

DYNAMIC MULTI-HOSPITAL KIDNEY EXCHANGE: COORDINATION COSTS, COMPLEMENTARITIES, AND THRESHOLD POLICIES

Ömer Faruk Şahin*

Duygu Sili†

Özgür Yılmaz‡

September 14, 2026

Abstract

Cross-hospital kidney exchange can increase transplantation, but it requires transportation, regulatory compliance, and surgical synchronization—costly activities whose benefits accrue primarily to patients. We study the socially optimal timing and organization of cooperation in a dynamic two-hospital kidney-exchange system. Two analytical benchmarks—complementary and substitutable pools—show how inventory imbalances and the option value of waiting determine when the gains from cross-hospital matching justify the coordination cost.

We then solve a literature-calibrated model using a net-position representation that is exact under an immediate-clearing restriction, whose welfare cost we quantify separately; the computed value function exhibits the curvature patterns predicted by the benchmarks. Within the reported pool-size window, the cooperation region vanishes well before coordination costs exhaust the two-way surplus, while the illustrative health-economic cost range lies well inside the active region. The welfare value of cooperation is governed jointly by arrival composition and cross-hospital asymmetry: hospitals with similar mixes gain little even at zero cost, while gains rise steeply—also at fixed hospital sizes—when compositions generate persistent complementary imbalances. Under the baseline composition the planner preserves scarce reciprocal pairs rather than consuming them across hospitals, and cooperation is just in time rather than inventory-smoothing under arrival-triggered clearing. The benchmark curvature patterns and the full model’s computed threshold-like policy geometry remain stable across a registry-based second calibration and the waiting-cost and temporal sensitivity exercises, but the quantitative value and composition of cooperation change sharply with the reciprocal-to-opportunity arrival ratio. The results provide a first-best benchmark for when costly cooperation among hospitals creates social value.

Keywords: kidney exchange, multi-hospital coordination, dynamic matching, Markov decision process, market design.

JEL Classification: I11, D47, C61, C78, D61.

*Stanford University, Stanford, CA, 94305 USA; *e-mail:* osahin25@stanford.edu.

†Universidad Carlos III de Madrid, 28903 Getafe (Madrid), Spain; *e-mail:* dsili@eco.uc3m.es.

‡Corresponding author. Koç University, İstanbul, 34450 Türkiye; *e-mail:* ozyilmaz@ku.edu.tr.

1 Introduction

Multi-hospital kidney exchange is a coordination problem with a wedge between private and social returns. Kidney exchange [Roth et al., 2004, 2007] converts otherwise-stranded medical possibilities into realized transplants, and exchange platforms now account for a substantial share of U.S. living-donor transplants [Agarwal et al., 2019, Ashlagi and Roth, 2021]. As of January 2026, more than 90,000 Americans were waiting for a kidney transplant.¹ One important source of unrealized transplantation is market fragmentation: Agarwal et al. [2019] estimate that the U.S. kidney-exchange market operates well below efficient scale and that remedying its fragmentation could increase transplants by roughly 30 to 63 percent. The expansion of exchange platforms has also made salient a friction whose interaction with stochastic inventories cannot be captured by the standard static framework: cross-hospital matching requires surgical synchronization, regulatory and administrative effort, and organ-transfer logistics, imposing real resource costs that may be only partly reimbursed while the primary welfare gains accrue to patients.² Hospitals may therefore prefer to clear easy-to-match pairs internally and contribute only their harder cases to the common pool—the strategic withholding studied by Ashlagi and Roth [2012, 2014] and Toulis and Parkes [2015].

Throughout the paper, c represents a real resource cost of cross-hospital exchange that is internalized by the planner. We characterize the first-best allocation given this cost; whether reimbursement can induce hospitals to implement that allocation is a separate question outside the model.

The less appreciated point—and the focus of this paper—is that even after abstracting from decentralized incentives, a benevolent planner who internalizes the real coordination cost faces a nontrivial dynamic allocation problem. Incompatible and over-demanded pairs (those with highly compatible donors, such as *O*-blood-type donors) arrive stochastically, so every match decision weighs the immediate surplus against the option value of a thicker future pool [Ünver, 2010, Baccara et al., 2020, Akbarpour et al., 2020]; waiting, in turn, prolongs suffering and exposes patients to attrition, whose welfare consequences are represented in reduced form through discounting (Section 3). Coordination costs cut across this trade-off: matching across hospitals today is costlier but may unlock exchanges that are otherwise infeasible. Crucially, the welfare value of cooperation among healthcare providers depends not only on its direct cost but on whether their stochastic inventories accumulate and are compositionally complementary. Existing dynamic matching models are predominantly single-pool, while multi-hospital work emphasizes strategic disclosure and incentive design. We instead study the first-best dynamic allocation problem when inter-hospital exchange entails an explicit resource cost. To our knowledge, this is the first continuous-time, first-best model of multi-hospital kidney exchange with explicit per-exchange coordination costs. We derive curvature and threshold results in two tractable benchmarks and examine their implications numerically in the full model.

Our central economic result is that costly inter-hospital cooperation is governed by two composition margins: whether reciprocal inventory accumulates, and whether that inventory accumulates

¹UNOS (2026), “U.S. surpasses 49,000 organ transplants while deceased organ donations dip,” press release, January 28, 2026 (unos.org): more than 90,000 people are on the kidney waiting list, and more than 100,000 await an organ of any kind.

²U.S. regulation permits reasonable procurement, packaging, and transport costs in kidney paired donation to be billed between transplant hospitals as kidney acquisition costs (42 CFR §413.416). Such billing shifts who bears the transport component of c ; it does not eliminate the underlying resource cost, and the administrative and surgical-synchronization components remain largely outside such billing.

on opposite sides of the two hospitals. First, we formulate the centralized planner’s problem as a continuous-time Markov decision process and develop two analytical benchmarks. In the Complementary Pools benchmark, reciprocal inventories create value by combining; in the Substitutable Pools benchmark, same-type inventories compete for common matching opportunities. The benchmarks yield opposite value-function curvature and monotone threshold policies, showing how inventory composition shapes the option value of waiting and the timing of cooperation. Second, we solve the full multi-hospital action environment subject to immediate intra-hospital clearing, using its exact two-dimensional net-position representation and a literature-based benchmark calibration. A separate unrestricted four-dimensional computation shows that imposing immediate clearing reduces welfare at empty pools by less than 0.4% under the reported calibration. We use discounted occupation measures to translate the resulting policy geometry into on-path cooperation rates and welfare gains. Counterfactual hospital-composition exercises and a registry-based second calibration show that the reciprocal-to-opportunity arrival ratio determines whether inventory accumulates, while cross-hospital asymmetry determines whether that inventory is complementary. These composition margins affect the use and welfare value of cooperation more strongly than coordination costs do within the illustrative health-economic range. Waiting-cost and temporal sensitivity exercises preserve the benchmark curvature patterns, although all structural statements about the full model are numerical findings for the reported calibrations.

The numerical analysis yields three findings. First, the computed cooperation region contracts monotonically with c and, in our computations, is empty within the reported pool-size window by $c \approx 1$, well before $c = 2$ mechanically exhausts the surplus from a two-way exchange—though the shutdown point is itself composition-dependent (Section 4.7). An illustrative health-economic bracket for c (Section 4) lies well inside the active range. Second, the value of cooperation is small under similar hospital arrival mixes but rises sharply when the reciprocal-to-opportunity ratio permits inventory accumulation and cross-hospital asymmetry creates complementary imbalances. This composition-based mechanism differs from the scale-fragmentation margin studied by [Agarwal et al. \[2019\]](#). Under the baseline composition, the planner accordingly uses abundant reciprocal pairs while preserving scarce ones with greater continuation value. Third, under arrival-triggered clearing (Assumption 7), inter-hospital reciprocal exchange is generally just in time rather than a means of smoothing inventories; a batched benchmark shows the underlying retention mechanism survives periodic clearing (Section 4.9). A rebalancing motive arises only in a narrow, low-cost region of strongly imbalanced same-type pools. Together, these findings show that inventory composition and future matching opportunities—not static matching surplus alone—determine the social value of inter-hospital cooperation.

These findings bear on policy by characterizing the first-best timing and composition of inter-hospital exchange when cross-hospital coordination consumes real resources. In the arrival environments we compute, an explicit waiting cost makes the planner more, not less, willing to cooperate—each inter-hospital match sheds pairs from costly storage—so the case for investing in cross-hospital coordination is strongest precisely where waiting is most costly, such as thin markets with long inter-arrival times (Section 4.11). This allocation provides a benchmark for studying decentralized implementation through reimbursement schemes, credit mechanisms, or priority rules. Because the present model abstracts from transfers, private information, and strategic hospital behavior, it does not establish which of these instruments can implement the first best.

Related literature. The paper sits at the intersection of four literatures. The first is multi-hospital kidney exchange and mechanism design. Building on Roth et al. [2004, 2007], Ashlagi and Roth [2012] and Ashlagi and Roth [2014] show that withholding even a few easy matches imposes large social costs, while modest incentives, such as priority for pooling, align hospital interests at low cost. Toulis and Parkes [2015] use a random-graph model of static multi-hospital kidney exchange to design a mechanism, which achieves efficiency and incentivizes moderate-to-large hospitals to report their pairs fully. Ashlagi et al. [2015] obtain worst-case efficiency guarantees with a strategyproof mechanism for multi-hospital exchange, Blom et al. [2024] design rejection-proof mechanisms under which no hospital gains by rejecting proposed exchanges, and Ashlagi and Roth [2021] survey the operational lessons of this agenda. Closest to us in setting, Hajaj et al. [2015] analyze a dynamic, multi-period, multi-center kidney exchange and design a credit-based, strategy-proof mechanism that induces centers to disclose their pairs over repeated clearings. Their question is decentralized-incentive design; ours is the complementary first-best question of how a planner should allocate inventories over time when inter-hospital exchange carries an explicit per-exchange resource cost, and the curvature and threshold structure of that allocation. The second is the literature on dynamic matching, beginning with Ünver [2010] and developed by Baccara et al. [2020] and Akbargpour et al. [2020], which establish that optimal policies in single-pool environments take a monotone threshold form: easy pairs are matched immediately because their immediate value exceeds the option value of waiting. We show that this threshold structure persists, but takes qualitatively new forms, once pools are held by institutionally separate hospitals linked by costly coordination. The third is the empirical and structural literature on frictions in kidney exchange. Although kidney shipping is routine and does not harm outcomes [Segev et al., 2005], transport and coordination expenses are often unreimbursed, discouraging cooperation [Agarwal et al., 2019]. We summarize these frictions parsimoniously through the cost parameter c and connect them to a tractable theoretical decomposition that delivers comparative statics in c directly.

The fourth is the health economics of organizational boundaries, care fragmentation, and cost incidence. U.S. health care is produced by organizationally fragmented providers, and production that spans institutional boundaries is costly to coordinate [Cebul et al., 2008]; care divided across more providers is associated with higher costs without corresponding quality gains [Agha et al., 2019]. Kidney exchange is a salient case: its technology requires production across institutional boundaries, while c parsimoniously captures the resulting frictions and allows the model to determine when bearing these costs is socially worthwhile. The two-layer distinction between the resource cost of coordination and its incidence connects to work on how payment rules allocate costs across institutions [Cutler, 1995], and quasi-experimental evidence that kidney exchanges improve patient outcomes [Teltser, 2019] supports the welfare relevance of the cooperation margin we study.

Section 2 introduces medical compatibility and the pair-type reductions, while Section 3 develops the dynamic model and the CP/SP benchmarks. Section 4 solves the full multi-hospital model, quantifies its policy and welfare implications across arrival compositions, and examines robustness to a registry-based calibration, explicit waiting costs, batched clearing, and alternative temporal calibrations. Section 5 concludes with economic implications, limitations, and directions for decentralized implementation.

2 Medical Compatibility and Paired Exchange

Two conditions must hold for a kidney transplant to be feasible. The first is **blood-type compatibility**. ABO types impose a partial order: O donors are universal, AB patients are universal recipients, and A (B) patients can accept from O or A (B) donors.³ The second is **tissue compatibility**. A *positive cross-match*, indicating that the recipient’s serum contains antibodies against the donor’s HLA antigens, rules out transplantation under standard immunosuppressive protocols. HLA desensitization makes some positive-crossmatch transplants feasible but is rarely cost-effective in the registry setting we study; we abstract from it. Such immunological barriers, often caused by prior transplants, transfusions, or pregnancies, are independent of blood type and frequently prevent otherwise feasible transplants [Roth et al., 2004, Saidman et al., 2006]. Consequently, kidney exchange pools typically include both ABO-incompatible and positive-crossmatch pairs, broadening the feasible match set and improving the efficiency of exchange programs [Saidman et al., 2006].

We denote the type of a pair by $X-Y$, where X is the patient’s blood type and Y the donor’s. Let $\mathcal{T}$ be the set of all pair types. Pairs are classified as **over-demanded** if the donor’s type is widely compatible (e.g., O donor); **under-demanded** if the donor’s type is hard to match (e.g., AB donor); **self-demanded** if patient and donor share the same type and **reciprocally demanded** when the patient has blood type A (B) and the donor has blood type B (A), forming the $A-B$ and $B-A$ pairs. Formally, the set of all pair types divides into four categories; Table 1 lists each category’s members and its role in the reduced model developed below.

Category	Blood-type pairs	Role in the reduced model
Over-demanded	$A-O, B-O, AB-O, AB-A, AB-B$	Matched on arrival (Assumption 6); enter as OD_{A-B} or OD_{B-A} upgrade events
Under-demanded	$O-A, O-B, O-AB, A-AB, B-AB$	Exogenous, always-available reservoir (Assumption 2); not part of the state
Self-demanded	$A-A, B-B, O-O, AB-AB$	Excluded (Assumption 3)
Reciprocally demanded	$A-B, B-A$	The state variables s_i^{AB}, s_i^{BA}

Table 1: Pair-type classification and each category’s role in the reduced model.

The restrictions introduced below serve three distinct purposes. Some define the simplified environment in which the analytical results are proved; some reduce the state or action space for numerical tractability; and others approximate operational features of current kidney-exchange programs. These roles are not mutually exclusive. To make the scope of the results transparent, we state the role of each assumption when it is introduced and summarize all dependencies at the end of Section 3.1.

HLA incompatibility is empirically important and is modeled explicitly in much of the kidney-exchange literature [Saidman et al., 2006, Ashlagi and Roth, 2014]. We abstract from it as follows.

Assumption 1. (No Cross-Pair HLA Incompatibility) *No patient is tissue-type incompatible with the donor of another pair.*

Role. This nonempirical tractability restriction makes compatibility type-determined: with it, whether two pairs can exchange depends on blood types alone, so counts of pairs by type form a sufficient state and the feasibility conditions of every Bellman equation are exact. The abstract CP

³Blood type depends on the presence of A and B antigens on red blood cells. A patient can receive a kidney if they possess all of the donor’s antigens.

and SP results do not depend on it, but their kidney-exchange interpretation—and the full model’s state space—do.

Under-demanded pairs are the pool’s most abundant category and tend to accumulate because they can generally exit only alongside an over-demanded partner [Roth et al., 2007, Ünver, 2010]. This motivates the following thick-market approximation.

Assumption 2. (*Under-Demanded Reservoir*) *At every decision epoch, a compatible under-demanded pair is available for each over-demanded arrival. Under-demanded pairs form an exogenous, non-binding reservoir and are not included in the state.*

Role. This approximation removes under-demanded inventories from the state and guarantees a baseline two-way match for every over-demanded arrival; the analytical benchmarks take the associated outside opportunities as primitives.

Self-demanded pairs can be matched with one another or absorbed into larger exchanges, sometimes altering feasibility under the three-pair cap in Assumption 5. Hence their exclusion removes genuine exchange opportunities and is not a loss-free simplification (cf. Roth et al., 2007).

Assumption 3. (*No Self-Demanded Pairs*) *Self-demanded pairs do not arrive.*

Role. This nonempirical tractability restriction removes self-demanded inventories and their associated actions from the state and action spaces; the CP and SP results do not depend on it.

Blood type AB is relatively rare in practice, but an $AB-O$ pair is more flexible than other over-demanded types: it can support either reciprocal direction and can participate in several larger exchange configurations.⁴ To preserve the reduced type system, we impose the following restriction.

Assumption 4. (*Exogenous $AB-O$ Split*) *Each arriving $AB-O$ pair is assigned exogenously to exactly one of the two compatibility classes, OD_{A-B} or OD_{B-A} , independently of the current state. The assignment probabilities are fixed ex ante and are equal to one-half in the calibration of Section 4.*

Role. This nonempirical dimensionality reduction removes the pair’s endogenous flexibility to serve whichever reciprocal type is scarcer. The abstract CP and SP results do not depend on it, but their kidney-exchange interpretation and all numerical results are conditional on the imposed split.

In line with standard practice, we restrict attention to exchanges involving at most three pairs. Cycles larger than three are rarely implemented because all surgeries must be performed simultaneously, creating substantial logistical and coordination challenges across hospitals.

Assumption 5. (*No Large Exchanges*) *The maximum exchange size is three.*

Role. This restriction both approximates current operational practice, in which cycles larger than three are uncommon, and limits the action space. It determines the kidney-exchange interpretation of the model’s rewards but is not essential to the abstract curvature arguments, which are stated in terms of generic rewards.

An over-demanded pair can typically form a three-way exchange with an under-demanded pair and a reciprocally demanded pair; examples include the three-way exchanges $(B - O, O - A, A - B)$ and $(A - O, O - B, B - A)$. We classify over-demanded pairs by whether they can form a three-way

⁴For example, together with a reciprocal pair it can support a four-way exchange such as $(AB-O, O-A, A-B, B-AB)$, while with two under-demanded pairs it can support a three-way exchange such as $(AB-O, O-A, A-AB)$. These exchanges generate surpluses of four and three, respectively.

exchange with $A-B$ or with $B-A$. The set of over-demanded pairs that can form a three-way exchange with $A-B$ is $\mathcal{OD}_{A-B} = \{B-O, AB-O, AB-A\}$. Similarly, the set of over-demanded pairs that can form a three-way exchange with $B-A$ is $\mathcal{OD}_{B-A} = \{A-O, AB-O, AB-B\}$. These are *compatibility classes* rather than a partition of arrivals: $AB-O$ appears in both because it can serve either direction, and each realized $AB-O$ arrival is assigned to exactly one class by the exogenous split introduced above—equivalently, the two entries may be read as labeled copies $AB-O^{AB}$ and $AB-O^{BA}$.

3 Multi-Hospital Dynamic Kidney Exchange

We study a continuous-time dynamic kidney exchange with two hospitals, h_1 and h_2 , coordinated by a centralized clearinghouse. Incompatible patient–donor pairs arrive stochastically, and the planner chooses whether to match them immediately—either within a hospital (*intra-hospital exchange*) or across hospitals (*inter-hospital exchange*)—or to defer matching at the expense of enlarging the pool. (A schematic of the three exchange channels that consume reciprocal pairs is in Section I of the online appendix, Figure OA.1.)

- *Stochastic arrival.* Incompatible pairs arrive to the two-hospital system according to a single Poisson process with aggregate rate λ . Each arriving pair is located at h_1 or h_2 with equal probability $\frac{1}{2}$, independently across arrivals, and its type $X-Y \in \mathcal{T}^R$ is drawn independently with probability p_{X-Y} , where $\sum_{X-Y} p_{X-Y} = 1$. Hence the event “a pair of type $X-Y$ arrives at hospital h_i ” occurs at rate $q_i^{X-Y} \lambda$, with $q_i^{X-Y} := \frac{1}{2} p_{X-Y}$. Let ρ denote the continuous-time discount rate. Under this arrival structure, the dynamic exchange problem can be formulated as a *Markov Decision Process (MDP)*.

The clearinghouse maximizes the expected discounted stream of net match surplus. Let $0 < T_1 < T_2 < \dots$ denote the arrival epochs of the aggregate Poisson process and let r_n denote the net surplus realized at T_n : the number of transplants performed by the exchange (if any) conducted at T_n , minus the coordination cost c if that exchange is inter-hospital, and $r_n = 0$ if no exchange is conducted. The objective is

$$v = \mathbb{E} \left[\sum_{n=1}^{\infty} e^{-\rho T_n} r_n \right].$$

Since $c \in [0, 2]$ (maintained throughout; see the inter-hospital exchange description below), $0 \leq r_n \leq 3$, and $\rho > 0$, the expectation is well defined.⁵

Over-demanded pairs are highly matchable: for example, an $A-O$ pair can immediately generate two transplants with an $O-A$ pair, whereas the latter produces no transplant on its own. Following the standard treatment in the literature [Ünver, 2010, Roth et al., 2007], we impose immediate matching. This observation motivates, but does not establish the optimality of the restriction.

