopoly markets and shows that such a market structure would indeed generate a low priority adherence index.

The second type of evidence on centers' incentives and discretion is from allocation out-of-sequence (AOOS). Local Organ Procurement Organizations (OPOs), which send organ offers, can sometimes bypass the priority ranking by making an offer to a selected center, which may then choose any of its patients to receive the organ. When AOOS happens, the preferred center effectively gets its patient batches connected and becomes a monopoly that can freely allocate the organ without the risk of losing it to rivals. Numerous studies have suggested that many high-priority patients are skipped in AOOS, and the organ is placed with a lower-ranked patient (e.g., Benkert et al., 2025; Liyanage et al., 2025; Henson et al., 2026), even though some high-priority patients may find the organ desirable (New York Times, 2025). In fact, the phenomenon of desirable offers being skipped is not restricted to monopoly markets or AOOS. Husain et al. (2019, p. 1) study a general set of kidney transplants and find that "a large number of deceased donor kidney offers are received by candidates but are declined on their behalf, resulting in what appears to be many missed opportunities for a transplant before death or removal from the waiting list."

There is also empirical evidence that speaks directly to the intensity of competition. Proposition 4.1 states that the expected allocation value decreases with the number of centers in the market. Using data from the national deceased liver transplant registry, Halldorson et al. (2013) find that, compared with donation service areas (DSAs) without competition (i.e., monopoly markets), DSAs with competition performed transplantation for patients with (i) a higher risk of graft failure or death and (ii) greater urgency, who are often ranked higher on the priority list. Moreover, these differences get larger as the competition (measured by center-level HHI) intensifies.

Multiple other studies have also found similar associations between higher local-market center density and worse transplant outcomes (Adler et al., 2014, 2015, 2016). This observation is consistent with our theory. Although deceased organs are sometimes offered to patients in other markets, the ranked list prioritizes patients in the local market for their geographical proximity, and thus the degree of competition for each organ is largely determined by the number of rival centers in the local market. Our model provides a theoretical explanation for this cross-sectional difference in allocation value through the batch-shortening effect, since competition fragments the ranked list across different transplant centers and makes it harder for centers to optimize organ allocation internally.

Finally, our proposed alternative mechanism, the priority-proposal rule, is also heavily evidence-based. First, this mechanism is already being tested in some local markets as a form of AOOS. New York Times (2025) reports that one OPO, LiveOn NY, invites a few preferred centers to propose their highest-ranked patient on the list for whom they would accept the organ, and the organ is then given to the highest-priority patient among those proposed. Second, broader adoption of such a mechanism has been suggested by researchers and practitioners. Henson et al. (2026) emphasize the importance of time constraints in organ placement and suggest that the priority-proposal rule may expedite the allocation process and potentially reduce the need for AOOS, a phenomenon that has generated significant public concern about the fairness and transparency of the system. Our paper directly answers their call for further study of the "logistical details, theoretical implications, and empirical effects of such mechanisms on allocation efficiency and fairness" (Henson et al., 2026, p. 1453). Our framework makes this comparison exact: the priority-proposal rule is strategically equivalent to the sequential-offer rule, and therefore preserves both priority adherence and allocation value while speeding up organ placement.

7 Conclusion

We develop a canonical model of organ allocation with informed intermediaries. In the U.S. transplant system, centers possess information about their patients and accept organ offers on their behalf. This delegation is a practical necessity. A centralized planner cannot observe, maintain, and process all clinically relevant information in real time. Transplant centers therefore play an essential role in using information that the allocation system itself does not fully possess.

At the same time, delegation creates a wedge between formal and real allocation. Centers are not simply implementing the priority ranking mechanically; they are making decisions for a portfolio of patients, based on their own clinical information, institutional objectives, and performance incentives. As a result, a center may bypass a higher-priority patient in order to place the organ with another patient with a higher allocation value. Such behavior need not reflect an attempt to game the system. It can instead be understood as centers doing the best they can for their own patients under the incentives and information structure they face.

Our analysis provides a framework for studying the consequences of this delegation. We introduce the *priority adherence index* to measure the wedge between formal and real allocation, and use it together with expected allocation value to quantify the tradeoff between priority adherence and allocation value. The model identifies three forces—*pool thinning*, *batch shortening*, and *strategic delay*—which explain how market structure, organ selectivity, and allocation procedures shape both the extent to which priority is followed and the value generated by the final allocation.

The broader lesson is that imperfect priority adherence should be interpreted as an allocation-design tradeoff rather than as evidence of misconduct by transplant centers. Center discretion can improve the use of private information, but it can also weaken adherence to the formal priority order. The policy challenge is therefore not to eliminate center discretion, but to design allocation rules that better channel this discretion toward system-wide objectives. Our analysis of the priority-proposal rule illustrates this approach: the rule preserves the same priority adherence and allocation value as the current sequential-offer rule while substantially speeding up organ placement.