Assumption 6. (Immediate Matching of Over-Demanded Pairs) *Every arriving over-demanded pair is matched on arrival and is never stored. It is matched either in a baseline two-way exchange with an under-demanded pair, available by Assumption 2, yielding two transplants, or in a three-way exchange that additionally consumes one compatible reciprocally demanded pair, yielding three transplants. The coordination cost c is deducted when that reciprocal pair is held at the other hospital.*

⁵Discounting is a reduced-form penalty for delay: postponing a unit transplant by T forgoes $1 - e^{-\rho T} = \int_0^T \rho e^{-\rho t} dt$, a discounted flow cost at rate ρ , not an undiscounted unit waiting cost; cf. Ünver [2010, footnote 4]. An explicit flow cost κ is analyzed in Section 3.4.

Role. This operational approximation removes over-demanded inventories from the state and generates the outside-opportunity structure used to map the CP and SP benchmarks to kidney exchange; the abstract benchmark proofs take those opportunities as primitives.

These assumptions simplify the model by restricting decisions to reciprocally demanded and over-demanded pairs (Assumptions 2 and 3) and by limiting exchanges to at most three pairs (Assumption 5). Under Assumption 2, an arriving over-demanded pair can always be matched in a two-way exchange with an under-demanded pair, generating a baseline surplus of two transplants. The dynamic problem therefore concerns only the choice of *which* reciprocal pair (if any) participates in the three-way exchange that the over-demanded pair can additionally enable. Assumption 6 thus forces no three-way exchange: the planner may always let the opportunity pass and settle for the baseline two-way exchange. In this restricted sense, over-demanded pairs “matter” only through their three-way interactions with reciprocally demanded pairs. Consequently, the type space reduces to the following set:

$$\mathcal{T}^R := \{A-B, B-A, OD_{A-B}, OD_{B-A}\}.$$

- *Inter-hospital exchange.* Hospitals may match incompatible patient–donor pairs internally (i.e., intra-hospital exchange) or coordinate to match across institutions (i.e., inter-hospital exchange). Inter-hospital matches entail a transport and coordination cost $c \in [0, 2]$, measured in transplant-equivalent welfare units: surpluses count transplants, so c is the number of transplant-equivalents absorbed by coordination, and the upper bound $c \leq 2$ keeps the inter-hospital two-way surplus $2 - c$ nonnegative (equivalently, $0 \leq r_n \leq 3$ in the objective above).

Homogeneous-cost benchmark. We use a single coordination cost c for every inter-hospital exchange. This specification aggregates transportation, administrative, regulatory, and surgical-synchronization costs and should be interpreted as a baseline rather than as a claim that realized costs are identical across exchanges. Fixed heterogeneity across exchange types, hospital directions, or transport modes can be incorporated through action-specific net rewards, preserving the basic comparison between immediate surplus and the option value of inventory. State-dependent or hospital-pair-specific costs, however, may alter the scalar- c comparative statics and threshold structure and would require a separate analysis. Because surplus is measured in transplant-equivalent units, c can be interpreted as the incremental resource cost of an inter-hospital exchange relative to the social value of one transplant; Section 4 motivates the illustrative range used in the numerical analysis.

Inter-hospital exchange is chosen only when its immediate net surplus *plus continuation value* dominates the available alternatives. Thus, an exchange with positive immediate surplus may be optimally declined because the relevant reciprocal inventory has greater continuation value; the Bellman equation (1) gives the exact comparison.

Coordination costs and waiting costs. The coordination cost c in our model is conceptually distinct from the cost of waiting in the pool. We follow standard practice in the dynamic matching literature [Ünver, 2010, Akbarpour et al., 2020] and treat patients as non-perishable, with the social cost of additional waiting captured implicitly through the discount factor ρ . In reduced form, delaying a unit transplant is penalized like a *discounted* flow cost at rate ρ (footnote 5); we do not claim equivalence to an undiscounted flow cost. An explicit waiting cost κ is analyzed later: Proposition 2 shows that the CP and SP curvature and threshold results survive it, and Section 4.11 examines the

full-model implications numerically. The coordination cost c , by contrast, is paid only when an inter-hospital action is selected, and reflects the differential cost of inter-hospital matching over intra-hospital matching: the transport, regulatory, and surgical-synchronization frictions specific to cross-institutional exchanges. Within a single hospital, c is irrelevant; with multiple hospitals, the hospitals arranging an exchange bear c while the transplant gains accrue to patients, and it is this incidence—not c itself—that creates the wedge between private and social returns to cross-hospital cooperation. The discount factor ρ and the coordination cost c therefore play distinct roles: ρ shapes the timing decision (match now versus wait), while c shapes the cross-institutional decision (match within or across hospitals).

An exchange policy is a rule that, upon each donor–patient pair arrival at hospital $h \in \{h_1, h_2\}$, determines whether to (i) defer any feasible exchange involving the arriving pair and retain it in the local pool, (ii) perform an intra-hospital exchange by matching the pair with compatible pairs within the same hospital, or (iii) perform an inter-hospital exchange by matching it with compatible pairs at the other hospital. A policy is **dynamically efficient** if it maximizes the expected discounted surplus of future exchanges, net of inter-hospital transport and coordination costs. Throughout, we reserve the term *mechanism* for decentralized settings with private information and incentives, which we do not model here.

By the memorylessness of Poisson arrivals, the process is Markovian: future dynamics depend only on the current pool composition, not on its history. Since over-demanded pairs are matched immediately (Assumption 6), the relevant state space consists only of reciprocally demanded pairs:

$$s = (s_1^{AB}, s_1^{BA}, s_2^{AB}, s_2^{BA}) \in \mathbb{N}_0^4,$$

where s_i^{AB} and s_i^{BA} denote the numbers of $A-B$ and $B-A$ pairs at hospital h_i . We use j to denote a reciprocally demanded pair—either type $A-B$ or $B-A$. For a given hospital h_i , let h_{-i} denote the other hospital. Similarly, for a reciprocally demanded type $j \in \{A-B, B-A\}$, let $-j$ denote the opposite reciprocally demanded type. The expected discount factor until the next arrival is

$$\mathbb{E}[e^{-\rho\tau}] = \frac{\lambda}{\lambda + \rho} = \delta,$$

where τ is the time of the next arrival. Section 4 explains the numerical choice of δ , its annual-time interpretation, and the corresponding sensitivity analysis.

For each h_i and type j , let e_i^j denote the corresponding unit vector, and recall that $-j$ denotes the reciprocal type opposite to j . The feasible actions, rewards, and transitions at each arrival are as follows. Upon the arrival of a reciprocal pair of type j at h_i , the planner chooses among three actions: (i) *retention*: the pair joins the pool, $s \mapsto s + e_i^j$; (ii) an *intra-hospital two-way exchange* with a pair of type $-j$ at h_i (feasible if $s_i^{-j} \geq 1$): $s \mapsto s - e_i^{-j}$, surplus 2; (iii) an *inter-hospital two-way exchange* with a pair of type $-j$ at h_{-i} (feasible if $s_{-i}^{-j} \geq 1$): $s \mapsto s - e_{-i}^{-j}$, surplus $2 - c$. Upon the arrival of an over-demanded pair $\mathcal{OD}_j$ at h_i , the planner chooses among three actions: (i) the *baseline two-way exchange* with an under-demanded pair: $s \mapsto s$, surplus 2; (ii) an *intra-hospital three-way exchange* consuming a reciprocal pair of type j at h_i (feasible if $s_i^j \geq 1$): $s \mapsto s - e_i^j$, surplus 3; (iii) an *inter-hospital three-way exchange* consuming a reciprocal pair of type j at h_{-i} (feasible if $s_{-i}^j \geq 1$): $s \mapsto s - e_{-i}^j$, surplus $3 - c$.

Consequently, both the intra- and the inter-hospital channel support two-way and three-way exchanges; the two channels differ not in the size of the exchange but in the net surplus generated, because inter-hospital exchanges incur the coordination cost c .

The model evaluates exchange opportunities only when a new pair or over-demanded opportunity arrives. We formalize this event-driven timing as follows.

Assumption 7. (*Arrival-Triggered Decisions*) *Decision epochs coincide with arrivals. At each decision epoch, every exchange must involve the arriving pair or opportunity; an exchange consisting exclusively of previously stored pairs cannot be initiated.*

Role. This substantive timing convention approximates a high-frequency, event-driven clearing regime in which a new arrival prompts immediate evaluation of the available exchange opportunities. It is used in all Bellman equations, benchmark proofs, and the counterexample in Appendix B. Platforms that clear periodically instead allow multiple arrivals to accumulate between match runs and may permit exchanges involving only previously stored pairs. Our results should therefore be interpreted as conditional on event-driven clearing rather than as an exact description of every platform’s match-run schedule.

3.1 The Full Multi-Hospital Model

We can now state the planner’s problem in full. At each event of the aggregate Poisson process exactly one arrival is realized: a pair of type $T \in \mathcal{T}^R$ at hospital h_i , with probability $q_i^T = \frac{1}{2} p_T$.⁶ Feasible actions, rewards, and transitions are as listed above; an action that would consume a pair not present in the pool is infeasible and is excluded from the maximization (Remark 1). Writing $j \in \{A-B, B-A\}$ for reciprocal types, the value function of the full model is the unique bounded solution of

$$\begin{aligned} v(s) = & \delta \sum_{i \in \{1,2\}} \sum_j \left[q_i^j \max \left\{ v(s + e_i^j), \quad 2 + v(s - e_i^{-j}), \quad 2 - c + v(s - e_{-i}^{-j}) \right\} \right. \\ & \left. + q_i^{O_{Dj}} \max \left\{ 2 + v(s), \quad 3 + v(s - e_i^j), \quad 3 - c + v(s - e_{-i}^j) \right\} \right], \end{aligned} \quad (1)$$

where, in each maximization, a branch that consumes a pair appears only if the corresponding stock is positive. One-step rewards are bounded by 3 and $\delta \in (0, 1)$, so the operator defined by the right-hand side of (1) is a δ -contraction on the space of bounded functions on $\mathbb{N}_0^4$; existence, uniqueness, and convergence of value iteration are recorded in Lemma 4 (Appendix C), stated once for all three Bellman operators of the paper. The CP and SP benchmarks of Sections 3.2 and 3.3 are stylized models that isolate the complementarity and substitutability forces at work in, respectively, the cross-type and same-type quadrants of the net-position representation (Lemma 1 below, under Assumption 8). They are *not* literal quadrant restrictions of (1): in the full model an arrival can move the state across quadrants, so a quadrant does not inherit the full model’s continuation values, and each benchmark retains selected event classes only—the quadrant’s own reciprocal arrivals and the over-demanded upgrades of Table 3. Reciprocal arrivals of the opposite type, which in the full model compete for the same stored pairs, are outside the benchmark event menus. Section 4.10 tests, quadrant by quadrant, whether the benchmark curvature properties nonetheless hold locally

⁶Symmetry across hospitals—equal arrival shares and a common type distribution—is a substantive simplification, imposed for tractability and used in the calibration of Section 4. The formulation extends verbatim to general probabilities q_i^T with $\sum_{i,T} q_i^T = 1$: the contraction argument below uses only $\delta \in (0, 1)$ and bounded rewards, and the benchmark theorems of Sections 3.2–3.3 already allow arbitrary arrival probabilities.

in the full model.

Remark 1. (Boundary convention) *An action that would consume a pair not present in the pool is infeasible and is excluded from every maximization; equivalently, the corresponding marginal value is set to $+\infty$, as in the definitions of the one-step differences below. We do not extend v by zero to states with negative components: such an extension would allow the Bellman operator to “match” nonexistent pairs at boundary states. Accordingly, all second-order inequalities in Lemmas 2 and 3 and in the supporting lemmas are asserted only at states where every difference they involve is defined.*

Kidney-exchange programs conduct frequent match runs and generally clear internally compatible pairs when they are identified. Delaying two immediately feasible transplants to preserve one pair for a future three-way exchange may also face ethical or regulatory constraints not captured by the reward function. These considerations motivate, but do not establish the optimality of, the following restriction.

Assumption 8. (Immediate Intra-Hospital Clearing) *Suppose a reciprocal pair of type j arrives at hospital h_i while a pair of type $-j$ is present in its pool, so that $s_i^{-j} \geq 1$. The planner must match the arriving pair on arrival, either in the intra-hospital two-way exchange or, when feasible and preferred, in an inter-hospital exchange. Retention is infeasible only in these states; all other action sets remain unrestricted.*

Role. This institutional approximation rules out coexistence of reciprocal types within a hospital and thereby enables the exact two-dimensional net-position representation in Lemma 1 and the constrained computation in Section 4. It is not required for the CP or SP benchmark results, and the unrestricted four-dimensional computation drops it.

An analogous immediacy restriction appears in Ünver [2010], where over-demanded pairs are matched upon arrival. The restriction is not innocuous as a matter of theory—for δ close to one the unrestricted planner would sometimes retain⁷—but it is quantitatively mild at our calibration; Remark 2 quantifies its cost.

Lemma 1. (Net-position reduction under immediate clearing) *Impose Assumption 8 on the full model (1), and let the initial state satisfy $\min\{s_i^{AB}, s_i^{BA}\} = 0$ for each i (e.g., empty pools). Then $\min\{s_i^{AB}, s_i^{BA}\} = 0$ holds at every subsequent state, and on this set the constrained-optimal value function depends on $(s_1^{AB}, s_1^{BA}, s_2^{AB}, s_2^{BA})$ only through the net positions $x_i := s_i^{AB} - s_i^{BA}$, $i = 1, 2$.*

Remark 2. (Cost of the clearing restriction) *Re-solving the model without Assumption 8 at the calibration of Section 4 quantifies what the assumption costs. Retention is strictly optimal only at high patience: writing $\tilde{\delta} = (\delta/2)/(1 - \delta/2)$ for the discount factor between a hospital’s own events, no violation occurs for $\tilde{\delta}$ below 0.98, and the first violations appear at $\tilde{\delta} = 0.99$. At the global $\delta = 0.999$ the worst-case action gap is 0.13, arising at depleted states such as $(s_i^{AB}, s_i^{BA}) = (0, 1)$, and imposing the assumption costs less than 0.4% of the planner’s value at empty pools for every reported c ; Table OA.1 in Section II of the online appendix reports the comparison for each value of c .*

The empty state provides a common ex ante benchmark for the constrained and unrestricted problems and corresponds to a newly initiated or fully cleared two-hospital pool. It is not intended to represent

⁷For a counterexample, in which retaining the arriving pair is strictly optimal and the pair bound $v^*(s+e_i^j) - v^*(s-e_i^{-j}) \leq 2$ fails, see Proposition 3 in Appendix B.

the long-run distribution of pool compositions. Accordingly, the reported percentage is the discounted welfare loss from this particular initial state, not a uniform bound over all states. The worst-case action gap of 0.13, concentrated at depleted states, separately measures the largest local discrepancy between immediate clearing and unrestricted retention under the reported calibration.

All statements in Section 4 accordingly characterize the constrained optimum within the immediate-clearing class.

Before presenting our theoretical analysis of the benchmarks, we summarize the assumptions used in the paper and their roles in Table 2.

Table 2: Roles and scope of the modeling assumptions

Restriction	Analytical role	Computational role	Institutional interpretation
Assumption 1 <i>No Cross-Pair HLA Incompatibility</i>	Makes compatibility type-determined: pair counts are a sufficient state and every feasibility condition in the Bellman equations is exact, so the kidney-exchange interpretation of all results is conditional on it. The abstract CP and SP proofs themselves do not use it.	Simplifies the compatibility graph and type space.	None. A nonempirical simplification; HLA incompatibility is empirically important.
Assumption 2 <i>Under-Demanded Reservoir</i>	Guarantees the baseline two-way match and the under-demanded member of every three-way exchange, so over-demanded arrivals enter the Bellman equations with deterministic rewards.	Removes under-demanded inventories from the state.	A thick-market approximation motivated by the abundance and accumulation of under-demanded pairs.
Assumption 3 <i>No Self-Demanded Pairs</i>	Not required for the abstract CP or SP results.	Removes self-demanded inventories and their associated actions.	None. A nonempirical and non-loss-free tractability restriction.
Assumption 4 <i>Exogenous AB-O Split</i>	Not required for the abstract CP or SP results; their kidney-exchange interpretation is conditional on the reduced type system.	Assigns each <i>AB-O</i> arrival to one class, removing its endogenous flexibility.	None. Although <i>AB</i> is rare, the split is a nonempirical dimensionality reduction.
Assumption 5 <i>No Large Exchanges</i>	Determines the exchange opportunities and rewards in the kidney-exchange mapping; the abstract proofs use generic rewards.	Limits the action space to exchanges involving at most three pairs.	Approximates practice, in which larger cycles are uncommon because of synchronization and coordination requirements.
Assumption 6 <i>Immediate Matching of Over-Demanded Pairs</i>	Generates the outside-opportunity interpretation; the benchmark proofs take these opportunities as primitives.	Prevents over-demanded pairs from becoming inventory.	An operational approximation motivated by their high immediate matchability and standard treatment in the literature.
Assumption 7 <i>Arrival-Triggered Decisions</i>	Used in all Bellman equations, benchmark proofs, and the counterexample in Appendix B.	Eliminates a separate clearing stage involving only stored pairs.	A substantive timing convention, not an exact description of every platform’s match-run schedule; batched clearing is examined in Section 4.9.
Assumption 8 <i>Immediate Intra-Hospital Clearing</i>	Required for Lemma 1 (exactness of the net-position reduction), but not for the CP or SP results; Proposition 3 establishes that it is a genuine restriction.	Enables the two-dimensional net-position computation by ruling out within-hospital coexistence. The unrestricted computation drops it.	Approximates frequent internal clearing and constraints on delaying immediately feasible transplants.
CP and SP benchmark restrictions	Define the two-dimensional models used in Lemmas 2–3 and Theorems 1–2; these are not global results for the full model.	Not imposed in the unrestricted four-dimensional computation; each benchmark retains selected event classes.	Analytical thought experiments, not literal quadrant restrictions or institutional descriptions.

Notes: The roles are not mutually exclusive. “Analytical role” identifies the results or mappings that use a restriction; it does not imply empirical necessity. The CP and SP benchmarks are separate analytical models, while the unrestricted four-dimensional computation retains the full event menu and drops Assumption 8.

The inherent complexity of multi-hospital kidney exchange arises from its four-dimensional state space, which prevents us from obtaining the global structural characterization we seek with existing analytical methods. We therefore study two reduced environments that capture inter-hospital

interactions: the *Complementary Pools (CP)* model and the *Substitutable Pools (SP)* model. In the CP model, each hospital holds a different reciprocal type, so the two pools create value by combining — the types are complements. In the SP model, both hospitals hold the same reciprocal type, so the two pools compete for the same opportunities — the types are substitutes. We present each model using abstract notation that highlights the underlying economic structure, and then connect the abstract objects back to the kidney exchange primitives. These cases highlight the key structural differences in inter-hospital matching behavior.

In each benchmark model below, we denote the optimal value function by v^* . Since the CP and SP models are analyzed independently in their own state spaces, no ambiguity arises from using the same notation across models. We begin with the CP model.

3.2 Complementary Pools (CP) Model

In the complementary model, we consider the dynamic matching problem of two abstract types that we denote by X and Y . We consider these types to be compatible, so it is feasible for the planner to match one of each of these types, in exchange for a reward of a_{XY} .

We also consider two other arrivals that we denote by O_X and O_Y . We treat these arrivals not as types that the planner can add to the pool, but rather as abstract *opportunities*, such that whenever such an opportunity O_X arrives, the planner can feasibly use one of the X type pairs by itself (even in the absence of any Y type pairs), and gain a reward of a_X . Similarly, whenever an opportunity O_Y arrives, the planner can use a Y type pair by itself for a reward of a_Y .

Although we keep these opportunities abstract here, Table 3 connects them to the over-demanded pairs of the full model, and we argue that, quadrant by quadrant, the trade-offs of the full model take the form studied here; Section 4.10 verifies numerically that the corresponding curvature properties hold locally in the full model.

The state space is $\mathcal{S} = \mathbb{N}_0^2$, where $s = (s_1, s_2)$ denotes the number of available X type pairs (s_1) and Y type pairs (s_2). We denote the arrival probability of X by p_X , of Y by p_Y , of O_X by p_X^O , and of O_Y by p_Y^O ; exactly one event occurs at each epoch, so $p_X + p_Y + p_X^O + p_Y^O = 1$.

Types X and Y can be matched *directly* only with each other, while the outside opportunities O_X and O_Y provide alternative uses for each type; this direct reciprocal compatibility makes cooperation between the two pools valuable, and the value function naturally exhibits complementarities. This complementarity is reflected in the following arrival events and the alternative decisions available at state $s = (s_1, s_2)$:

- If a type X pair arrives, the planner may
 - match it with a type Y pair, yielding surplus a_{XY} , or
 - keep it in the pool.
- If an opportunity O_X arrives, the planner may
 - match a waiting type X pair with the opportunity, yielding reward a_X , or
 - let the opportunity pass, yielding reward 0.

Symmetrically, if a type Y pair arrives, the planner may match it with a waiting type X pair for surplus a_{XY} or keep it in the pool. If an opportunity O_Y arrives, the planner may use a type Y pair for reward a_Y or let the opportunity pass.