Several important extensions remain. Future work could incorporate OPO incentives more explicitly, following Chan and Roth (2024), and study how data infrastructure, performance regulation, and alternative allocation procedures interact. More broadly, our framework can be used to evaluate ongoing reforms in the transplant system by asking not only whether they improve priority adherence or allocation value in isolation, but also how they balance the information benefits of delegation against the fairness and transparency concerns that delegation inevitably creates. Our mechanism-design analysis suggests promising opportunities to improve both objectives. Realizing these opportunities will require further theoretical exploration, richer data infrastructure, and continued collaboration between academics and policymakers.

References

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(3): 729–747.

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(4): 425–442.

Adler, Joel T, Heidi Yeh, James F Markmann, and Louis L Nguyen. 2015. "Market Competition and Density in Liver Transplantation: Relationship to Volume and Outcomes." *Journal of the American College of Surgeons*, 221(2): 524–531.

Adler, Joel T, Heidi Yeh, James F Markmann, and Louis L Nguyen. 2016. "Temporal Analysis of Market Competition and Density in Renal Transplantation Volume and Outcome." *Transplantation*, 100(3): 670–677.

Adler, Joel T, Rosh KV Sethi, Heidi Yeh, James F Markmann, and Louis L Nguyen. 2014. "Market Competition Influences Renal Transplantation Risk and Outcomes." *Annals of Surgery*, 260(3): 550–557.

Agarwal, Nikhil, and Eric Budish. 2021. "Market Design." In *Handbook of Industrial Organization.* Vol. 5, 1–79. Elsevier.

Agarwal, Nikhil, Itai Artisans Ashlagi, Paulo Somaini, and Daniel Waldinger. 2018. "Dynamic Incentives in Wait List Mechanisms." Vol. 108, 341–347, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy*, 105(1): 1–29.

Ball, Ian, Matthew O. Jackson, and Deniz Kattwinkel. 2022. "Comment on Jackson and Sonnenschein (2007) "Overcoming Incentive Constraints by Linking Decisions"." *Econometrica*, 90(4): 3–7.

Benkert, Abigail R, Joseph B Lerman, Jacob N Schroder, Chetan B Patel, Adam D DeVore, Kunal J Patel, Jeffrey E Keenan, Carmelo A Milano, and Oliver K Jawitz. 2025. "Out of Sequence Donor Heart Allocation: A United Network for Organ Sharing Registry Analysis." *Journal of Cardiac Failure.*

Chan, Alex, and Alvin E Roth. 2024. "Regulation of Organ Transplantation and Procurement: A Market-Design Lab Experiment." *Journal of Political Economy*, 132(11): 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(61): 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.

Frankel, Alexander. 2014. "Aligned Delegation." *American Economic Review*, 104(1): 66–83.

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

Guan, Grace, Sanjit Neelam, Joachim Studnia, Xingxing S Cheng, Marc L Melcher, Michael A Rees, Alvin E Roth, Paulo Somaini, and Itai Ashlagi. 2025. "Insights From Refusal Patterns for Deceased Donor Kidney Offers." *Transplantation*, 10–1097.

Guo, Yingni, and Eran Shmaya. 2023. "Regret-Minimizing Project Choice." *Econometrica*, 91(5): 1567–1593.

Halldorson, Jeffrey B, Harry J Paarsch, Jennifer L Dodge, Alberto M Segre, Jennifer Lai, and John Paul Roberts. 2013. "Center Competition and Outcomes Following Liver Transplantation." *Liver Transplantation*, 19(1): 96–104.

Health Resources and Services Administration. 2025. "How Organ Allocation Works." https://www.hrsa.gov/optn/patients/organ-transplants/how-organ-allocation-works, Accessed: May 2026.

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, S Ali, Kristen L King, Stephen Pastan, Rachel E Patzer, David J Cohen, Jai Radhakrishnan, and Sumit Mohan. 2019. "Association Between Declined Offers of Deceased Donor Kidney Allograft and Outcomes in Kidney Transplant Candidates." *JAMA Network Open*, 2(8): e1910312–e1910312.

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(4): 1122–1130.

Jackson, Matthew O., and Hugo F. Sonnenschein. 2007. "Overcoming Incentive Constraints by Linking Decisions." *Econometrica*, 75(1): 241–257.

King, Kristen L, S Ali Husain, Miko Yu, Joel T wagons 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(6): e2316936–e2316936.

Kwong, Allison J, W Ray Kim, John R Lake, David P Schladt, Dzhuliyana Handarova, Jesse Howell, Benjamin Schumacher, Samantha Weiss, Jon J Snyder, and Ajay K Israni. 2025. "OPTN/SRTR 2023 Annual Data Report: Liver." *American Journal of Transplantation*, 25(2): S193–S287.

Liyanage, Luckmini N, Daniyar Akizhanov, Suhani S Patel, Dorry L Segev, Allan B Massie, Darren E Stewart, and Sommer E Gentry. 2025. "Contemporary Prevalence and Practice Patterns of Out-of-Sequence Kidney Allocation." *American Journal of Transplantation*, 25(2): 343–354.

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(11): 3071–3078.

Miralles, Antonio. 2012. "Cardinal Bayesian Allocation Mechanisms Without Transfers." *Journal of Economic Theory*, 147(1): 179–206.

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

New York Times. 2025. "Organ Transplant System 'in Chaos' as Waiting Lists Are Ignored." *The New York Times.*

Ng, Han. 2025. "Federal Oversight and Strategic Choices of Kidney Transplant Centers." *Working Paper.*

Organ Procurement and Transplantation Network. 2026. "OPTN Policies." *https://www.hrsa.gov/sites/default/files/hrsa/optn/optn_policies.pdf*, Accessed: June 2026.

Sweat, Kurt R. 2025. "Endogenous Priority in Centralized Matching Markets: The Design of the Heart Transplant Waitlist." *Working Paper.*

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:

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

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

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

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 $\mathrm{Index}(F(0), n)$ strictly increases in $F(0)$.

We next show that $\mathrm{Index}(F(0), n)$ 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:

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

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 $p\in(0,1)$:

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

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 from the second-to-last batch if x_i is the highest value among those weakly after i within that 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.[10] 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

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

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 $n = m + 1$.

□

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 rival 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. □

D Equilibrium for Selective Organs and $k \geqslant 2$ Centers

Theorem D.1. *Assume $F(0) \in (0, 1)$ and $k \geqslant 2$. Let $w(c)$ be the probability of the affiliation map c. For each center $j \in \{1, \ldots, k\}$ and each nonempty set of its affiliated patients $I_j \subseteq \{1, \ldots, n\}$, let $\mathcal{C}(j, I_j) := \{c : c^{-1}(j) = I_j\}$ be the set of affiliation maps consistent with* *center j having exactly the patients in I_j. Furthermore, for each $i \in I_j$, let $\alpha_i(j, I_j)$ be the probability that center j makes its first acceptance at patient i, and $\delta_i(j, I_j)$ the probability that j's rival centers reject the organ for all their patients who precede i. An equilibrium consists of a profile $\{\alpha_i(j, I_j), \delta_i(j, I_j)\}_{i \in I_j}$ for each valid pair (j, I_j), such that:*

$$\alpha_i(j, I_j) = \Pr\big[x_i > 0 \text{ and } \delta_i(j, I_j)x_i \geqslant \delta_{i'}(j, I_j)x_{i'} \ \ \forall i' \in I_j\big] = \int_0^1 f(u) \prod_{i' \neq i,\, i' \in I_j} F\Big(\tfrac{\delta_i(j, I_j)u}{\delta_{i'}(j, I_j)}\Big)\, \mathrm{d}u,$$

$$\delta_i(j, I_j) = \frac{\sum_{c \in \mathcal{C}(j, I_j)} w(c) \Big(1 - \sum_{i' < i,\, c(i') \neq j} \alpha_{i'}\Big(c(i'), c^{-1}\big(c(i')\big)\Big)\Big)}{\sum_{c \in \mathcal{C}(j, I_j)} w(c)}.$$

For a realized type $\theta_j = \{(i, x_i) : i \in I_j\}$, center j accepts for patient $i \in I_j$ if $x_i > 0$ and $\delta_i(j, I_j)x_i \geqslant \delta_{i'}(j, I_j)x_{i'}$ for all $i' > i$ with $i' \in I_j$, and rejects otherwise.

When $k = 2$, center j's patient set I_j determines the entire affiliation map: all patients outside I_j must be affiliated with the other center. Thus $\mathcal{C}(j, I_j)$ is a singleton, and the theorem reduces to Theorem 4.2.

The proof follows the same logic as the proof of Theorem 4.2, given in Appendix B. The only additional step is that, when $k > 2$, center j's patient set I_j no longer pins down a unique affiliation map. Hence, in computing the reach probability $\delta_i(j, I_j)$, center j must average over all affiliation maps $c \in \mathcal{C}(j, I_j)$, weighted by their conditional probabilities. Apart from this averaging over consistent affiliation maps, the equilibrium argument is identical to the two-center case. We therefore omit the proof.