Abstract reward	Kidney-exchange counterpart	Value	Convention
<i>CP model: $X = A-B$ pairs at h_1, $Y = B-A$ pairs at h_2 (quadrant $x_1 > 0 > x_2$)</i>			
a_{XY}	inter-hospital two-way reciprocal exchange	$2 - c$	total
a_X	OD_{A-B} at h_1 : upgrade of its baseline two-way exchange to an intra-hospital three-way exchange	$3 - 2 = 1$	incremental
a'_X	OD_{A-B} at h_2 : upgrade to an inter-hospital three-way exchange	$(3 - c) - 2 = 1 - c$	incremental
a_Y, a'_Y	symmetric, for OD_{B-A} and the Y pool	$1, 1 - c$	incremental
<i>SP model: $X_1 = A-B$ pairs at h_1, $X_2 = A-B$ pairs at h_2 (quadrant $x_1, x_2 > 0$)</i>			
a_{ii}	OD_{A-B} at h_i consumes the pair at its own hospital	$3 - 2 = 1$	incremental
$a_{i,-i}$	OD_{A-B} at h_i consumes the pair at the other hospital	$(3 - c) - 2 = 1 - c$	incremental

Table 3: Mapping of abstract benchmark rewards to kidney-exchange surpluses. Opportunity rewards are *incremental*: by Assumption 6 the over-demanded pair receives the baseline two-way surplus of 2 regardless of the planner’s decision, so “letting the opportunity pass” carries reward 0 in the benchmarks and the reward of using it is the upgrade $3 - 2 = 1$ (intra) or $(3 - c) - 2 = 1 - c$ (inter). Reciprocal–reciprocal exchanges forgo no baseline, so a_{XY} equals the total surplus $2 - c$. The benchmarks retain selected event classes only: reciprocal arrivals of the opposite type, which in the full model compete for the same stored pairs, are outside the benchmark event menus (Section 3.1).

The planner thus faces two trade-offs: upon an O_X arrival, match a waiting X pair now for a_X or preserve it for a future exchange worth a_{XY} ; upon a Y arrival, execute the exchange now for a_{XY} or keep the X pair for future opportunities worth a_X and a_Y separately. Both are resolved by comparing immediate rewards with each pair’s contribution to the present discounted value of all future exchanges.

These decisions depend on the relative values of the payoff from acting immediately and the option value of holding a pair for future exchanges. To formalize these optimal decisions, we introduce notation for one-step differences in the value function. The marginal value of an additional X type pair is $\Delta_1 v(s_1, s_2) := v(s_1, s_2) - v(s_1 - 1, s_2)$ if $s_1 \geq 1$, $\Delta_1 v(s_1, s_2) := \infty$ otherwise. Similarly, the marginal value of an additional Y type pair is $\Delta_2 v(s_1, s_2) := v(s_1, s_2) - v(s_1, s_2 - 1)$ if $s_2 \geq 1$, $\Delta_2 v(s_1, s_2) := \infty$ otherwise. The joint gain from simultaneously having one additional X and one additional Y pair is the *diagonal* difference (written Δ_{1+2} , not to be confused with the mixed difference $\Delta_1 \Delta_2$)

$$\Delta_{1+2} v(s_1, s_2) \begin{cases} := v(s_1, s_2) - v(s_1 - 1, s_2 - 1) & \text{if } s_1, s_2 \geq 1 \\ := \infty & \text{o.w.} \end{cases}$$

We can now state the Bellman equation:

$$\begin{aligned} v(s_1, s_2) = & \delta \left[p_X \cdot \left(v(s_1 + 1, s_2) + \max \{ a_{XY} - \Delta_{1+2} v(s_1 + 1, s_2), 0 \} \right) \right. \\ & + p_Y \cdot \left(v(s_1, s_2 + 1) + \max \{ a_{XY} - \Delta_{1+2} v(s_1, s_2 + 1), 0 \} \right) \\ & + p_X^O \cdot \left(v(s_1, s_2) + \max \{ a_X - \Delta_1 v(s_1, s_2), 0 \} \right) \\ & \left. + p_Y^O \cdot \left(v(s_1, s_2) + \max \{ a_Y - \Delta_2 v(s_1, s_2), 0 \} \right) \right]. \end{aligned} \quad (2)$$

Each $\max\{\cdot, 0\}$ is a decision. An arriving X pair is matched immediately when a_{XY} exceeds the joint option value $\Delta_{1+2} v(s_1 + 1, s_2)$ of holding both pairs; a waiting X pair is matched with an arriving opportunity when a_X exceeds its marginal value $\Delta_1 v(s_1, s_2)$; otherwise the planner waits. The Y side

is symmetric.

Remark 3. *In the kidney-exchange mapping of Table 3, the reward of an opportunity depends on the hospital at which the over-demanded pair arrives (1 intra, $1 - c$ inter). Formally, this splits O_X into two classes with probabilities $p_X^{O, \text{own}}$ and $p_X^{O, \text{cross}}$ and rewards a_X and a'_X . The extension leaves all structural results unchanged: every step of the induction in Appendix C treats arrival terms separately, so an additional opportunity class with a different reward adds one more case of identical form.*

One-period rewards are bounded and $\delta \in (0, 1)$, so T_{CP} is a sup-norm δ -contraction: a unique bounded value function exists and value iteration converges uniformly. This standard result is recorded, once for all three Bellman operators of the paper, as Lemma 4 in Appendix C (Blackwell, 1965; Puterman, 1994; Stokey et al., 1989).

Our analysis of the value function proceeds by characterizing its second-order structural properties (see Appendix A).

Lemma 2. *The optimal value function v^* of the CP model is componentwise concave, supermodular, and diagonally concave, at every state where the corresponding differences are defined.*

These properties imply diminishing marginal value of additional pairs in larger states (s_1, s_2) , even though X and Y pairs are complements. A key consequence is that the optimal exchange rule in the complementary model takes the form of a threshold policy: there exist critical levels of s_1 and s_2 such that once pool sizes exceed these thresholds, it is optimal to match rather than wait. The proof is by induction on the value iteration.

Theorem 1. (Optimal Policy in the CP Model) *The optimal policy is characterized by state-dependent threshold functions $\bar{t}_{XY}(s_2)$ and $\bar{t}_{O_X}(s_2)$, taking values in $\mathbb{N}_0 \cup \{+\infty\}$ (a threshold equals $+\infty$ when the corresponding action region is empty), such that for each state $s = (s_1, s_2)$.⁸*

- upon a type X arrival, the two-way exchange with a waiting type Y pair is optimal if and only if $s_1 \geq \bar{t}_{XY}(s_2)$; and
- upon an O_X arrival, matching a waiting type X pair with it is optimal if and only if $s_1 \geq \bar{t}_{O_X}(s_2)$.

These threshold functions are monotone: $\bar{t}_{XY}(s_2)$ is nonincreasing in s_2 and $\bar{t}_{O_X}(s_2)$ is nondecreasing in s_2 . The results for Y arrivals and O_Y opportunities are symmetric, with the roles of s_1 and s_2 interchanged: upon a Y arrival, the exchange with a waiting X pair is optimal if and only if $s_2 \geq \bar{t}_{YX}(s_1)$, and analogously for O_Y , with symmetric thresholds $\bar{t}_{YX}$ and $\bar{t}_{O_Y}$.

The policy is monotone in the own-pool size: once s_1 crosses the relevant threshold, acting—exchanging with a waiting Y pair, or matching a waiting X pair with an arriving opportunity—is optimal and remains optimal for all larger s_1 ; below the threshold the arriving pair is retained (transition $s \mapsto (s_1 + 1, s_2)$) or the opportunity is let pass. The thresholds themselves are functions of the partner-pool size s_2 , with the monotonicities stated in the theorem. Under the mapping of Table 3, opportunities arrive in two classes with different rewards ($a_X = 1$ for the intra-hospital upgrade, $a'_X = 1 - c$ for the inter-hospital one); Theorem 1 applies to each class separately, with class-specific thresholds satisfying $\bar{t}_{O_X}(s_2) \leq \bar{t}'_{O_X}(s_2)$ —the higher-reward class acts at the lower threshold. The same logic applies symmetrically to Y arrivals and O_Y opportunities.

⁸At payoff ties the planner acts (performs the exchange or matches the waiting pair); under this convention the “if and only if” action regions are exact. The same convention applies in Theorem 2.

3.3 Substitutable Pools (SP) Model

We now consider a substitution model, where we have again two abstract types, this time denoted by X_1 and X_2 . Unlike the complementary model, these two types cannot be matched with each other. Instead, they can only be used when an appropriate opportunity arrives.

We consider the arrival of some opportunities that we denote by O_1 and O_2 . Whenever such an opportunity O_i arrives, the planner can use it to consume one of the X_1 or X_2 types to get a reward of a_{ij} , where a_{ij} denotes the reward of using the opportunity O_i to consume the type X_j . Crucially, each opportunity can be used with *either* type, so X_1 and X_2 compete for the same opportunities, acting as substitutes for one another.

Although we keep these types and opportunities abstract here, Table 3 connects them to the over-demanded pairs of the full model, and we argue that the trade-offs of the full model, when restricted to the case where both hospitals hold the same reciprocal type, take the form studied here; Section 4.10 verifies the corresponding curvature properties numerically.

The state space is $\mathcal{S} = \mathbb{N}_0^2$, where $s = (s_1, s_2)$ denotes the number of available X_1 types (s_1) and X_2 types (s_2). We denote the arrival probability of type X_1 with p_1 , type X_2 with p_2 , opportunity O_1 with p_1^O , and opportunity O_2 with p_2^O ; exactly one event occurs at each epoch, so $p_1 + p_2 + p_1^O + p_2^O = 1$.

Since both types serve the same role—each can be used to capture opportunities—the pools are substitutes: having more of one type reduces the incremental value of additional units of the other type. This substitutability is reflected in the following arrival events and the alternative decisions available at state $s = (s_1, s_2)$:

- If a type X_1 pair arrives, it is kept in the pool. No decision is required, since X_1 cannot be matched with X_2 . The state transitions to $(s_1 + 1, s_2)$. Symmetrically, an X_2 arrival transitions the state to $(s_1, s_2 + 1)$.
- If an opportunity O_i arrives (for $i \in \{1, 2\}$), the planner may
 - use the opportunity to consume a type X_1 pair, yielding reward a_{i1} , or
 - use the opportunity to consume a type X_2 pair, yielding reward a_{i2} , or
 - let the opportunity pass, yielding reward 0.

The planner thus faces the following tradeoff: whenever an opportunity O_i arrives, it must decide whether to use the opportunity at all, and if so, which type to consume. Using the opportunity to consume type X_j yields an immediate reward a_{ij} but sacrifices the future option value of having that pair available for later opportunities. Letting the opportunity pass preserves both pools but forgoes the immediate reward. This tradeoff is resolved by comparing the immediate reward against the marginal value of the pair that would be consumed, in terms of its contribution to the present discounted value of all future exchanges.

To formalize these optimal decisions, we introduce notation for one-step differences in the value function. The marginal value of an additional X_1 type pair is

$$\Delta_1 v(s_1, s_2) \begin{cases} := v(s_1, s_2) - v(s_1 - 1, s_2) & \text{if } s_1 \geq 1 \\ := \infty & \text{o.w} \end{cases}$$

Similarly, the marginal value of an additional X_2 type pair is

$$\Delta_2 v(s_1, s_2) \begin{cases} := v(s_1, s_2) - v(s_1, s_2 - 1) & \text{if } s_2 \geq 1 \\ := \infty & \text{o.w.} \end{cases}$$

We can now state the Bellman equation:

$$\begin{aligned} v(s_1, s_2) = & \delta \left[p_1 \cdot v(s_1 + 1, s_2) + p_2 \cdot v(s_1, s_2 + 1) \right. \\ & + p_1^O \cdot \left(v(s_1, s_2) + \max \{ a_{11} - \Delta_1 v(s), a_{12} - \Delta_2 v(s), 0 \} \right) \\ & \left. + p_2^O \cdot \left(v(s_1, s_2) + \max \{ a_{21} - \Delta_1 v(s), a_{22} - \Delta_2 v(s), 0 \} \right) \right]. \end{aligned} \quad (3)$$

In this equation, type arrivals simply augment the pool: a type X_1 arrival moves the state from (s_1, s_2) to $(s_1 + 1, s_2)$ with no decision required, and similarly for X_2 . The decisions arise solely upon opportunity arrivals, where each $\max \{ a_{i1} - \Delta_1 v(s), a_{i2} - \Delta_2 v(s), 0 \}$ operator represents an optimal choice among three options.

Upon an O_i arrival the planner compares three net gains: $a_{i1} - \Delta_1 v(s)$ (consume the own pair), $a_{i2} - \Delta_2 v(s)$ (consume the partner pair), and 0 (pass). The largest wins; passing is optimal exactly when the marginal value of each pool exceeds its opportunity reward.

Existence, uniqueness, and convergence for (3) are covered by Lemma 4 in Appendix C. The structural properties of the value function in the SP model imply diminishing marginal values across pools.

Lemma 3. *The optimal value function v^* of the SP model is componentwise concave, submodular, and anti-diagonally concave, at every state where the corresponding differences are defined.*

This structure leads to a threshold characterization of the optimal policy. Fix an opportunity type O_i and define the action gains at state s , $G_{i1}(s) := a_{i1} - \Delta_1 v^*(s)$, and $G_{i2}(s) := a_{i2} - \Delta_2 v^*(s)$, together with $A_{i1}(s) := G_{i1}(s)$ and $A_{i2}(s) := G_{i1}(s) - G_{i2}(s) = (a_{i1} - a_{i2}) + (\Delta_2 v^* - \Delta_1 v^*)(s)$. We say the *single-crossing condition* holds for O_i if

$$A_{i2}(s) \geq 0 \implies A_{i1}(s) \geq 0 \quad \text{for all } s \text{ with } s_1 \geq 1; \quad (\text{SC})$$

that is, wherever consuming the own-type pair weakly beats consuming the cross-type pair, it also weakly beats inaction.

Theorem 2. (Optimal Policy in the SP Model) *Fix $i \in \{1, 2\}$. There exist state-dependent threshold functions $\tau_{i1}(s_2)$ and $\tau_{i2}(s_2)$, taking values in $\mathbb{N}_0 \cup \{+\infty\}$, such that, upon an opportunity O_i arrival at state $s = (s_1, s_2)$:*⁹

- *not using the opportunity is optimal if and only if $s_1 < \min\{\tau_{i1}(s_2), \tau_{i2}(s_2)\}$;*
- *consuming a type X_1 pair is optimal if and only if $s_1 \geq \tau_{i1}(s_2)$; and*
- *consuming a type X_2 pair is optimal if and only if $\tau_{i2}(s_2) \leq s_1 < \tau_{i1}(s_2)$ and a type X_2 pair is available.*

The cross-type threshold $\tau_{i2}(s_2)$ is nonincreasing in s_2 ; under the single-crossing condition (SC), the own-type threshold $\tau_{i1}(s_2)$ is nondecreasing in s_2 .

⁹At payoff ties the planner acts, and a tie between consuming a type X_1 pair and consuming a type X_2 pair is resolved in favor of X_1 ; under these conventions the displayed action regions are exact and mutually exclusive.

We next present a sufficient condition for (SC) to hold.

Proposition 1. *If $a_{i1} \geq \sup_{\{s: s_1 \geq 1\}} \Delta_1 v^*(s)$, then $A_{i1}(s) \geq 0$ at every state with $s_1 \geq 1$, and hence (SC) holds: the single-crossing condition is satisfied whenever the own-type reward weakly exceeds the largest marginal value of a waiting pair.*

Absent (SC), the own-type threshold need not be monotone in s_2 ; Section 4.10 verifies (SC) at every interior grid state for all reported cost values, so the monotone characterization applies throughout our calibration.

The two monotonicities capture the substitution logic. A larger partner pool makes consuming a partner pair cheaper (concavity), so $\tau_{i2}(s_2)$ falls; at the same time it raises each X_1 pair’s relative option value—the planner prefers to draw on the abundant pool and preserve X_1 for states in which X_2 is unavailable—so, under single crossing, $\tau_{i1}(s_2)$ rises. The online appendix tabulates the resulting action regions of Theorem 2 by threshold configuration (Section IV, Table OA.2).

3.4 Robustness to an explicit waiting cost

The reduced-form treatment of waiting through ρ raises a natural question: would carrying an explicit waiting cost change the structural results of this section? It would not. Waiting costs enter the Bellman operators as a state-affine term, and the second-order structure that drives Lemmas 2–3 is invariant to affine terms:

Proposition 2. *Fix $\kappa \geq 0$ and add to the CP or SP benchmark a holding cost accruing at rate κ per waiting pair, so that the per-event reward gains the term $-\theta(s_1 + s_2)$ with $\theta := \kappa(1 - \delta)/\rho \geq 0$. Then the optimal value function satisfies all the properties of Lemma 2 (respectively Lemma 3), and the threshold characterizations of Theorems 1–2 hold verbatim, with κ -dependent threshold functions.*

The proof is short: the holding cost adds a state-affine term to every value-iteration iterate, affine terms have identically zero second differences, and the induction arguments behind Lemmas 2–3 use only second-order hypotheses and generic rewards; the threshold theorems then follow from the lemmas exactly as before. (Section 4.10 verifies the proposition numerically.)

Proposition 2 concerns the benchmark curvature and threshold results; it does not establish that immediate intra-hospital clearing remains optimal when waiting is costly. Because a stored reciprocal pair can provide future disposal capacity, waiting costs may instead increase the value of retention. Section 4.11 and Section III of the online appendix examine this interaction numerically.

We now connect the stylized models to the full model. The sign pattern of (x_1, x_2) splits the state space into four quadrants: in same-type quadrants the two hospitals’ stocks compete for the same over-demanded arrivals—locally analogous to the trade-off of the SP model—while in cross-type quadrants the stocks can be matched with each other, with over-demanded arrivals as outside opportunities—locally analogous to the trade-off of the CP model. Classifying two stocks as complements when they can be matched with each other and as substitutes when they compete for a common third type is consistent here, unlike in a general three-type example in which the same pair of types could qualify as both: the compatibility graph is bipartite, so A – B pairs (at either hospital) are complementary to B – A pairs, same-type stocks at different hospitals are substitutes, and no pairwise relation conflicts. This yields a sharp local prediction for the sign of every cross-difference

of the value function. The argument is heuristic—the full-model Bellman operator moves the state across quadrants, so local pairwise relations do not by themselves imply global curvature—and Section 4.10 therefore tests the prediction numerically, confirming it at every interior grid state.

The numerical analysis below first presents results for the CP and SP benchmarks—directly visualizing the threshold policies of Theorems 1 and 2—and then solves the constrained full model.

4 Numerical Analysis

4.1 Calibration and computational approach

Immediate surpluses throughout this section are those of the full model (1): intra-hospital two- and three-way exchanges generate 2 and 3, and their inter-hospital counterparts $2 - c$ and $3 - c$; over-demanded arrivals carry the baseline two-way surplus 2 in every branch, so the incremental rewards are 1 (own three-way), $1 - c$ (cross three-way), and 0 (pass), matching the abstract benchmark rewards via Table 3. Also, v^* denotes the optimal value function of the *constrained full model*: the full multi-hospital matching environment of Section 3.1 under the immediate intra-hospital clearing restriction (Assumption 8). The numerical analysis solves this model in net-position form, with state $x = (x_1, x_2) \in \{-50, \dots, 50\}^2$, where $x_i := s_i^{AB} - s_i^{BA}$ is the net reciprocal imbalance at h_i . Under the immediate intra-hospital clearing rule of Assumption 8 this reduction is exact (Lemma 1), and the net-position solver enforces the assumption by construction: retention and intra-hospital clearing induce the same transition in x , the solver credits the exchange surplus, and a coexistence state cannot even be represented—so every figure and table in this section is computed under Assumption 8. It is a restriction, not an optimality property: solving the unrestricted four-dimensional model at the same calibration, retaining an arriving pair alongside its stored reciprocal is strictly optimal at some depleted states, with a worst-case action gap of 0.13; the welfare cost of imposing clearing stays below 0.4% of the value at empty pools for every reported c (Remark 2, Appendix B). All results in this section characterize the constrained optimum within the immediate-clearing class. The sign pattern of (x_1, x_2) partitions the state space into four quadrants that organize the entire analysis: $x_1 x_2 < 0$ (Quadrants II and IV) corresponds to the CP regime and $x_1 x_2 > 0$ (Quadrants I and III) corresponds to the SP regime. The numerical solutions use discount factor $\delta = \lambda/(\lambda + \rho) = 0.999$, where λ is the aggregate arrival rate. The numerical exercise is a *literature-based benchmark calibration*: every parameter is taken from, or derived from, values reported in the transplantation and kidney-exchange literatures, rather than estimated from registry microdata.¹⁰ Table 4 summarizes the parameter values, their sources or construction, and their economic roles; the derivation of the over-demanded weights and a robustness check appear in the table notes. The conditional event probabilities implied by the weights, $p_{A-B} = 0.4568$, $p_{B-A} = 0.2741$, $p_{A-B}^O = 0.0749$, $p_{B-A}^O = 0.1943$, sum to one and preserve $p_{A-B}/p_{B-A} = 5/3$; reciprocal types account for 73.1% of tracked events.

¹⁰The model's key input is the arrival composition of exchange-pool pair types—the reciprocal $A-B$ and $B-A$ shares and the over-demanded shares net of crossmatch. For the United States this composition is not published in the public OPTN/SRTR releases, and patient-level registry data require a data-use agreement. It is, however, an *aggregate*: a donor-by-recipient blood-group table of the pairs entering an exchange pool, which some national programmes publish directly. Section 4.8 therefore re-solves the model under a second calibration constructed from such a published table, and the robustness check reported below indicates that the qualitative findings are insensitive to moderate perturbations of the weights. Estimating the weights from U.S. patient-level data, together with the explicit waiting cost κ of Section 3.4, is left to future work.

Parameter	Baseline value	Source / construction	Economic role
Reciprocal weights w_{A-B}, w_{B-A}	0.05, 0.03	Terasaki et al. [1998] , as cited and used by Ünver [2010]	Arrival rates of the two state variables; the 5:3 ratio drives the abundant–scarce asymmetry
Over-demanded weights $w_{O\mathcal{D}_{A-B}}, w_{O\mathcal{D}_{B-A}}$	0.0082, 0.0213	Derived from U.S. blood-type frequencies and a positive-crossmatch probability (notes)	Arrival rates of the three-way upgrade opportunities
$AB-O$ split	equal ($\frac{1}{2}, \frac{1}{2}$)	Assumption 4	Assigns $AB-O$ arrivals to the two over-demanded classes
Discount factor δ	0.999	Chosen; cf. 0.999995 in Ünver [2010] (see <i>Temporal calibration</i> below)	Patience; option value of waiting
Coordination cost c	varied over [0, 2]	Comparative-statics parameter, not estimated	Per-exchange resource cost of inter-hospital coordination
State space	$ x_i \leq 50$ (net position); stock cap 14 (4D check)	Computational	Truncation; robustness to the cap in Section 4.3

Notes: Over-demanded weights are derived from U.S. blood-type frequencies $(q_O, q_A, q_B, q_{AB}) = (0.45, 0.40, 0.11, 0.04)$ under blood-type-independent pairing, with ABO-compatible pairs entering the pool through a positive crossmatch of probability $p_c = 0.11$ [[Zenios et al., 2001](#)]: $w_{O\mathcal{D}_{A-B}} = p_c(q_B q_O + q_{AB} q_A + \frac{1}{2} q_{AB} q_O) = 0.0082$ and $w_{O\mathcal{D}_{B-A}} = p_c(q_A q_O + q_{AB} q_B + \frac{1}{2} q_{AB} q_O) = 0.0213$, splitting $AB-O$ arrivals equally. Re-solving all models with the weights of an earlier draft, (0.0635, 0.0381, 0.00626, 0.021), changes the CP benchmark policy at 6 of 2601 grid states, the SP benchmark policy at none, and leaves every reported sign pattern and threshold structure unchanged.

Table 4: Literature-based benchmark calibration.

Economic units and coordination-cost range. Because surpluses are measured in transplant-equivalent units, c has a direct empirical counterpart:

$$c = \frac{\text{real resource cost of one inter-hospital exchange}}{\text{social value of one additional transplant}}.$$

Both sides can be bracketed. For the denominator, [Held et al. \[2016\]](#) estimate the social welfare gain of a kidney transplant at roughly \$1.1 million per recipient (6.5 additional quality-adjusted life years relative to dialysis, plus taxpayer savings of about \$146,000). For the numerator, the incremental resources that an inter-hospital exchange consumes over an intra-hospital one are organ packaging and shipping, courier and logistics services, and the administrative burden of cross-institutional synchronization; [Lentine et al. \[2026\]](#) document that U.S. programs manage the cost components of paired donation—registry fees, donor evaluation, nephrectomy billing, and organ shipping—in widely heterogeneous ways, and facility-level cost variation of the order of \$10,000 per case is documented even within single components [[Mast et al., 2011](#)]. Not every documented component belongs in this numerator, which must measure the *incremental* resource cost of an inter-hospital exchange relative to an intra-hospital one: donor evaluation and the nephrectomy itself are incurred wherever the exchange takes place, and registry fees attach to platform participation rather than to a particular cross-hospital match. The incremental numerator therefore consists of organ packaging and shipping, the additional coordination and scheduling labor of synchronizing two centers, and incremental regulatory or administrative filings; billing arrangements must further be separated from resource use, since costs billed between hospitals (42 CFR §413.416) shift incidence without

changing the social cost. A transparent decomposition of these increments is not available in the literature, so we read the bracket that follows as an *illustrative health-economic range* rather than an estimate. Taking a deliberately generous incremental numerator of \$50,000 together with a conservative denominator of \$500,000 yields $c = 0.1$; plausible midpoints (a few thousand dollars of shipping and coordination labor against \$1.1 million) yield c below 0.01. On this illustrative bracket, c lies in roughly $[0.005, 0.1]$ —at the bottom of the range we report. Larger values of c should be read as stress tests—or as descriptions of settings where coordination is institutionally much harder than its direct resource cost suggests—rather than as central estimates, and the numerical analysis accordingly treats $c = 0.1$ as a generous benchmark within the illustrative range and reports the full range for completeness.

Temporal calibration. Three forces are folded into the reduced-form ρ : pure social time preference (conventionally 3–5% per year), the hazard of patient attrition through death or delisting (several percent per year on kidney waitlists), and the patient-welfare cost of waiting itself. A defensible composite rate is therefore of order $\rho \approx 0.08$ –0.15 per year once the waiting cost is carried explicitly by κ (Proposition 2 and Section 4.11). Our computations use $\delta = \lambda/(\lambda + \rho) = 0.999$, which at the illustrative $\lambda = 1,000$ arrivals per year corresponds to $\rho \approx 1$ per year—an intentionally conservative computational horizon, not an empirical discount claim: it discounts future matching opportunities far more heavily than the defensible range, which at the same λ implies $\delta \approx 0.9999$. (For comparison, Ünver [2010] calibrates $\lambda/(\lambda + \rho) = 0.999995$, at $\lambda = 10,000$ pairs per year and $\rho = 0.05$; five-nines discounting requires on the order of 10^6 value-iteration steps at our tolerances and makes the truncated state space bind.) Because the paper’s structural findings operate through option values, heavier discounting works *against* them—retention regions and positive thresholds emerge despite, not because of, patience—though not necessarily for every policy-region or comparative-static statement; the structural properties of Lemmas 2–3 hold for every $\delta \in (0, 1)$, only threshold levels shift, and the immediate-clearing restriction of Assumption 8 binds only for δ near one (Remark 2). Section 4.11 verifies the conservativeness claim directly: for every computed $\delta \geq 0.999$ the structural findings are unchanged and the symmetric-hospital welfare gains only shrink (Table 10); the same sensitivity analysis varies δ jointly with the per-event holding cost θ , since a thinner market changes both the discounting between events and the waiting cost per event.

The value function is computed via value iteration until the sup-norm difference falls below 10^{-9} (10^{-12} for the benchmark verification runs).

The benchmark panels condition on the event classes of the corresponding quadrant model and renormalize their probabilities (for the SP model the two same-type classes are symmetric); opportunity rewards enter incrementally, with the baseline two-way surplus carried in every branch, consistent with Table 3; the state grid is capped at 50 with a reflecting upper boundary. The section proceeds as follows. Section 4.2 examines the structural properties of the two benchmark models in a single figure. The next three subsections then report the paper’s three substantive findings: the active range of coordination costs (Section 4.3); an asymmetry between abundant and scarce reciprocal types (Section 4.4); and the result that rebalancing via inter-hospital matches is confined to a narrow deep-imbalance region in same-type states (Section 4.5). Section 4.6 converts these state-space patterns into economically weighted magnitudes along the optimal path; Section 4.7 characterizes when coordination becomes quantitatively important by relaxing hospital symmetry;

Section 4.8 re-solves the model under a second calibration taken from a national programme’s published pool composition; Section 4.9 checks the arrival-triggered timing convention against batched clearing; Section 4.10 provides a direct numerical check of supermodularity and submodularity in the appropriate quadrants of the constrained full model; and Section 4.11 closes by recomputing the headline results under a calibrated waiting cost, varying the discount factor jointly with it. Threshold-vs- c comparative-static plots for the benchmark models are deferred to the Online Appendix (Section XIII).

4.2 Benchmark models: policy regions and marginal values

Figure 1 shows the policy regions and marginal value structure of the two benchmark models at $c = 0.1$. The top row reports the CP model. The left panel partitions the state space into the four policy regions of Theorem 1: “keep / two-way” below both thresholds, “exchange” above $\bar{t}_{XY}$, “use opportunity” above $\bar{t}_{OX}$, and “both” above both. The right panel plots the joint marginal value $\Delta_{1+2}v^*(s) := v^*(s) - v^*(s - e_1 - e_2)$, which is the cost of a two-way exchange consuming one X pair and one Y pair. The white level-set curve traces the exchange threshold $\bar{t}_{XY}$: exchange is optimal precisely when the joint marginal value falls below $a_{XY} = 2 - c$. Near the origin the joint marginal value is high (warm colors); as both pools grow it decreases (cool colors), reflecting concavity. Numerically, the cross-difference $\Delta_1\Delta_2v^*$ is positive at every interior grid point (minimum $+1.2 \times 10^{-7}$ at a value-iteration tolerance of 10^{-12}), numerically consistent with the supermodularity established in Lemma 2. Quantitatively, the four regions partition the 51×51 grid as follows at $c = 0.1$: both actions are optimal at 94.3% of grid states, only the opportunity at 3.7%, only the exchange at 0.2%, and neither—the keep/two-way region—at 1.9%, concentrated near the origin. The dominance of the “both” region reflects how low the calibrated thresholds are; the threshold variation of economic interest is confined to the near-origin band, which the inset in the policy panel magnifies.

The bottom row reports the SP model. The left panel shows the three-region policy of Theorem 2 for an O_1 arrival (the O_2 case is symmetric): “not using” the opportunity below both thresholds, matching from the partner pool in the narrow band $\tau_2 \leq s_1 < \tau_1$, and matching from the own pool once $s_1 \geq \tau_1$. Because type X_1 pairs arrive frequently under the calibrated parameters, the own-type threshold τ_1 binds at $s_1 = 1$ over essentially the entire state space: the own-pool region covers 98.0% of grid states, the cross-pool band 1.9%, and inaction only the empty state. The right panel plots the marginal value $\Delta_1v^*(s)$, which is nonincreasing in both s_1 (concavity in the own pool) and s_2 (submodularity across pools). The cross-difference is non-positive in the interior (minimum -9.6×10^{-5} ; maximum 0 to machine precision), numerically consistent with the submodularity established in Lemma 3. Together, the two rows visualize the opposite curvature regimes that the rest of this section traces through the constrained full model.

4.3 Finding 1: The active range of coordination costs

The first substantive finding is that inter-hospital cooperation collapses sharply as the coordination cost c rises. Figure 2 shows the optimal action upon a reciprocal-pair arrival at $c = 0.1$; Section V of the online appendix repeats the map at the other cost extreme, $c = 2.0$ (Figure OA.2). Red = two-way exchange, blue = inter-hospital three-way exchange, green = intra-hospital three-way exchange.

At $c = 0.1$, the inter-hospital channel is active: blue regions appear in the lower-left of Quadrant III at h_1 and in the lower-left of Quadrant IV at h_2 , where one hospital holds a large excess and an




Figure 1: Benchmark models at $c = 0.1$. Top row: CP model. Left panel shows the four policy regions; right panel shows the joint marginal value $\Delta_{1+2}v^*$ with the exchange threshold $\bar{\tau}_{XY}$ as a white level set. Bottom row: SP model. Left panel shows the three-region policy with thresholds τ_{11} (white, own-type) and τ_{12} (dashed, cross-type), matching the notation of Theorem 2; right panel shows the own-type marginal value Δ_1v^* . In each panel, the inset magnifies the near-origin band where the thresholds vary; the outlined rectangle marks the magnified region. The positive cross-difference in the CP model is numerically consistent with the supermodularity established in Lemma 2; the negative cross-difference in the SP model is numerically consistent with the submodularity of Lemma 3.

inter-hospital three-way exchange becomes worthwhile. The intra-hospital three-way exchange (green) dominates the half-plane in which the local hospital holds excess of the matching reciprocal type, while two-way (red) prevails in the half-plane where it does not. This asymmetry reflects the arrival ratio $p_{A-B}/p_{B-A} = 5/3$: when one reciprocal type arrives more frequently, the abundant-type hospital's sharing threshold is effectively zero, while the scarce-type hospital sets a strictly positive sharing threshold to preserve future matching opportunities, consistent with the dynamic-matching framework of Ünver [2010].

At $c = 2.0$, inter-hospital three-way exchange yields net surplus $3 - c = 1$, well below the intra-hospital surplus of 3, and the inter-hospital region vanishes entirely: hospitals revert to autarky (online appendix Figure OA.2).

The transition between the two regimes is monotone in c but has a sharp *activation* structure. At $c = 0$ the planner treats the two hospitals as a single pool and achieves, within the immediate-clearing class, the welfare of a fully centralized clearinghouse. As c rises, the interior inter-hospital trigger region contracts monotonically—from 5,748 grid states at $c = 0$ to 3,980 at $c = 0.1$ and




Figure 2: Reciprocal-channel policy at $c = 0.1$: h_1 (left) and h_2 (right). Red = two-way exchange, blue = inter-hospital three-way, green = intra-hospital three-way. Panels show the OD action aligned with the local excess. The corresponding maps at $c = 2.0$, where the inter-hospital region vanishes entirely, are in the online appendix (Figure OA.2).

114 at $c = 1.0$ —and is empty on this window for every $c \geq 1.1$.¹¹ The OD channel shuts down analytically at $c = 1$, where its incremental surplus $1 - c$ becomes nonpositive; the reciprocal channel survives only marginally longer because, at $\delta = 0.999$, waiting for a costless intra-hospital match is cheap, so a cross-hospital surplus of $2 - c$ must be close to 2 to dominate the option of waiting. Within the reporting window, cooperation is therefore extinguished near the incremental three-way surplus ($c \approx 1$), well below the two-way ceiling $c = 2$ at which the surplus would be exhausted mechanically, and the contraction is fastest at low c : the step from $c = 0$ to $c = 0.1$ removes 1,768 states, more than any later step. The extensive-margin comparative static is accordingly: reductions in the real cost expand the set of cooperation states only once they bring c into the active range $c \lesssim 1$. Within that range, the geometric cooperation region expands most rapidly at low coordination costs, although the occupation-measure calculations in Section 4.6 show that the corresponding on-path welfare effect is small under symmetric arrivals. The illustrative bracket introduced at the start of this section, $c \in [0.005, 0.1]$, lies inside this active low-cost region. The shutdown point itself is composition-dependent: for hospital pairs whose arrival mixes generate persistent complementary imbalances, cooperation remains valuable at costs where it has shut down under symmetry (Section 4.7). Grid-state counts weight all pool configurations equally and are therefore statements about the policy’s geometry, not about welfare incidence; Section 4.6 reports how often these states are actually visited and what the cooperation they support is worth. This comparative static concerns policies that reduce the resource cost itself (logistics, shared transport infrastructure, organ-shipment technology); reimbursement only shifts the incidence of c and cannot move the first-best frontier (Sections 1 and 5).

¹¹ Counts refer to the window $|x_i| \leq 45$ of the 101×101 grid. They are robust to the grid cap, and the emptiness statement is a window statement: trigger states do exist at deeper imbalances, but they carry essentially no discounted occupation mass from empty pools (Section 4.6). Section XI of the online appendix reports the cap-robustness series (caps ± 40 – ± 75), the deep-imbalance counts, and the replication routines.

4.4 Finding 2: Abundant–scarce asymmetry

The second finding is a striking asymmetry in the inter-hospital OD channel: the optimal policy uses the channel aggressively for the abundant reciprocal type, while for the scarce type it is inactive at every evaluated grid state under the baseline parameterization, even when the immediate net surplus is strictly positive.

Two distinct inter-hospital channels operate in the model. The *reciprocal channel*, examined in Section 4.3, exchanges opposite-type reciprocal pairs directly across hospitals (surplus $2 - c$). The *OD channel* routes an arriving over-demanded pair to consume a reciprocal pair at the other hospital in a three-way exchange (surplus $3 - c$). Figure 3 decomposes the OD-channel policy by OD type and hospital.

The left panel (OD_{A-B} , consumes an $A-B$ pair from the *abundant* type) is dominated by green and blue. Whenever the partner hospital holds excess of the abundant type, the planner sends the arriving OD_{A-B} across hospitals to consume those abundant pairs: at $c = 0.1$ the inter-hospital three-way exchange (blue) is optimal at essentially 100% of the relevant cross-type states. The right panel (OD_{B-A} , consumes a $B-A$ pair from the *scarce* type) is by contrast almost entirely red. Even when inter-hospital three-way exchange is feasible and immediately profitable—for example, OD_{B-A} arriving at h_1 in Quadrant IV with h_2 holding excess $B-A$, where $3 - c = 2.9 > 2$ —the planner exercises this option at none of the evaluated grid states (0%).




Figure 3: Optimal policy at $c = 0.1$ decomposed by OD type, for arrivals at h_1 . Left panel (OD_{A-B} , abundant $A-B$ type): inter-hospital channel active. Right panel (OD_{B-A} , scarce $B-A$ type): inter-hospital channel inactive at every evaluated state. Red = two-way (OD + UD), blue = inter-hospital three-way, green = intra-hospital three-way. Because the arrival process is symmetric across hospitals, the h_2 panels are mirror images; the full four-panel decomposition is in Section VI of the online appendix (Figure OA.3).

Table 5 quantifies this asymmetry. The inter-hospital OD channel for the abundant type is used at nearly every relevant grid state at low cost, drops to zero once net surplus turns non-positive at $c = 1$, and is shut down at $c = 2$. For the scarce type, the channel is inactive at every evaluated grid state for every cost value considered; along the optimal path from empty pools its discounted share of events is likewise exactly zero at every reported c (Section 4.6). The planner is dynamically reluctant to consume a scarce pair across hospitals because its option value—its future contribution to within-hospital three-way matches—exceeds the immediate incremental surplus.

	$c = 0.1$	$c = 0.5$	$c = 1.0$	$c = 2.0$
OD _{A-B} at h_1 (abundant, QII)	99.9%	97.6%	0.0%	0.0%
OD _{B-A} at h_1 (scarce, QIV)	0.0%	0.0%	0.0%	0.0%

Table 5: Fraction of cross-type grid states at which inter-hospital three-way exchange is optimal, by OD type and coordination cost. “Cross-type” refers to the quadrant in which the matching reciprocal type is present only at the partner hospital. Rows report OD arrivals at h_1 ; because the arrival process is symmetric across hospitals, the corresponding h_2 cases are mirror images (cf. Figure 3).

The economic message is clear. Inter-hospital coordination is not symmetric across pair types: it is a tool for moving abundant pairs across institutions, not for redistributing scarce ones. This first-best asymmetry identifies a feature that any future analysis of decentralized implementation would need to reproduce.

4.5 Finding 3: Just-in-time coordination and the limits of pool smoothing

A natural conjecture is that a forward-looking planner would use inter-hospital matches to smooth cross-hospital asymmetries in pool composition, sinking coordination costs today to economize on coordination tomorrow. Our third finding is that, within the immediate-clearing class of Assumption 8, this motive is almost entirely absent: in cross-type states the optimal policy never smooths, and a rebalancing motive survives only in a narrow deep-imbalance region of the same-type quadrants at low c .

Figure 4 shows the inter-hospital reciprocal two-way trigger region at $c = 0.1$. Red indicates states at which an arriving reciprocal pair is immediately exchanged across hospitals; gray is the continuation region. The trigger region is concentrated in Quadrants II ($x_1 < 0, x_2 > 0$) and IV ($x_1 > 0, x_2 < 0$), where hospitals hold opposite excesses. Its boundary is convex, resembling the convex exchange threshold $\bar{t}_{XY}(s_2)$ that Theorem 1 establishes for the CP benchmark and consistent with the supermodularity and concavity of Lemma 2; the full-model frontier itself is a computed object, not a theorem. The asymmetry between the two complementary quadrants again reflects $p_{A-B}/p_{B-A} = 5/3$: in Quadrant II the trigger region extends further because the scarce $B-A$ pairs at h_1 are more valuable for exchange.

The key observation concerns reachability. Starting from empty pools and following the optimal policy, the discounted occupation measure (Section 4.6) concentrates the trigger set’s on-path mass on its frontier—the trigger states with at least one non-trigger neighbor: at $c = 0.1$ the 408 frontier states carry discounted mass 3.2×10^{-3} , against 2.1×10^{-4} for the 4,632 interior states, so 94% of the trigger set’s mass sits on its boundary layer, and the interior is entered only transiently, when a single arrival jumps the state past the frontier. Inter-hospital reciprocal exchange is performed essentially *at* the boundary, preempting entry into the convex trigger set, rather than *within* it as a corrective action. The frontier therefore acts as a just-in-time threshold, not a smoothing pool composition tool. Combined with the asymmetry of Section 4.4, this delivers a clean economic interpretation: in cross-type states, coordination costs incurred today are not optimally sunk to avoid coordination costs tomorrow.

The limits of no-smoothing. The no-smoothing conclusion is exact in the cross-type quadrants but not global. In the same-type Quadrant III (both hospitals hold $B-A$ excess), when an $A-B$ pair




Figure 4: Inter-hospital reciprocal two-way trigger region at $c = 0.1$. Red = exchange, gray = continuation. The convex frontier acts as a preemptive boundary: starting from empty pools, the optimal policy exchanges at the frontier, so the interior of the trigger region carries almost no discounted occupation mass (see text).

arrives at the hospital with the deeper pool, matching it *across* hospitals—consuming the thinner hospital’s B – A pair at surplus $2 - c$ —strictly dominates the free intra-hospital match at 348 grid states at $c = 0.1$ (277 of them in the trimmed interior, concentrated at $x_1 \in [-32, -3]$, $x_2 \in [-15, -1]$, with the mirror region symmetric), with a payoff advantage of up to 0.23 and a value gap of up to 0.68. The two magnitudes compare different objects: the *payoff advantage* is the action-value gap at a single state—the Bellman value of the cross-hospital match minus that of the best alternative action at the same state—while the *value gap* is the sup-norm difference between the value functions computed with and without the cross-hospital action in the choice set. The planner is paying c to concentrate depletion in the thinner pool: emptying it hastens the system’s transition into the complementary configuration, in which the partner hospital’s abundant A – B inflow begins clearing the deep B – A stock through the reciprocal channel. The region shrinks rapidly in c (29 boundary states at $c = 0.5$, none at $c = 1$) and leaves the rest of the analysis intact: the trigger region of Figure 4 changes at 4 of 10,201 cells, the OD-channel shares of Table 5 are unchanged, and the curvature verification of Section 4.10 is unaffected. On path the exception is rare as well as narrow: its discounted share of events from empty pools is 0.030% at $c = 0.1$ and zero at $c = 0.5$ (Section 4.6). Pool smoothing is thus not a general feature of the optimal policy, but neither is it categorically absent: it appears exactly where deep same-type imbalance meets cheap coordination. One qualification applies throughout: these timing statements are made under arrival-triggered clearing (Assumption 7), whereas operating programmes often clear in periodic matching runs. Section 4.9 examines batched clearing directly in a benchmark exercise and finds that the retention mechanism underlying the just-in-time pattern survives periodic clearing; for the full model, Finding 3 should nonetheless be read as conditional on the arrival-triggered convention.

4.6 On-path magnitudes

Grid-state counts describe the policy’s geometry but weight all pool configurations equally, however rarely they occur. To convert the geometry into economic magnitudes, we compute the exact discounted occupation measure of the computed finite-state chain: starting from empty pools and

following the computed optimal policy, each state receives weight proportional to the discounted expected number of visits (normalized by $1 - \delta$ so the weights sum to one), obtained from a sparse linear system rather than by simulation. Weighting each arrival type’s action by this measure yields the discounted share of arrivals resolved by each action; evaluating the policy’s reward stream—and, separately, its transplant count—by the same linear method yields welfare and transplant levels, which we compare with a counterfactual *autarky* planner restricted to intra-hospital actions only. Table 6 reports the results.

	$c = 0.1$	$c = 0.5$	$c = 1.0$
Inter-hospital share of arrivals	0.300%	0.238%	0.000%
via reciprocal two-way exchange	0.066%	0.000%	0.000%
of which: same-type rebalancing exception	0.030%	0.000%	0.000%
via OD three-way, abundant type	0.234%	0.238%	0.000%
via OD three-way, scarce type	0.000%	0.000%	0.000%
Welfare gain over autarky	+0.15%	+0.07%	0.00%
Transplant gain over autarky	+0.18%	+0.17%	0.00%

Notes: Discounted occupation measure from empty pools at $\delta = 0.999$ under the constrained full model’s computed optimal policy; shares are percentages of all arrivals (reciprocal and over-demanded). Autarky resolves the model with inter-hospital actions removed; gains are percentages of the autarky value at empty pools. The evaluated policies reproduce the value-iteration solutions to 2×10^{-7} , and the $c = 2$ solution coincides with autarky, as it should.

Table 6: On-path magnitudes of inter-hospital cooperation under the baseline symmetric calibration.

Three magnitudes stand out. First, under this calibration inter-hospital cooperation is rare on path: 0.300% of arrivals at $c = 0.1$, falling to zero by $c = 1$ —the on-path counterpart of the active range of Section 4.3, and the reason the deep-imbalance trigger states discussed there are economically irrelevant at high c . Second, the composition of that cooperation is exactly the asymmetry of Section 4.4: the OD channel operates only through the abundant type, the scarce-type share is zero at every reported c , and the same-type rebalancing exception (a subset of the reciprocal two-way row) is an order of magnitude smaller still. Third, cooperation’s welfare value at low cost is positive but modest under this calibration—+0.15% of the autarky value at $c = 0.1$, roughly two additional transplants per thousand arrivals—and declines to zero as c approaches the active-range boundary.

Nor is the coordination cost itself the binding constraint here: re-solving the symmetric model at $c = 0.05, 0.01$, and 0 raises the welfare gain over autarky only to +0.17%, +0.18%, and +0.18% (inter-hospital shares 0.36%, 0.43%, and 0.88% of arrivals). Even *costless* coordination is worth less than a fifth of one percent to hospitals with identical arrival distributions. Given that the illustrative cost bracket places c at the bottom of our range, what governs the value of coordination must be something other than c ; Section 4.7 shows that it is the composition of arrivals across hospitals.

Three remarks discipline the interpretation. These magnitudes illustrate the coordination margin for hospitals with *similar* arrival distributions; Section 4.7 characterizes how the margin changes as that similarity is relaxed. Second, the levels also inherit the state-space and no-rebalancing features of the constrained model; they quantify the cooperation margin within the immediate-clearing class, consistent with the rest of this section. Third, the empty pool is one initial condition among many. Re-reading the policy-evaluation vectors at nonempty starts, the welfare gain over autarky ranges from +0.09% (same-side excess, $x_1 = x_2 = 3$, where inventories are substitutable and cooperation has little to add) to +0.34% (opposite-side excess, $x_1 = -x_2 = 8$, the complementary configuration) across

initial states with $|x_i| \leq 8$ —the ordering the composition mechanism predicts. Under a long burn-in of the symmetric chain the pools drift into same-side excess within the computed grid, where the incremental value of cooperation is approximately zero; a mature symmetric programme thus gains, if anything, less than the empty-pool benchmark suggests, which is again the composition logic of Section 4.7 rather than an artifact of the empty start. Because the unbounded controlled chain need not possess a stationary distribution, these burn-in figures are descriptive discounted-occupation calculations from specified initial states, not stationary-distribution estimates.

4.7 When does coordination matter? The role of compositional asymmetry

The symmetry of the baseline calibration is a substantive restriction: both hospitals draw arrivals from the same distribution, so persistent complementary imbalances—one hospital long on $A-B$ while the other is long on $B-A$ —never form, and the complementary quadrants that the CP benchmark describes are visited only transiently. This subsection relaxes exactly that restriction. The formulation of Section 3.1 already allows hospital-specific arrival probabilities; we parameterize cross-hospital compositional asymmetry by $a \in [0, 1]$, tilting reciprocal arrivals while holding the aggregate type distribution fixed:

$$q_1^{A-B} = \frac{1+a}{2} p_{A-B}, \quad q_2^{A-B} = \frac{1-a}{2} p_{A-B}, \quad q_1^{B-A} = \frac{1-a}{2} p_{B-A}, \quad q_2^{B-A} = \frac{1+a}{2} p_{B-A},$$

with over-demanded arrivals kept symmetric so that the reciprocal channel is isolated. At $a = 0$ the hospitals are identical (the baseline); at $a = 1$ hospital h_1 receives all $A-B$ and h_2 all $B-A$ reciprocal arrivals—for example, two centers drawing on donor populations with mirrored blood-group compositions. Because $p_{A-B} \neq p_{B-A}$, this tilt varies more than composition: hospital h_1 's share of all events is $\frac{1}{2} + \frac{a}{2} (p_{A-B} - p_{B-A})$, which rises to 59.1% at $a = 1$, so the sweep moves hospital size together with hospital composition. We therefore read it as a full-specialization stress experiment and add a size-neutral exercise below that isolates composition alone. For each (a, c) we re-solve the model, recompute the discounted occupation measure from empty pools, and compare with an autarky planner facing the *same* asymmetric arrivals. Panel A of Table 7 reports the results.

Three features stand out. First, the gains rise steeply and convexly in a : from +0.15% under symmetry to +3.2% at mild asymmetry, +12.3% at $a = 0.5$, and +68%—with 73% more transplants—in the full-specialization stress experiment, which we read as the stress-test endpoint of this experiment rather than a forecast of gains from real hospital cooperation. Persistent complementary imbalance, not the level of the coordination cost, is what makes cross-hospital exchange valuable; this is the CP benchmark's configuration realized in an institutional environment, and it connects the comparative static directly to observable heterogeneity in donor populations across centers. Second, the *active range* of Section 4.3 is itself composition-dependent: under symmetry, cooperation is worthless by $c = 1$, but at $a = 1$ the reciprocal channel still yields +26% at $c = 1.0$ (online appendix Table OA.3), because with persistent imbalances the surplus $2 - c$ is harvested at scale rather than occasionally. Third, symmetric variations do not substitute for asymmetry: re-solving the symmetric model with the reciprocal share of arrivals moved from 73% to 49%, 62%, or 84% changes the gain only within [+0.08%, +0.28%], and moving the reciprocal ratio from 5:3 to 1:1 or 3:5 (holding the share fixed) keeps it within [+0.06%, +0.40%]—each interacting with the over-demanded mix in sign-ambiguous ways, but always an order of magnitude below even mild cross-hospital asymmetry.¹²

¹²The routine the asymmetry-sweep routine in the replication code reproduces the table and both sweeps. The share

<i>Panel A: hospital-specific tilt a (hospital sizes vary with composition), $c = 0.1$</i>				
Asymmetry a		Welfare gain	Transplant gain	Inter share
0.00		+0.15%	+0.18%	0.30%
0.25		+3.2%	+3.6%	4.1%
0.50		+12.3%	+13.6%	12.9%
0.75		+33.3%	+35.9%	22.0%
1.00		+68.0%	+72.6%	31.1%
<i>Panel B: size-neutral tilt (each hospital receives exactly half of all arrivals), $c = 0.1$</i>				
Tilt	f_1 vs. f_2	Welfare gain	Transplant gain	Inter share
0.00	62.5% / 62.5%	+0.15%	+0.18%	0.30%
0.25	71.9% / 53.1%	+1.6%	+1.8%	2.3%
0.50	81.2% / 43.8%	+7.2%	+8.0%	8.5%
0.75	90.6% / 34.4%	+16.3%	+17.8%	15.2%
1.00	100% / 25%	+33.4%	+35.9%	22.0%

Notes: Gains are relative to an autarky planner facing the same arrival process, at empty pools; inter share is the discounted share of arrivals resolved by an inter-hospital action. Panel A varies the hospital-specific tilt a ; the $a = 0$ row reproduces Table 6, and the $a = 1$, $c = 0.1$ entries are unchanged under grid cap ± 75 . Panel B holds each hospital's total arrival rate at exactly one half and tilts only the within-hospital reciprocal mix (f_i is hospital h_i 's A – B share among its reciprocal arrivals); the aggregate type distribution is unchanged. Panel A's welfare gains at $c = 0.5$ and $c = 1.0$ are in Section VII of the online appendix (Table OA.3).

Table 7: Welfare value of cooperation under cross-hospital compositional asymmetry: the a -sweep (Panel A) and the size-neutral tilt (Panel B).

The size-neutral exercise confirms that composition, not size, carries the effect. We hold each hospital's total arrival rate at exactly one half—over-demanded arrivals symmetric, hospital h_1 's reciprocal arrivals keeping mass $(p_{A-B} + p_{B-A})/2$ —and tilt only the mix: writing f_i for the A – B share of hospital h_i 's reciprocal arrivals, the symmetric point is $f_1 = f_2 = 62.5\%$, and the maximal tilt feasible at equal sizes is $f_1 = 100\%$ against $f_2 = 25\%$. (Full two-sided specialization is unattainable at equal sizes precisely because the aggregate ratio is 5:3—the reason the a -sweep moves hospital sizes.) Panel B of Table 7 reports the sweep at $c = 0.1$.

The steep, convex gradient survives intact: the maximal size-neutral tilt yields +33%, on the order of the a -sweep at $a = 0.75$, so the gains of Panel A are driven by compositional asymmetry itself rather than by the accompanying size variation. Two parameterizations should not be conflated when reading these magnitudes against data. In the a -sweep, $(1 + a)/2$ is the share of *aggregate* A – B arrivals allocated to h_1 ; the observable statistic—the A – B share *within* h_1 's reciprocal arrivals—is $(1 + a) p_{A-B} / [(1 + a) p_{A-B} + (1 - a) p_{B-A}]$, which rises from 62.5% at $a = 0$ to 73.5% at $a = 0.25$ and 100% at $a = 1$. The size-neutral tilt maps to that statistic directly: at tilt 0.25 the within-hospital shares are 71.9% versus 53.1% (Table 7). As for what values are realistic: published national aggregates demonstrate that reciprocal-type composition varies substantially across settings—the U.S.-based construction implies an A – B : B – A ratio of 5:3, while the UK scheme's registrations run 52:80, an inversion (Section 4.8)—but they do not identify the distribution of within-programme hospital asymmetry. We therefore interpret the intermediate tilts as mechanism-revealing scenarios rather than estimated hospital-pair differences, and the full-specialization row as a stress-test endpoint.

The broader numerical finding of this section is therefore joint: the value of coordination depends both on the reciprocal-to-opportunity composition of arrivals—which determines whether recip-

and ratio sweeps hold hospitals symmetric and renormalize the event distribution; their non-monotone effects reflect the interaction between reciprocal supply and the over-demanded opportunities that consume it, as in Section 4.8.

rocal inventory accumulates at all (Section 4.8)—and on whether hospitals’ arrival compositions generate *persistent complementary imbalances*, which determine whether the accumulated inventory is on the same side of the market or on opposite sides. Costly coordination is quantitatively important exactly when both conditions hold; policies that reduce c are most valuable for compositionally dissimilar hospital pairs, and nearly irrelevant for similar ones.

4.8 A contemporary registry-based calibration

The baseline weights are constructed from reported frequencies rather than measured on a live exchange pool. As a check on that construction, we re-solve the model under a second calibration taken directly from a national programme’s published statistics. [NHS Blood and Transplant \[2025\]](#) reports, for the UK Living Kidney Sharing Scheme, the number of pairs *registered* in the scheme between 1 April 2021 and 31 March 2025 by donor and recipient blood group ($n = 1,406$). That table is an aggregate—no patient-level data are involved—and it is the closest publicly available proxy for the object our arrival process needs, the blood-group composition of the pairs entering an exchange pool; the caveats closing this subsection explain in what respects registrations and arrivals differ. Mapping its cells into the pair-type classes of Table 1 and the over-demanded classes of Section 3 yields the conditional event probabilities in the right-hand column of Table 8.¹³

	Baseline (constructed)	UK scheme (registered pairs, 2021–2025)
p_{A-B}	0.457	0.158
p_{B-A}	0.274	0.242
p_{A-B}^O	0.075	0.191
p_{B-A}^O	0.194	0.409
Reciprocal share of events	73.1%	40.0%
p_{A-B}/p_{B-A}	1.67	0.65

Table 8: Two calibrations of the arrival composition.

The two compositions differ in two economically distinct ways. The ratio of the reciprocal types *inverts* (1.67 versus 0.65): in the UK scheme it is $B-A$ rather than $A-B$ that is the abundant reciprocal type, so the labels “abundant” and “scarce” attach to the opposite types. More consequentially, over-demanded arrivals rise from 27% to 60% of tracked events, so each reciprocal type now faces more upgrade opportunities than it receives arrivals.

Under this calibration, the benchmark curvature and full-model local sign predictions are preserved: the CP and SP benchmark curvatures keep their strict signs, the quadrant sign pattern of Section 4.10 is reproduced, and the single-crossing condition (SC) holds at 100% of interior grid states for every reported c , though with a much narrower margin, because marginal values are far larger when reciprocal pairs are scarce. The active-range finding is if anything sharper: the interior trigger region is empty from $c = 0.9$ —again well below the two-way ceiling (exact ranges, margins, and state counts in Section XII of the online appendix).

What changes are the magnitudes and the abundant–scarce asymmetry. With over-demanded arrivals outnumbering reciprocal ones, both reciprocal pools are consumed about as fast as they fill: 92.9% of the discounted occupation mass lies within $|x_i| \leq 3$ and 18.9% at the empty state

¹³The cell-by-cell mapping, the treatment of under- and self-demanded entrants, and a cross-vintage stability check are described in Section X of the online appendix.

itself. The nearest state at which any inter-hospital action is optimal is $\max_i |x_i| = 8$, and the entire set of such states carries discounted mass 1.2×10^{-6} . On path, therefore, the planner essentially never coordinates, and the welfare and transplant gains over autarky are zero to three decimals at every reported c . Correspondingly, the sharp asymmetry of Section 4.4 weakens: at $c = 0.1$ the inter-hospital over-demanded channel is optimal at 52.5% of the relevant cross-type states for one reciprocal type and 45.7% for the other, rather than 99.9% and 0%.

We read this as a comparative static rather than a contradiction, and it identifies the composition statistic that governs the value of coordination: the ratio of reciprocal arrivals to the over-demanded opportunities that consume them. When that ratio is high, one reciprocal type accumulates, imbalances persist, and cross-hospital exchange has something to work with—the configuration the CP benchmark describes and the one in which our baseline magnitudes arise. When it is low, pools stay near the empty state, imbalances never develop, and coordination is optimally left unused however cheap it is. Taken together with the asymmetry results of Section 4.7, this result shows that the reciprocal-to-opportunity arrival ratio and cross-hospital asymmetry jointly determine whether cooperation has substantial value; average matching rates alone are therefore insufficient for evaluating cooperation policies.

Two caveats keep this second calibration in its place as a robustness exercise rather than a replacement. The registration counts are a selected sample: which incompatible pairs enter a national scheme depends on referral practice and on the local availability of ABO-incompatible and HLA-incompatible transplantation, and 286 of the 1,897 pairs in the scheme’s matching runs are compatible pairs seeking a better match (the matching-run population counts participation in quarterly runs and includes these compatible pairs, so it differs from the 1,406 registered incompatible pairs behind Table 8). And our over-demanded arrivals are, by Assumption 6, matched on arrival, whereas the registered over-demanded pairs of a real scheme wait: reading registration counts as arrival rates therefore overstates the rate at which reciprocal inventory is consumed, and biases this calibration toward thin pools. The baseline construction remains our primary calibration for that reason.

4.9 Batched clearing

The registry calibration also motivates a robustness check on the *timing* convention: the UK scheme identifies exchanges in *quarterly matching runs* rather than at each arrival,¹⁴ whereas our planner decides at every arrival (Assumption 7). To gauge what this convention contributes to the just-in-time finding, we solve a batched variant of the CP benchmark in which arrivals accumulate for k events and all decisions—including the disposition of queued over-demanded opportunities—are taken only at the clearing run, with any number of matches allowed per run. Batching is a modest delay cost, not a change in policy structure: relative to event-driven clearing, welfare falls by 0.19% at $k = 5$, 0.43% at $k = 10$, and 1.16% at $k = 25$ (comparable at the calibrated waiting cost: 0.18%, 0.40%, 1.10%), where $k = 25$ corresponds to quarterly runs at a 100-arrivals-per-year scale. More importantly, the batched planner does *not* clear greedily at runs: in 1,094 of 1,521 interior evaluation states it matches fewer pairs than feasible, preserving the same option-value retention logic that generates just-in-time behavior in the event-driven model. The batched benchmark thus confirms that the retention mechanism survives periodic clearing, although the full-model just-in-time geometry of

¹⁴NHS Blood and Transplant, UK Living Kidney Sharing Scheme: matching runs are performed four times per year (January, April, July, October); see odt.nhs.uk.

Section 4.5 remains conditional on arrival-triggered decisions.

4.10 Numerical verification of structural properties

We next test numerically whether the supermodularity and submodularity properties of Lemmas 2 and 3 appear locally in the constrained full model. Figure 5 plots the cross-difference of the value function in pool coordinates (x_1, x_2) at $c = 0.1$ (left) and $c = 0$ (right). Red indicates non-negative cross-difference (supermodularity); blue indicates non-positive (submodularity).




Figure 5: Cross-difference of the constrained full model’s value function in pool coordinates. Red (+) = supermodularity (complementary quadrants II and IV); blue (–) = submodularity (substitutable quadrants I and III). Under the baseline parameterization, the computed cross-differences display the quadrant-specific signs suggested by the analytical benchmarks. The two panels share a common color scale; magnitudes are larger at $c = 0$ because costless coordination amplifies the interaction between hospitals.

Under the baseline parameterization, the computed cross-differences display the quadrant-specific signs suggested by the analytical benchmarks, switching at the quadrant boundaries. In Quadrants II and IV (opposite excesses), the cross-difference is non-negative with range $[+0.000, +0.039]$, consistent with local supermodularity in the complementary quadrants. In Quadrants I and III (aligned excesses), the cross-difference is non-positive, with ranges $[-0.102, 0.000]$ in Quadrant I and $[-0.040, -0.003]$ in Quadrant III, consistent with local submodularity in the substitutable quadrants. At $c = 0$ the magnitudes increase because costless coordination amplifies the interaction between hospitals; at $c = 2$ the cross-differences are numerically zero everywhere, since hospitals operate independently. Section VIII of the online appendix tabulates these ranges together with the benchmark verifications (Table OA.4); in the CP benchmark the cross-difference is strictly positive and in the SP benchmark weakly negative with the maximum equal to zero at machine precision, so no boundary artifact remains. As an additional numerical check, we repeated the benchmark curvature verification with an explicit holding cost: the sign patterns of Proposition 2 are preserved to machine precision ($\theta \in \{0.01, 0.1\}$), at $\delta = 0.9$ with grid cap 120 and reporting window 60, and at $\delta = 0.99$ with cap 400 and window 80) on grids whose truncation boundary is kept far from the reporting window.¹⁵

We additionally check the single-crossing condition (SC) that underlies the monotonicity of the own-type threshold τ_{11} in Theorem 2. Across the interior of the grid, (SC) is satisfied at 100% of grid states

¹⁵The caveat matters: with $\theta > 0$ a reflecting grid boundary acts as free disposal, and the resulting curvature distortion propagates a distance of order $1/(1 - \delta)$ from the boundary. Naive truncated computations at δ close to one can therefore appear to violate the proposition; the violations vanish as the boundary is moved away.

for every reported cost value $c \in \{0.1, 0.5, 1.0, 2.0\}$, with a wide margin. With the incremental own-type reward $a_{11} = 1$ (Table 3), at $c = 0.1$ the maximum computed marginal value is $\Delta_1 v^* = 0.0006$, implying $\min A_{11} = a_{11} - \max \Delta_1 v^* = 0.9994$; across all reported cost values, $\max \Delta_1 v^* \leq 0.0044$, so $\min A_{11} \geq 0.9956 > 0$. This is the situation of Proposition 1, so the sufficient condition holds throughout the computed interior and the monotone characterization of Theorem 2 applies at this calibration.

The three findings now have a unified interpretation. In complementary states, the value function is supermodular and the planner uses inter-hospital exchange aggressively at low cost, with cooperation extinguished once c exceeds the incremental three-way surplus ($c \approx 1$), well below the two-way ceiling (Section 4.3). In substitutable states, the value function is submodular, intra-hospital matching dominates, and the inter-hospital channel is silent outside the deep-imbalance rebalancing region of Section 4.5. Under the baseline parameterization, in both regimes, at every evaluated grid state and along the optimal path (Section 4.6), the planner does not route an over-demanded arrival across hospitals to consume a scarce reciprocal pair: the pair’s option value exceeds the immediate incremental surplus (the reciprocal two-way channel, which by construction pairs the two types, is the scarce type’s only cross-hospital outlet). Taken together, the results show that reductions in real coordination costs expand the first-best cooperation region primarily in complementary states, whereas the option value of scarce reciprocal inventory limits cooperation elsewhere.

4.11 Robustness to a calibrated waiting cost

Proposition 2 (Section 3.4) establishes that an explicit per-pair waiting cost leaves the structural results intact. This subsection asks the quantitative question: at a calibrated waiting cost, embedded in an internally consistent temporal calibration, do the numbers of this section change? We proceed in two steps: first a comparative static in the per-event holding cost θ alone, holding the discount factor at its baseline value, and then joint variation of (δ, θ) , since the discount rate implied by our computational δ at any given arrival rate is itself a calibration object (see the temporal-calibration paragraph at the start of this section).

The calibration follows the transplant-equivalent units of Section 3. A pair waiting in the pool remains on dialysis, at a Medicare cost of roughly \$90,000 per patient-year [United States Renal Data System, 2018]; adding the forgone quality of life during a year of waiting at the valuations underlying Held et al. [2016] raises the flow burden to at most about \$250,000 per pair-year. Dividing by the \$1.1 million social value of one transplant expresses the waiting cost in transplant-equivalents, $\kappa \in [0.08, 0.23]$ per pair-year, and dividing by the arrival rate converts it to the per-event holding cost the solver uses, $\theta = \kappa/\lambda$: with the baseline illustration $\lambda = 1,000$ arrivals per year, $\theta \in [0.8, 2.3] \times 10^{-4}$ per pair per event. Table 9 first varies θ alone at the baseline $\delta = 0.999$ —a comparative static that isolates the holding-cost channel. Its final column, $\theta = 10^{-3}$, is a diagnostic upper value rather than a description of a thinner market: a market with $\lambda = 100$ arrivals per year changes the discounting between events as well as the holding cost per event, and the internally consistent thin-market computation appears below. All magnitudes in this subsection are computed under Assumption 8; the welfare cost of that restriction has been quantified only at $\theta = 0$ (Appendix B), and the corresponding approximation error under the calibrated waiting cost has not been quantified, so the levels below are conditional on the immediate-clearing class.

Two conclusions follow. First, within the immediate-clearing policy class, all qualitative comparative-

	$\theta = 0$ (baseline)	$\theta = 10^{-4}$ (calibrated range)	$\theta = 2.5 \times 10^{-4}$	$\theta = 10^{-3}$ (diagnostic)
Trigger region, $c = 0.1$ ($ x_i \leq 45$)	3,980	4,102	4,214	4,658
OD channel, abundant type ($c = 0.1$)	99.9%	99.9%	99.9%	99.9%
OD channel, scarce type ($c = 0.1$)	0.0%	0.0%	0.0%	0.0%
On-path inter-hospital share ($c = 0.1$)	0.30%	0.31%	2.7%	3.9%
Welfare gain over autarky ($c = 0.1$)	+0.15%	+0.17%	+0.35%	+1.8%
at full asymmetry $a = 1$	+68%	+73%	+80%	+136%
Retention advantage, Appendix B counterexample	+0.091	+0.092	+0.093	+0.100

Notes: Grid caps ± 150 for $\theta > 0$ (at $\theta = 0$ the results are cap-invariant; Section 4.3). With holding costs, deep-inventory values matter, so levels retain some grid sensitivity: between caps ± 50 and ± 150 the trigger count moves by up to 3.4% and, at the two highest θ , the on-path share and welfare gain by a few tenths of a percentage point, with every direction of comparison unchanged. All columns hold $\delta = 0.999$ fixed and are comparative statics in the holding cost alone; internally consistent (δ, θ) pairs are reported in the text. The retention-advantage row is computed in the counterexample environment of Appendix B, not under the baseline arrival composition. At high θ , percentage gains are amplified by the shrinking autarky denominator (holding costs accumulate in the comparator), so the level results in the text are the primary outcome. The $a = 1$ rows are stress-test endpoints, not headline results, and conservative for $\theta > 0$: truncation caps the autarky comparator’s unbounded accumulation of holding costs. The single-crossing margin only widens with θ (the marginal value of a stored pair falls), so Theorem 2’s monotone characterization continues to apply throughout.

Table 9: The holding-cost comparative static at the baseline discount factor.

static results survive the calibrated waiting cost (the trigger-frontier overlay in Section IX of the online appendix, Figure OA.4, shows the policy geometry): the policy geometry moves by a few percent, the abundant–scarce asymmetry of Section 4.4 is exactly preserved at every θ , the single-crossing condition holds with a wider margin, and the retention advantage in the Appendix B counterexample moves from 0.091 to 0.093—the adverse direction identified in Section 3.4, at a negligible magnitude. Second, the direction of the κ -comparative static is itself economically informative: in every arrival environment we compute, waiting costs make the planner *more* willing to cooperate, not less, because an inter-hospital match sheds pairs from costly storage—the disposal-value mechanism of Section 3.4 operating in cooperation’s favor. (Proposition 2 preserves curvature and thresholds for any $\kappa \geq 0$; it does not by itself fix the direction in which the cooperation region moves, so the direction is a numerical finding for the reported environments.) On path, the inter-hospital share rises with θ , the symmetric-hospital welfare gain roughly doubles over the calibrated range, and the asymmetric-pair gains of Section 4.7 strengthen. A health-economic reading is that the case for investing in cross-hospital coordination is strongest precisely where waiting is most costly—thin markets with long inter-arrival times—which the internally consistent thin-market computation below makes concrete.

Joint temporal calibration and sensitivity in δ . The columns of Table 9 hold δ fixed; an internally consistent calibration moves (λ, ρ, κ) together. Decomposing ρ as in the temporal-calibration paragraph at the start of this section—pure time preference of 3–5% plus attrition of several percent per year—gives $\rho \approx 0.12$ per year, hence at the dense-market illustration $\lambda = 1,000$: $\delta = \lambda/(\lambda + \rho) \approx 0.9999$ with calibrated $\theta \approx 1.5 \times 10^{-4}$; and at a thin market $\lambda = 100$: $\delta \approx 0.9988$ with $\theta \approx 1.5 \times 10^{-3}$. Table 10 traces the δ -dimension at $\theta = 0$; two computations then complete the joint picture. Three tiers should be kept distinct throughout: the baseline $\delta = 0.999$ is a computational benchmark; the coherent dense-market parameters support *structural verification only*, since welfare levels there are not identified at feasible grid depths; and the coherent thin-market parameters are an

approximate quantitative exercise whose cap sensitivity we report. The paper’s welfare magnitudes are accordingly benchmark computations plus sensitivity analysis, not levels evaluated at a preferred joint calibration.

δ	Implied ρ at $\lambda=1,000$	Welfare gain	Inter share	OD abundant/scarce	Trigger region, $c = 1.1$
0.99	$\approx 10/\text{yr}$	+0.95%	2.5%	90.6% / 45.7%	3,362 states
0.999	$\approx 1/\text{yr}$ (baseline)	+0.15%	0.30%	99.9% / 0.0%	empty
0.9995	$\approx 0.5/\text{yr}$	+0.08%	0.15%	100% / 0.0%	empty
0.9999	$\approx 0.1/\text{yr}$ (defensible)	+0.02%	0.03%	100% / 0.0%	empty

Notes: Symmetric hospitals, $c = 0.1$, $\theta = 0$; welfare gain and inter share from empty pools against matched autarky; OD columns as in Table 9; the last column counts inter-hospital trigger states at $c = 1.1$ in the window $|x_i| \leq 45$. Welfare-gain and share entries agree to reported precision between grid caps ± 50 and ± 100 at every δ .

Table 10: Sensitivity of the symmetric-hospital results to the discount factor.

Three patterns are worth recording. First, every structural finding is stable in the empirically relevant direction: for all $\delta \geq 0.999$ the abundant–scarce asymmetry is exact and the trigger region at $c = 1.1$ is empty—the active-range boundary of Section 4.3 does not move outward as discounting weakens. At $\delta = 0.99$ —included only to bracket, since it corresponds to $\rho \approx 10$ per year at the dense-market scale—the option-value effects weaken exactly as the theory predicts: the scarce OD channel activates and cooperation survives past $c = 1.1$. The findings’ comparative static in δ thus runs through the option-value mechanism the paper describes. Second, the symmetric-hospital welfare gain *falls* as δ rises toward the defensible range—from +0.95% at $\delta = 0.99$ to +0.02% at $\delta = 0.9999$ —so the baseline does not understate symmetric gains by discounting too heavily: under every temporal calibration we can compute, they remain below one percent, reinforcing the composition message of Section 4.7. Third, the internally consistent thin-market pair $(\delta, \theta) = (0.9988, 1.5 \times 10^{-3})$ yields an inter-hospital share of roughly 5% and a welfare gain of +2.3 to +2.7% over autarky (across grid caps 50–100), with the abundant–scarce asymmetry preserved. The composition of that gain makes the health-economic mechanism transparent. In transplant-equivalent levels per 1,000 arrivals (cap 100, with cap 50 in parentheses), the net gain of +27.7 (+25.1) decomposes into +2.3 (+2.2) of additional transplant surplus plus +30.1 (+28.9) of avoided waiting costs, less 4.7 (5.9) of coordination costs incurred: the benefit of cooperation in a thin, waiting-costly market is primarily avoided waiting rather than additional transplantation—the disposal mechanism stated in levels rather than percentages, and a decomposition that does not depend on the autarky comparator’s holding-cost base. Once discounting and holding costs move together, the thin-market case for coordination is therefore stronger, not weaker, than the fixed- δ diagnostic suggests. At the dense-market coherent point $(0.99988, 1.5 \times 10^{-4})$, by contrast, level computations are not identified at feasible grid caps: the discounted horizon of order $1/(1 - \delta)$ exceeds any computable pool depth, so we report structure only (the asymmetry and active-range entries match the $\delta = 0.9999$ row) and let the δ -series bracket the levels.¹⁶

¹⁶Replication: the temporal-calibration routines in the replication code. At $\theta > 0$ the thin-market entries move by a few tenths of a percentage point across caps 50–100, with every direction unchanged; exact levels at five-nines discounting are future work (Section 5).

5 Conclusion

We developed a dynamic two-hospital kidney exchange model in which a benevolent planner, facing stochastic arrivals, chooses at each event between matching within a hospital, matching across hospitals at a real coordination cost c , and waiting. Two analytical benchmarks isolate the forces at work: with complementary pools the value function is supermodular and a larger partner pool accelerates cooperation, while with substitutable pools it is submodular and optimal delay lengthens under a single-crossing condition; in both, the optimal policy is a threshold rule that trades the immediate surplus of an exchange against the option value of a thicker future pool. The constrained full-action model, solved numerically in its exact net-position representation, exhibits the local curvature signs and threshold-like policy geometry predicted by the analytical benchmarks and yields three substantive findings: coordination costs extinguish cooperation near the incremental three-way surplus for compositionally similar hospitals; the welfare at stake is governed jointly by the composition of arrivals and by cross-hospital asymmetry—a fraction of a percent under symmetry even at zero cost, rising by orders of magnitude toward the full-specialization stress-test endpoint and its size-neutral counterpart; and under arrival-triggered decisions, inter-hospital exchange is generally used just in time rather than as a broad inventory-smoothing device, apart from a narrow same-type rebalancing exception.

The first-best allocation these results characterize is the natural benchmark for policy, and its comparative statics separate two instruments often discussed together. Policies that reduce the real resource cost of coordination—logistics, shared transport infrastructure, organ-shipment technology—move the planner’s cooperation frontier outward. Reimbursement rules, by contrast, shift the incidence of c across hospitals without changing the social cost, and so can at most implement, never improve, the first best. Which decentralized arrangements—reimbursement, credit mechanisms, or priority schemes—can implement this allocation under private information and strategic hospital behavior is a question our model does not answer: we characterize the allocation such mechanisms would need to support and leave the incentive analysis to future work.

The robustness programme separates what is structural from what is quantitative. The curvature and threshold structure and the active-range logic survive a registry-based second calibration, an explicitly calibrated waiting cost (Proposition 2), and discount factors varied toward the internally consistent temporal calibration, and the retention mechanism survives periodic clearing in the batched benchmark of Section 4.9; the abundant–scarce option-value asymmetry remains qualitatively present under the registry composition but is sharply attenuated there.

An empirical extension would unbundle the reduced-form ρ into pure time preference and attrition and estimate the waiting cost κ from registry data, so that the model could be disciplined against observed welfare-loss estimates such as those in Agarwal et al. [2019]. Section 4.11 takes a first step by calibrating κ against dialysis costs and transplant values, but exact level computations at the implied five-nines discounting remain beyond truncated-grid methods and call for different computational tools. Future research should also extend the framework to larger hospital networks, richer compatibility structures, longer cycles and chains, and batched clearing in the full model, and embed the coordination cost in a dynamic mechanism-design analysis of reimbursement and credit schemes. Empirical validation with registry data would test whether observed programme behavior tracks the threshold logic characterized here.

Appendix A Second-Order Properties

We write e_i for the i -th unit vector in $\mathbb{R}^2$. For a function $f : \mathbb{N}_0^2 \rightarrow \mathbb{R}$ and index $i \in \{1, 2\}$, we define the first-order difference operator $\Delta_i f(s) := f(s) - f(s - e_i)$ whenever $s_i \geq 1$. We also define the joint difference operator $\Delta_{1+2} f(s) := f(s) - f(s - e_1 - e_2)$ and the two shift difference operators $\Delta_{1,-2} f(s) := f(s) - f(s - e_1 + e_2)$ and $\Delta_{2,-1} f(s) := f(s) - f(s - e_2 + e_1)$

Definition 1. A function $f : \mathbb{N}_0^2 \rightarrow \mathbb{R}$ is **concave in component i** if $\Delta_i f(s + 2e_i) \leq \Delta_i f(s + e_i)$ for all s , or equivalently $\Delta_i \Delta_i f(s) \leq 0$ for all s . It is **componentwise concave** if this holds for every $i \in \{1, 2\}$.

Definition 2. A function $f : \mathbb{N}_0^2 \rightarrow \mathbb{R}$ is **supermodular** if $\Delta_1 f(s + e_2) \geq \Delta_1 f(s)$ for all s , or equivalently $\Delta_1 \Delta_2 f(s) \geq 0$ for all s .

Definition 3. A function $f : \mathbb{N}_0^2 \rightarrow \mathbb{R}$ is **submodular** if $\Delta_1 f(s + e_2) \leq \Delta_1 f(s)$ for all s , or equivalently $\Delta_1 \Delta_2 f(s) \leq 0$ for all s .

Definition 4. A function $f : \mathbb{N}_0^2 \rightarrow \mathbb{R}$ is **diagonally concave** if

$$\Delta_1 \Delta_{1+2} f(s) \leq 0 \quad \text{and} \quad \Delta_2 \Delta_{1+2} f(s) \leq 0$$

wherever the displayed differences are defined; equivalently, the joint difference $\Delta_{1+2} f$ is non-increasing in each pool size, so the marginal $\Delta_i f$ is non-increasing along the main diagonal $s - e_1 - e_2 \rightarrow s$.

Definition 5. A function $f : \mathbb{N}_0^2 \rightarrow \mathbb{R}$ is **anti-diagonally concave** if

$$\Delta_1 \Delta_{1,-2} f(s) \leq 0 \quad \text{and} \quad \Delta_2 \Delta_{2,-1} f(s) \leq 0$$

wherever the displayed differences are defined; equivalently, for each $i \neq j$ the marginal $\Delta_i f$ is non-increasing along the anti-diagonal $s - e_i + e_j \rightarrow s$.

The two conditions run in different directions and are not interchangeable. The CP value function is diagonally concave (Lemma 2); the SP value function is anti-diagonally concave (Lemma 3). Neither model requires both.

Case-analysis method for max operators. Throughout the proofs, we frequently encounter inequalities of the form $\max\{a_1, \dots, a_k\} + \max\{b_1, \dots, b_k\} \leq \max\{c_1, \dots, c_k\} + \max\{d_1, \dots, d_k\}$. Since the left-hand side equals $a_{i^*} + b_{j^*}$ for some specific maximizers i^*, j^* , and the right-hand side satisfies $\text{RHS} \geq c_m + d_n$ for any choice of indices m, n , the inequality holds provided that for every pair (i, j) , there exists a pair (m, n) such that $a_i + b_j \leq c_m + d_n$. The pair (m, n) need not equal (i, j) ; we are free to choose different argument matchings on the right-hand side. For $\geq$ inequalities (supermodularity), the logic is reversed: for each right-hand side selection, we exhibit a left-hand side selection that dominates.

Appendix B Proof of Lemma 1, and the restrictiveness of Assumption 8

Proof of Lemma 1. Coexistence at h_i could be created only by retaining an arriving reciprocal pair while its reciprocal type is in stock, which Assumption 8 excludes; every other feasible transition (an exchange, which consumes the arriving pair or a stored one, or retention with $s_i^{-j} = 0$) preserves $\min\{s_i^{AB}, s_i^{BA}\} = 0$. On the no-coexistence set the map $s \mapsto (x_1, x_2)$ is a bijection, and every branch of

(1) under Assumption 8 reads off the net positions alone: an arriving type- j pair at h_i moves x_i one step in the j direction with surplus 2 if the opposite type is in stock there and 0 otherwise, or moves x_{-i} one step in the j direction with surplus $2 - c$ (feasibility of each branch is the sign condition on the corresponding x); an over-demanded arrival leaves the state unchanged (baseline surplus 2) or moves x_i (own three-way, surplus 3) or x_{-i} (cross three-way, surplus $3 - c$) one step against j , feasibility again being the sign condition on the corresponding x . Formally, write $\Phi(s) = (x_1, x_2)$ for this bijection, and let T denote the constrained Bellman operator on the no-coexistence set and $\tilde{T}$ the net-position operator used in Section 4. The branch inventory above shows that $(Tf) \circ \Phi^{-1} = \tilde{T}(f \circ \Phi^{-1})$ for every bounded f : the two operators are conjugate under Φ . Both are δ -contractions (the argument of Lemma 4 applies verbatim to both), so by uniqueness of fixed points $v^* = \tilde{v}^* \circ \Phi$, where $\tilde{v}^*$ is the net-position value function. Hence the constrained-optimal value depends on the state only through (x_1, x_2) .

We next show that Assumption 8 is a genuine restriction: without it, retaining an arriving reciprocal pair alongside its stored reciprocal can be strictly optimal, so the pair bound $v^*(s + e_i^j) - v^*(s - e_i^{-j}) \leq 2$ fails in general.

Proposition 3. *There exist arrival probabilities—including profiles with every arrival type occurring with strictly positive probability—and discount factors $\delta \in (0, 1)$ under which, at a state with $s_i^{-j} \geq 1$, retaining an arriving reciprocal pair of type j is strictly better than matching it on arrival, and $v^*(s + e_i^j) - v^*(s - e_i^{-j}) > 2$.*

Proof. Set $c = 2$. Every inter-hospital branch of (1) is then weakly dominated: the inter-hospital two-way exchange yields $0 + v^*(s - e_{-i}^{-j}) \leq v^*(s + e_i^j)$ and the inter-hospital three-way yields $1 + v^*(s - e_{-i}^j) \leq 2 + v^*(s)$, because v^* is nondecreasing in each component (a standard coupling argument). The two hospitals therefore decouple, and each solves a single-hospital problem on its own event clock, with conditional arrival probabilities p_T and per-own-event discount $\tilde{\delta} = (\delta/2)/(1 - \delta/2)$, which ranges over $(0, 1)$ as δ does. We work with incremental rewards, subtracting the state-independent value of the over-demanded baseline stream: matching an over-demanded arrival with a stored reciprocal pair yields +1, the intra-hospital reciprocal exchange yields +2, and no comparison between actions is affected. Write $(a, b) = (s^{AB}, s^{BA})$ for the focal hospital's stocks and choose

$$p_{A-B} = p_{A-B}^O = \frac{1}{2}, \quad p_{B-A} = p_{B-A}^O = 0, \quad \tilde{\delta} = 0.99.$$

Step 1: value with no stored B–A pair. From a state $(n, 0)$ no B–A pair ever appears, so the reachable states are $(n, 0)$, $n \geq 0$. Let V_n denote the optimal incremental value there. An arriving A–B pair can only be retained; an over-demanded arrival is matched with a stored pair whenever one is available (verified below). The Bellman equations are

$$V_0 = \frac{\tilde{\delta}}{2}(V_1 + V_0), \quad V_n = \frac{\tilde{\delta}}{2}(V_{n+1} + 1 + V_{n-1}), \quad n \geq 1.$$

Incremental values are bounded, so the explosive root of the associated recurrence is excluded and $V_n = K + Cr^n$ with

$$K = \frac{\tilde{\delta}}{2(1 - \tilde{\delta})}, \quad r = \frac{1 - \sqrt{1 - \tilde{\delta}^2}}{\tilde{\delta}} \in (0, 1),$$

and the boundary condition at $n = 0$, combined with the characteristic identity $\frac{\tilde{\delta}}{2}(1 + r^2) = r$, gives $C = -r/(1 - r)$. Hence $V_n - V_{n-1} = r^n$, so at an over-demanded arrival using a stored pair beats passing by $1 - r^n > 0$, confirming the presumed policy. At $\tilde{\delta} = 0.99$: $r = 0.867609$, and $V_0 = 42.946632$, $V_1 = 43.814241$, $V_2 = 44.566986$.

Step 2: a feasible retention policy with one stored B–A pair. From states $(n, 1)$ consider the stationary policy: retain arriving A–B pairs while $n \leq 1$; at $n = 2$, match the next arriving A–B pair with the stored B–A pair (surplus 2, moving to $(2, 0)$); and match every over-demanded arrival with a stored A–B pair when one is available. Its value W_n at $(n, 1)$ solves the linear system

$$W_0 = \frac{\tilde{\delta}}{2}(W_1 + W_0), \quad W_1 = \frac{\tilde{\delta}}{2}(W_2 + 1 + W_0), \quad W_2 = \frac{\tilde{\delta}}{2}(3 + V_2 + W_1),$$

whose solution at $\tilde{\delta} = 0.99$ is $W_0 = 44.145741$, $W_1 = 45.037574$, $W_2 = 45.839257$.

Step 3: comparison. Consider the state $(0, 1)$ with an A–B pair arriving. Immediate clearing yields $2 + v^*(0, 0) = 2 + V_0 = 44.946632$. Retention leads to $(1, 1)$, whose optimal value satisfies $v^*(1, 1) \geq W_1$ because the retention policy above is feasible there. Hence

$$v^*(1, 1) - (2 + v^*(0, 0)) \geq W_1 - 2 - V_0 = 0.090942 > 0 :$$

retention is strictly optimal, and the pair bound fails, $v^*(1, 1) - v^*(0, 0) \geq 2.090942 > 2$.

Step 4: strictly positive probabilities. The Bellman operator is a $\tilde{\delta}$ -contraction whose fixed point is continuous, in sup norm, in the arrival probabilities. The inequality in Step 3 is strict, so it persists under $p_{A-B} = p_{A-B}^O = (1 - 2\varepsilon)/2$, $p_{B-A} = p_{B-A}^O = \varepsilon$ for all sufficiently small $\varepsilon > 0$, and the state $(0, 1)$ is then reached from the empty pool by a single B–A arrival. Value iteration confirms the construction: at $\varepsilon = 0.01$ the retention advantage at $(0, 1)$ is 0.0747, and at $\varepsilon = 0$ it is 0.090942 with $v^*(1, 1) = W_1$, so the retention policy is not merely feasible but optimal there. $\square$

The mechanism is the inventory option of Remark 2: the stored B–A pair waits for a future A–B arrival while the A–B stock serves the over-demanded channel; the mutual surplus of 2 is delayed, not destroyed, and the delay cost vanishes as the discount factor approaches one.

Appendix C Proofs of the results in Sections 3.2 - 3.4

Although we present the main structure of the proofs in this section, some involve tedious chains of inequalities. To streamline the exposition, we defer the supporting lemmas (Lemmas OA.1–OA.9) to Section XIV of the online appendix.

Lemma 4. (Existence, uniqueness, and convergence) *Let T be the Bellman operator of the full model (1), of the CP benchmark (2), or of the SP benchmark (3), acting on the space of bounded real-valued functions on its state space equipped with the sup norm. Then T is a δ -contraction; a unique bounded fixed point v^* exists; and for every bounded v_0 the value-iteration sequence $v_k = T v_{k-1}$ converges uniformly to v^* at geometric rate δ .*

Proof. One-step rewards are bounded ($r_n \in [0, 3]$ in the full model with $c \in [0, 2]$; by the largest of the finitely many constant branch rewards in the benchmarks). Each operator is monotone, and adding a constant k to its argument raises the image by exactly δk , so Blackwell’s sufficient conditions apply

[Blackwell, 1965]; the conclusions follow from Banach's fixed-point theorem [Puterman, 1994, Stokey et al., 1989]. $\square$

Proof of Lemma 2. We give an induction argument to establish the structural properties of the value function in the CP model. We proceed by induction on the value iteration. All second-order differences below are asserted at every state where every term is defined (Remark 1). The supporting lemmas display their case analyses at interior states, where every branch of every max is feasible; the boundary paragraph closing this proof extends the induction to the remaining states at which a target difference is defined.

Initial step: $v_0(s_1, s_2) = 0$ for all $s_1, s_2 \geq 0$. All second-order differences of v_0 are zero, since v_0 is constant on this domain: $\Delta_1 \Delta_1 v_0 = 0$, $\Delta_2 \Delta_2 v_0 = 0$, $\Delta_1 \Delta_2 v_0 = 0$, $\Delta_1 \Delta_{1+2} v_0 = 0$, $\Delta_2 \Delta_{1+2} v_0 = 0$.

Inductive Hypothesis ($n - 1$): For all $s = (s_1, s_2)$, the value function v_{n-1} satisfies

$$\Delta_1 \Delta_2 v_{n-1} \geq 0, \quad (\text{H1})$$

$$\Delta_1 \Delta_1 v_{n-1} \leq 0, \quad (\text{H2})$$

$$\Delta_1 \Delta_{1+2} v_{n-1}, \Delta_2 \Delta_{1+2} v_{n-1} \leq 0, \quad (\text{H3})$$

$$\Delta_2 \Delta_2 v_{n-1} \leq 0. \quad (\text{H4})$$

We show that if v_{n-1} satisfies (H1)–(H4), then v_n also satisfies these properties. Consequently, by Lemma 4 (existence and uniqueness), the optimal value function $v^* = \lim_{n \rightarrow \infty} v_n$ inherits these properties. In particular, v^* is componentwise concave, supermodular, and diagonally concave. The Bellman update for the CP model is:

$$\begin{aligned} v_n(s_1, s_2) = & \delta \left[p_X \max \{ v_{n-1}(s_1+1, s_2), v_{n-1}(s_1, s_2-1) + a_{XY} \} \right. \\ & + p_Y \max \{ v_{n-1}(s_1, s_2+1), v_{n-1}(s_1-1, s_2) + a_{XY} \} \\ & + p_X^O \max \{ v_{n-1}(s_1, s_2), v_{n-1}(s_1-1, s_2) + a_X \} \\ & \left. + p_Y^O \max \{ v_{n-1}(s_1, s_2), v_{n-1}(s_1, s_2-1) + a_Y \} \right], \end{aligned} \quad (4)$$

which describes the one-step dynamic programming update. Here the first line corresponds to a type X arrival, the second to a type Y arrival, the third to an opportunity O_X arrival, and the fourth to an opportunity O_Y arrival.

Throughout, we make frequent use of the Bellman equation (4), which describes the one-step dynamic programming update. Here we focus on the structural implications of this Bellman update.

(1) *Componentwise concavity:* We first show concavity in s_2 . The proof for s_1 is analogous. Consider an arbitrary state (s_1, s_2) and its neighboring states in the s_2 -direction. Using the notation of forward differences, concavity in s_2 requires $\Delta_2 \Delta_2 v_n(s_1, s_2) = v_n(s_1, s_2) - 2v_n(s_1, s_2-1) + v_n(s_1, s_2-2) \leq 0$. By substituting the Bellman equation and canceling the common discount factor δ , after rearrangement, the inequality $\Delta_2 \Delta_2 v_n(s_1, s_2) \leq 0$ can be written as: $p_X \Phi_X + p_Y \Phi_Y + p_X^O \Phi_{O_X} + p_Y^O \Phi_{O_Y} \leq 0$, where each Φ term denotes the contribution of the corresponding arrival type to the second-order difference. In particular, we define for each arrival type an expression: Φ_X denotes the second-order difference of the X -arrival max expression in the s_2 direction, and similarly for the other terms. Because p_X, p_Y, p_X^O, p_Y^O are all positive, it suffices to prove that each Φ term is individually non-positive.

The model is symmetric under the joint swap $(s_1, X, O_X, a_X) \leftrightarrow (s_2, Y, O_Y, a_Y)$, and the hypothesis set (H1)–(H4) is closed under this swap. Note, however, that the swap maps a contribution to $\Delta_2\Delta_2$ into a contribution to $\Delta_1\Delta_1$; it does not deliver the Y -side contributions to $\Delta_2\Delta_2$, which we therefore verify directly. For a Y arrival ($\max\{v_{n-1}(s_1, s_2+1), v_{n-1}(s_1-1, s_2) + a_{XY}\}$), the four cases of the analysis parallel to Lemma OA.1 (Section XIV of the online appendix) reduce, after cancellation of a_{XY} , to: (H4) at (s_1, s_2+1) ; (H4) at (s_1-1, s_2) ; the chain $\Delta_2 v_{n-1}(s_1, s_2+1) \leq \Delta_2 v_{n-1}(s_1, s_2) \leq \Delta_2 v_{n-1}(s_1-1, s_2-1)$ ((H4) then (H3)); and $\Delta_2 v_{n-1}(s_1-1, s_2) \leq \Delta_2 v_{n-1}(s_1, s_2)$ ((H1)). For an O_Y arrival ($\max\{v_{n-1}(s), v_{n-1}(s-e_2) + a_Y\}$), the four cases reduce to (H4) at s , (H4) at $s-e_2$, two applications of (H4), and an exact equality of matched branches. Hence $\Phi_Y \leq 0$ and $\Phi_{O_Y} \leq 0$, and it remains to prove $\Phi_X \leq 0$ and $\Phi_{O_X} \leq 0$.

We next isolate the contribution of a type X arrival to the desired second-order inequality. To verify that $\Phi_X \leq 0$ (and similarly for the other terms), it suffices to check the inequality for each possible combination of choices made by the max operators. Since each max selects one of its two arguments, there are four ordered pairs of choices on each side. By the standard property that for any real numbers A, B, C, D , $\max\{A, B\} + \max\{C, D\} \geq \max\{A+C, A+D, B+C, B+D\}$, the original inequality holds if it holds for every combination where the same arguments are chosen on both sides. Thus, it suffices to check these four cases. The detailed verification is carried out in Lemma OA.1, where each of the four argument configurations is shown to yield a non-positive contribution.

The contribution of opportunity O_X arrivals, Φ_{O_X} , can be analyzed similarly. Again there are four sub-cases for which branch each maximum takes, and each configuration is shown to yield a non-positive contribution. The detailed verification is carried out in Lemma OA.2 (Section XIV of the online appendix).

We have shown that all four arrival contributions to $\Delta_2\Delta_2$ are non-positive (Φ_Y and Φ_{O_Y} directly above, Φ_X and Φ_{O_X} via the supporting lemmas), so $\Delta_2\Delta_2 v_n(s_1, s_2) \leq 0$: v_n (and hence v^*) is concave in the s_2 dimension. Concavity in the s_1 dimension follows by the joint swap: it maps the four verified contributions to $\Delta_2\Delta_2$ into the four contributions to $\Delta_1\Delta_1$ (the Y - and O_Y -cases above into the X - and O_X -cases, and Lemmas OA.1 and OA.2 into the Y - and O_Y -cases), with the swapped hypothesis set again (H1)–(H4).

(2) *Supermodularity*: By rewriting $\Delta_1\Delta_2 v_n(s_1, s_2)$, we have $[v_n(s_1, s_2) - v_n(s_1, s_2-1)] - [v_n(s_1-1, s_2) - v_n(s_1-1, s_2-1)] \geq 0$. Using the Bellman equation and canceling δ , this condition can be split into contributions from each arrival type, each multiplied by its arrival probability. Once again, by symmetry it suffices to focus on X -arrivals and O_X -arrivals.

We next isolate the contribution of a type X arrival to the cross-difference inequality. Lemma OA.3 (Section XIV of the online appendix) rewrites the corresponding Bellman terms and verifies that the type- X -arrival contribution satisfies the supermodularity inequality. In particular, this step shows that the X (and symmetrically the Y) arrival contributions preserve $\Delta_1\Delta_2 v_n \geq 0$.

We now verify supermodularity for opportunity arrivals. Lemma OA.4 (Section XIV of the online appendix) derives the corresponding reduced inequality and shows that the O_X -arrival contribution satisfies the supermodularity condition. Together with Lemma OA.3 for type X arrivals (and its symmetric counterpart for type Y arrivals), this establishes that all arrival types preserve $\Delta_1\Delta_2 v_n \geq 0$. By symmetry, the same argument applies to O_Y arrivals. Hence $\Delta_1\Delta_2 v_n(s_1, s_2) \geq 0$ for all states, and in the limit the optimal value function v^* is supermodular.

(3) *Diagonal concavity*: We show the case $i = 1$ (the $i = 2$ case is symmetric).

$$\Delta_1 \Delta_{1+2} v_n(s_1, s_2) = v_n(s_1, s_2) + v_n(s_1-2, s_2-1) - v_n(s_1-1, s_2) - v_n(s_1-1, s_2-1) \leq 0.$$

Using the Bellman equation and eliminating δ , this inequality again decomposes into components for each arrival type. Because of symmetry, it suffices to consider a type X arrival and an opportunity O_X arrival.

We next verify that type X arrivals preserve diagonal concavity. Lemma OA.5 (Section XIV of the online appendix) analyzes the contribution of a type X pair arriving to the second-order difference $\Delta_1 \Delta_{1+2} v_n$. By rewriting the Bellman terms and checking all four argument configurations, the lemma shows, under the inductive hypothesis, that this contribution is non-positive. The Y -arrival contribution to $\Delta_1 \Delta_{1+2}$ is verified directly (the swap does not deliver it; see below): with branches $g(s) = v_{n-1}(s_1, s_2+1)$ and $h(s) = v_{n-1}(s_1-1, s_2) + a_{XY}$, the four cases reduce, after cancellation of a_{XY} , to (H3) at (s_1, s_2+1) ; (H3) at (s_1-1, s_2) ; two applications of (H3); and an exact equality of matched branches. Hence the Y contribution is non-positive as well.

Next, for an opportunity O_X arrival, the contribution is analyzed in Lemma OA.6 (Section XIV of the online appendix). There the Bellman terms are reduced to an inequality that holds under the inductive hypothesis, establishing that this contribution is non-positive as well. The O_Y contribution to $\Delta_1 \Delta_{1+2}$ is verified directly: with $g(s) = v_{n-1}(s)$ and $h(s) = v_{n-1}(s-e_2) + a_Y$, the four cases reduce to (H3) at s ; (H3) at $s-e_2$; the chain $\Delta_{1+2} v_{n-1}(s) \leq \Delta_{1+2} v_{n-1}(s-e_1) \leq \Delta_{1+2} v_{n-1}(s-e_1-e_2)$ (both halves of (H3)); and $\Delta_1 v_{n-1}(s-e_2) \leq \Delta_1 v_{n-1}(s-e_1-e_2)$ ((H2)).

Since every arrival type contributes a non-positive amount to $\Delta_1 \Delta_{1+2} v_n$, their weighted sum is non-positive, yielding $\Delta_1 \Delta_{1+2} v_n(s_1, s_2) \leq 0$. The condition $\Delta_2 \Delta_{1+2} v_n \leq 0$ then follows by the joint swap $(s_1, X, O_X, a_X) \leftrightarrow (s_2, Y, O_Y, a_Y)$, which maps the four verified contributions to $\Delta_1 \Delta_{1+2}$ into the four contributions to $\Delta_2 \Delta_{1+2}$. Thus diagonal concavity is preserved in the value iteration, and by induction it holds for the optimal value function v^* .

Boundary states. At states where a branch of some max is infeasible (Remark 1), the branch is simply absent from the operator. For the $\leq$ -inequalities (concavity and diagonal concavity) an absent branch removes cases only on the left-hand side of the case analysis, and for the $\geq$ -inequality (supermodularity) only on the right-hand side, so the surviving cases form a subset of those verified above. Whenever a surviving case's stated matching would itself select an infeasible branch, an alternative matching closes the case with one-step instances of (H1)–(H4): for example, the concavity analysis at $s_2 = 2$ needs only (H4) and (H1); in the supermodularity analysis at $s_2 = 1$ the right-hand configuration selecting the exchange branch in both maxima is itself infeasible, so its case is void; and in the analysis of Lemma OA.6 at $s_1 = 2$ the surviving mixed case closes by the exact-equality matching of its Case 4. The induction therefore holds at every state where the target difference is defined, which is the domain asserted in Lemma 2.

Having verified concavity, supermodularity, and diagonal concavity for v_n (assuming they held for v_{n-1}), the induction is complete. Therefore, the optimal value function v^* in the CP model is concave in each component, supermodular, and diagonally concave.

Proof of Theorem 1. Let v denote the optimal value function. By Lemma 2, v is componentwise concave, supermodular, and diagonally concave.

Exchange threshold. For a two-way exchange using one type X pair and one type Y pair, define the net gain $G(s_1, s_2) := a_{XY} - \Delta_{1+2}v(s_1+1, s_2) = a_{XY} - [\nu(s_1+1, s_2) - \nu(s_1, s_2-1)]$. Exchange is optimal if and only if $G(s_1, s_2) \geq 0$: both branches of the maximization in the Bellman update carry the common factor δ , so δ does not enter the comparison, and exchange is optimal if and only if $a_{XY} \geq \Delta_{1+2}v(s_1+1, s_2)$.

Monotonicity of G in s_1 . Diagonal concavity (Lemma 2) gives $\Delta_1\Delta_{1+2}v \leq 0$, equivalently $\Delta_1v(s) \leq \Delta_1v(s - e_1 - e_2)$. The same inequality, rearranged, says $\Delta_{1+2}v(\cdot, s_2)$ is non-increasing in s_1 . Hence G is non-decreasing in s_1 . To see this explicitly, note that

$$\begin{aligned} \Delta_{1+2}v(s_1+2, s_2) - \Delta_{1+2}v(s_1+1, s_2) &= [\nu(s_1+2, s_2) - \nu(s_1+1, s_2-1)] \\ &\quad - [\nu(s_1+1, s_2) - \nu(s_1, s_2-1)] \\ &= \Delta_1v(s_1+2, s_2) - \Delta_1v(s_1+1, s_2-1). \end{aligned}$$

This expression involves Δ_1v evaluated at two states that differ in both coordinates, so componentwise concavity alone (which controls Δ_1v along fixed s_2) is not sufficient. Instead, note that $\Delta_1\Delta_{1+2}v \leq 0$ is precisely the diagonal concavity condition established in Lemma 2, which directly yields the required monotonicity.

Monotonicity of G in s_2 . Diagonal concavity (Lemma 2) gives $\Delta_2\Delta_{1+2}v \leq 0$. Expanding,

$$\Delta_{1+2}v(s_1+1, s_2+1) - \Delta_{1+2}v(s_1+1, s_2) = \Delta_2v(s_1+1, s_2+1) - \Delta_2v(s_1, s_2),$$

and diagonal concavity signs this expression as non-positive. Hence $\Delta_{1+2}v(s_1+1, \cdot)$ is non-increasing in s_2 and G is non-decreasing in s_2 . To verify, observe that

$$\begin{aligned} \Delta_{1+2}v(s_1+1, s_2+1) - \Delta_{1+2}v(s_1+1, s_2) &= [\nu(s_1+1, s_2+1) - \nu(s_1, s_2)] \\ &\quad - [\nu(s_1+1, s_2) - \nu(s_1, s_2-1)] \\ &= \Delta_2v(s_1+1, s_2+1) - \Delta_2v(s_1, s_2). \end{aligned}$$

Again, the two evaluation points differ in both coordinates, so supermodularity (which controls Δ_2v along fixed s_1) is not sufficient. The condition $\Delta_2\Delta_{1+2}v \leq 0$ from diagonal concavity directly signs this expression as nonpositive.

Since G is nondecreasing in both s_1 and s_2 , for each fixed s_2 there exists a smallest stock level at which exchange becomes optimal. Define $\bar{t}_{XY}(s_2) = \inf \{s_1 \geq 0 : a_{XY} \geq \Delta_{1+2}v(s_1+1, s_2)\}$. By the monotonicity of G in s_1 , exchange is optimal if and only if $s_1 \geq \bar{t}_{XY}(s_2)$. Because G is also nondecreasing in s_2 , the threshold $\bar{t}_{XY}(s_2)$ is nonincreasing in s_2 : a larger pool of complementary Y pairs lowers the bar for exchange. Symmetrically, an exchange is also characterized by $s_2 \geq \bar{t}_{YX}(s_1)$.

Opportunity threshold. For opportunity O_X arrivals, the planner compares using the opportunity with the two-way baseline. Define $H(s_1, s_2) := a_X - \Delta_1v(s_1, s_2) = a_X - [\nu(s_1, s_2) - \nu(s_1-1, s_2)]$. Using the opportunity is optimal if and only if $H(s_1, s_2) \geq 0$; as before, the common factor δ multiplies both branches and does not enter the comparison.

Monotonicity of H in s_1 . Componentwise concavity (Lemma 2) gives $\Delta_1\Delta_1v \leq 0$, so $\Delta_1v(\cdot, s_2)$ is nonincreasing in s_1 . Hence H is nondecreasing in s_1 .

Monotonicity of H in s_2 . Supermodularity (Lemma 2) gives $\Delta_1\Delta_2v \geq 0$, so $\Delta_1v(s_1, \cdot)$ is nondecreasing

in s_2 . Hence H is nonincreasing in s_2 : when there are many complementary Y pairs, each X pair becomes more valuable (it could participate in a future exchange with Y), so more X pairs are needed before the planner is willing to consume one on an opportunity.

Define $\bar{t}_{O_X}(s_2) = \inf \{s_1 \geq 0 : a_X \geq \Delta_1 v(s_1, s_2)\}$. By the monotonicity of H in s_1 , using the opportunity is optimal if and only if $s_1 \geq \bar{t}_{O_X}(s_2)$. Because H is nonincreasing in s_2 , the threshold $\bar{t}_{O_X}(s_2)$ is nondecreasing in s_2 .

The exchange threshold $\bar{t}_{XY}$ is nonincreasing in the complementary pool size (a larger complementary pool makes exchange easier to trigger), while the opportunity threshold $\bar{t}_{O_X}$ is nondecreasing in the complementary pool size (a larger complementary pool raises the option value of holding pairs, making the planner more reluctant to match them with opportunities). Hence the complementary model policy is of (monotone) threshold type, as stated.

Proof of Lemma 3. We now provide the proof of the properties of the value function. The approach parallels that of the proof of Lemma 2, with appropriate sign changes for inter-type effects. We proceed by induction on the value iteration. All second-order differences below are asserted at every state where every term is defined (Remark 1). The supporting lemmas display their case analyses at interior states, where every branch of every max is feasible; the boundary paragraph closing this proof extends the induction to the remaining states at which a target difference is defined.

Initial step. $v_0(s_1, s_2) = 0$ for all $s_1, s_2 \geq 0$. Consequently, all second-order differences of v_0 are zero: $\Delta_1 \Delta_1 v_0 = 0$, $\Delta_2 \Delta_2 v_0 = 0$, $\Delta_1 \Delta_2 v_0 = 0$, $\Delta_1 \Delta_{1,-2} v_0 = 0$, $\Delta_2 \Delta_{2,-1} v_0 = 0$.

Inductive Hypothesis ($n - 1$). For all s , v_{n-1} satisfies

$$\Delta_1 \Delta_2 v_{n-1} \leq 0, \tag{S1}$$

$$\Delta_1 \Delta_1 v_{n-1} \leq 0, \tag{S2}$$

$$\Delta_1 \Delta_{1,-2} v_{n-1}, \Delta_2 \Delta_{2,-1} v_{n-1} \leq 0, \tag{S3}$$

$$\Delta_2 \Delta_2 v_{n-1} \leq 0. \tag{S4}$$

Assuming (S1)–(S4) hold for v_{n-1} , we show they hold for v_n ; the properties then pass to the limit v^* by the uniform convergence in Lemma 4, so the optimal value function v^* in the SP model satisfies: concavity in each dimension; submodularity; and anti-diagonal concavity conditions. The Bellman update for the SP model is:

$$\begin{aligned} v_n(s_1, s_2) = & \delta \left[p_1 v_{n-1}(s_1+1, s_2) + p_2 v_{n-1}(s_1, s_2+1) \right. \\ & + p_1^O \max \{ v_{n-1}(s_1, s_2), v_{n-1}(s_1-1, s_2) + a_{11}, v_{n-1}(s_1, s_2-1) + a_{12} \} \\ & \left. + p_2^O \max \{ v_{n-1}(s_1, s_2), v_{n-1}(s_1, s_2-1) + a_{22}, v_{n-1}(s_1-1, s_2) + a_{21} \} \right]. \end{aligned} \tag{5}$$

We label the three arguments of each max as A (not using), B (consuming own type), C (consuming other type). Throughout, we make frequent use of the Bellman equation (5). Here we focus on the structural implications of this Bellman update.

The Bellman equation (5) has four terms. The type arrivals (X_1 and X_2) are simple: they contribute $v_{n-1}(s_1+1, s_2)$ and $v_{n-1}(s_1, s_2+1)$ respectively, which are just shifts of v_{n-1} and automatically inherit all second-order properties. The non-trivial contributions come from the opportunity arrivals O_1

and O_2 , each of which involves a three-way max expression.

Using the Bellman equation, each second-order inequality decomposes by arrival type:

$$p_1 \varphi_{X_1} + p_2 \varphi_{X_2} + p_1^O \varphi_{O_1} + p_2^O \varphi_{O_2} \leq 0 \quad (\text{or } \geq 0 \text{ for supermodularity}),$$

where each φ term denotes the contribution of the corresponding arrival type. Because all arrival probabilities are positive, it suffices to show that each arrival type individually preserves each property. Since p_1 and p_1^O are symmetric with p_2 and p_2^O (but for the other type), it suffices to prove that the contributions from type X_1 arrivals, O_1 arrivals, and O_2 arrivals individually preserve the required properties.

(1) *Componentwise concavity.* Consider any state (s_1, s_2) and its neighbors in the s_2 direction. Concavity, after arrangement, requires $[v_n(s_1, s_2) - v_n(s_1, s_2-1)] - [v_n(s_1, s_2-1) - v_n(s_1, s_2-2)] \leq 0$, which we interpret as the sum of four non-positive terms, each corresponding to one of the four terms in the Bellman update above. The type-arrival contributions are shifts of v_{n-1} and inherit concavity from (S2) and (S4) directly. For the opportunity arrivals, the detailed verification is carried out in Lemma OA.7 (Section XIV of the online appendix), where the O_2 -arrival contribution to $\Delta_2 \Delta_2 v_n$ is shown to be non-positive by examining all nine cases. Every case of that proof treats the branch constants a_{22}, a_{21} as arbitrary reals—they cancel or transfer across the matching in each case—so the same lemma applies verbatim to the O_1 max, whose branches have the same displacements with constants (a_{12}, a_{11}) : the O_1 -arrival contribution to $\Delta_2 \Delta_2 v_n$ is non-positive as well. Concavity in s_1 , (S2), then follows by the joint swap $(s_1, X_1, O_1, a_{11}, a_{21}) \leftrightarrow (s_2, X_2, O_2, a_{22}, a_{12})$, under which the model maps to an SP model with permuted parameters and the hypothesis set (S1)–(S4) is closed. Thus (S2) and (S4) are maintained through the value iteration.

(2) *Submodularity.* It suffices to prove that for type X_1 arrivals, the contribution $\Delta_1 \Delta_2 v_{n-1}(s_1+1, s_2) \leq 0$ by the inductive assumption (S1). For type X_2 arrivals, $\Delta_1 \Delta_2 v_{n-1}(s_1, s_2+1) \leq 0$ by (S1).

We now verify the submodularity condition for opportunity arrivals. Specifically, consider an opportunity O_1 arriving. The detailed verification is provided in Lemma OA.8 (Section XIV of the online appendix), where the Bellman terms are rewritten into the reduced max-comparison form and shown to satisfy the required inequality. Hence, the O_1 -arrival contribution holds; by symmetry, the same argument applies to O_2 arrivals, and the remaining arrival types follow analogously.

(3) *Anti-Diagonal concavity.* The condition $\Delta_1 \Delta_{1,-2} v_n(s_1, s_2) \leq 0$ requires $v_n(s_1, s_2) + v_n(s_1-2, s_2+1) \leq v_n(s_1-1, s_2) + v_n(s_1-1, s_2+1)$. When a type X_1 pair arrives, it suffices to show $\Delta_1 \Delta_{1,-2} v_{n-1}(s_1+1, s_2) \leq 0$, which follows by the inductive assumption (S3). Similarly for type X_2 arrivals.

The verification for opportunity arrivals is provided in Lemma OA.9 (Section XIV of the online appendix), where the reduced inequality form is shown to satisfy the required condition. Hence, $\Delta_1 \Delta_{1,-2} v_n(s_1, s_2) \leq 0$ for all (s_1, s_2) , and by symmetry, $\Delta_2 \Delta_{2,-1} v_n(s_1, s_2) \leq 0$ holds as well. This completes the induction.

Boundary states. At states where a branch of an opportunity max is infeasible (Remark 1), that branch is absent from the operator. For the $\leq$ -inequalities proved here an absent branch removes cases only on the left-hand side of each case analysis, so the surviving cases form a subset of the nine verified in the supporting lemmas; in every surviving case the matching used places branches only at states

where they remain feasible, and where a stated matching would fail, a one-step instance of (S1)–(S4) closes the case. The induction therefore holds at every state where the target difference is defined, which is the domain asserted in Lemma 3.

Therefore, in the SP model, the optimal value function v^* is concave in each pool size, submodular across types, and satisfies the required anti-diagonal concavity conditions.

Proof of Theorem 2. Let v denote the optimal value function. By Lemma 3, v is componentwise concave, submodular, and anti-diagonally concave. In particular:

- Componentwise concavity: $\Delta_1 v(s)$ is nonincreasing in s_1 and $\Delta_2 v(s)$ is nonincreasing in s_2 .
- Submodularity: $\Delta_1 v(s)$ is nonincreasing in s_2 and $\Delta_2 v(s)$ is nonincreasing in s_1 .
- Anti-diagonal concavity: $\Delta_2 \Delta_{2,-1} v \leq 0$, which (as shown below) implies that $\Delta_{1,-2} v$ is nondecreasing in s_2 . Recall $\Delta_{1,-2} v(s) = v(s) - v(s - e_1 + e_2)$.

For an opportunity O_i arrival at state $s = (s_1, s_2)$, define the gains relative to the baseline of not using the opportunity: $G_1(s) := a_{i1} - \Delta_1 v(s) = a_{i1} - [v(s_1, s_2) - v(s_1 - 1, s_2)]$, $G_2(s) := a_{i2} - \Delta_2 v(s) = a_{i2} - [v(s_1, s_2) - v(s_1, s_2 - 1)]$, and their difference $D(s) := G_1(s) - G_2(s) = (a_{i1} - a_{i2}) - (\Delta_1 v(s) - \Delta_2 v(s))$. Optimality requires choosing the action with the largest payoff, which is equivalent to:

- Not using the opportunity if and only if $G_1 < 0$ and $G_2 < 0$.
- Consuming X_1 if and only if $G_1 \geq 0$ and $G_1 \geq G_2$ (i.e., $D \geq 0$).
- Consuming X_2 if and only if $G_2 \geq 0$ and $G_2 \geq G_1$ (i.e., $D \leq 0$).

At a tie $G_1 = G_2 \geq 0$ both characterizations hold, and the tie convention of Theorem 2 selects X_1 ; the derivation of the action regions below shows that such ties cannot occur in the interior of the X_2 region.

Monotonicity in s_1 . By concavity, $\Delta_1 v$ decreases with s_1 ; hence G_1 increases with s_1 . By submodularity, $\Delta_2 v$ decreases with s_1 ; hence G_2 also increases with s_1 . Thus, for fixed s_2 , as s_1 grows, the optimal action moves from not using the opportunity to either cross-type or matching from the own pool.

Monotonicity of $\Delta_1 v - \Delta_2 v$ in s_2 . Observe that

$$\Delta_1 v(s) - \Delta_2 v(s) = [v(s) - v(s - e_1)] - [v(s) - v(s - e_2)] = v(s - e_2) - v(s - e_1) = \Delta_{1,-2} v(s - e_2).$$

Anti-diagonal concavity (S3) gives $\Delta_2 \Delta_{2,-1} v(s) \leq 0$, which expanded reads $\Delta_2 v(s) \leq \Delta_2 v(s + e_1 - e_2)$ for all s . Substituting $s \mapsto s - e_1 + 2e_2$ and rearranging,

$$\Delta_2 v(s - e_1 + 2e_2) \leq \Delta_2 v(s + e_2),$$

which is equivalent to $\Delta_{1,-2} v(s + e_2) \geq \Delta_{1,-2} v(s)$; that is, $\Delta_{1,-2} v$ is nondecreasing in s_2 . Since $\Delta_1 v(s) - \Delta_2 v(s) = \Delta_{1,-2} v(s - e_2)$, it follows that $\Delta_1 v - \Delta_2 v$ is nondecreasing in s_2 .

Monotonicity of the gains in s_2 .

- By submodularity, $\Delta_1 v$ decreases with s_2 ; thus G_1 increases with s_2 .
- By concavity in the own pool, $\Delta_2 v$ decreases with s_2 ; thus G_2 increases with s_2 .
- Because $\Delta_1 v - \Delta_2 v$ is nondecreasing in s_2 , we have $D = G_1 - G_2$ is nonincreasing in s_2 .

Monotonicity of D in s_1 . Using $\Delta_1 v(s) - \Delta_2 v(s) = \Delta_{1,-2} v(s - e_2)$ and anti-diagonal concavity (S3) ($\Delta_1 \Delta_{1,-2} v \leq 0$), the difference $\Delta_1 v - \Delta_2 v$ is nonincreasing in s_1 ; hence $D = (a_{i1} - a_{i2}) - (\Delta_1 v - \Delta_2 v)$ is

nondecreasing in s_1 . This delivers the three-region structure in s_1 : once $s_1 \geq \tau_1(s_2)$, both $A_1 \geq 0$ and $D \geq 0$ continue to hold for all larger s_1 .

Monotonicity of τ_1 (own-type threshold). Now consider the own-type threshold $\tau_1(s_2)$. Its defining conditions are: $A_1(s_1, s_2) := a_{i1} - \Delta_1 v(s_1, s_2) \geq 0$, $A_2(s_1, s_2) := (a_{i1} - a_{i2}) + (\Delta_2 v - \Delta_1 v)(s_1, s_2) \geq 0$. By submodularity, $\Delta_1 v$ is nonincreasing in s_2 ; hence A_1 is nondecreasing in s_2 . For A_2 , note that $\Delta_2 v - \Delta_1 v = -\Delta_{1,-2} v(s - e_2)$. Since anti-diagonal concavity implies $\Delta_{1,-2} v$ is nondecreasing in s_2 (established above), the quantity $\Delta_2 v - \Delta_1 v$ is nonincreasing in s_2 , and therefore A_2 is nonincreasing in s_2 . Thus, as s_2 increases, condition A_1 becomes easier to satisfy while condition A_2 becomes harder: the set $\{s_1 : A_1 \geq 0\}$ expands with s_2 , which lowers the least admissible s_1 , whereas $\{s_1 : A_2 \geq 0\}$ contracts, which raises it. The two effects move the infimum in opposite directions, so the curvature properties alone do not determine the monotonicity of τ_1 . Under the single-crossing condition (SC) (displayed in Section 3.3; the index i is suppressed here), the constraint A_1 is slack wherever A_2 binds, so $\tau_1(s_2) = \inf\{s_1 \geq 0 : A_2(s_1, s_2) \geq 0\}$, which is nondecreasing in s_2 because A_2 is nonincreasing in s_2 . Proposition 1 gives a sufficient condition for (SC), and Section 4.10 verifies it for the calibration. Hence the infimum $\tau_1(s_2) = \inf\{s_1 \geq 0 : A_1(s_1, s_2) \geq 0 \text{ and } A_2(s_1, s_2) \geq 0\}$ is nondecreasing in s_2 .

Monotonicity of τ_2 (cross-type threshold). For the cross-type threshold $\tau_2(s_2)$, the defining conditions can be rewritten as: $C_1(s_1, s_2) := a_{i2} - \Delta_2 v(s_1, s_2) \geq 0$, $C_2(s_1, s_2) := (a_{i2} - a_{i1}) + (\Delta_1 v - \Delta_2 v)(s_1, s_2) \geq 0$. By submodularity, $\Delta_2 v$ is nonincreasing in s_1 ; thus C_1 is nondecreasing in s_1 . By concavity in its own component, $\Delta_2 v$ is nonincreasing in s_2 ; thus C_1 is also nondecreasing in s_2 . For C_2 , using $\Delta_1 v - \Delta_2 v = \Delta_{1,-2} v(s - e_2)$ and the fact established above that $\Delta_{1,-2} v$ is nondecreasing in s_2 , it follows that $\Delta_1 v - \Delta_2 v$ is nondecreasing in s_2 , so C_2 is also nondecreasing in s_2 . Therefore, for fixed s_1 , if both conditions hold at some s_2 , they continue to hold for larger s_2 . Consequently, the set of s_1 satisfying both conditions expands as s_2 increases, and its infimum $\tau_2(s_2)$ cannot increase. Thus $\tau_2(s_2) = \inf\{s_1 \geq 0 : C_1(s_1, s_2) \geq 0 \text{ and } C_2(s_1, s_2) \geq 0\}$ is nonincreasing in s_2 .

Optimal decision rule. The three action regions now follow from the monotonicity facts established above. First, if $s_1 \geq \tau_1(s_2)$, then $A_1 \geq 0$ and $D \geq 0$ persist for all larger s_1 (both are nondecreasing in s_1), i.e., $G_1 \geq 0$ and $G_1 \geq G_2$: consuming X_1 is optimal. Second, on $\tau_2(s_2) \leq s_1 < \tau_1(s_2)$ (when $\tau_2 < \tau_1$), the conditions $C_1 \geq 0$ hold and persist since C_1 is nondecreasing in s_1 ; moreover $D < 0$ strictly throughout this interval—if $D(s_1, s_2) \geq 0$ held at some s_1 in the interval, then $A_2 = D \geq 0$ and $A_1 = G_2 + D \geq 0$ (as $G_2 = C_1 \geq 0$ there), so $s_1 \geq \tau_1(s_2)$, a contradiction—hence $G_2 \geq 0$ and $G_2 > G_1$: consuming X_2 is optimal, and the X_1 -tie cannot occur there, so the regions are exact and mutually exclusive under the stated tie rule. Third, if $s_1 < \min\{\tau_1(s_2), \tau_2(s_2)\}$, both gains are negative: if $G_1 \geq 0$ held, then either $D \geq 0$, giving $s_1 \geq \tau_1$, or $D < 0$, giving $G_2 > G_1 \geq 0$ and $C_2 = -D > 0$, hence $s_1 \geq \tau_2$ —a contradiction either way; and if $G_2 \geq 0$ held with $G_1 < 0$, then $D < 0$, so $C_1 = G_2 \geq 0$ and $C_2 = -D > 0$, giving $s_1 \geq \tau_2$, again a contradiction. The opportunity is therefore not used. In summary:

- If $s_1 < \min\{\tau_1(s_2), \tau_2(s_2)\}$, both gains are negative and the opportunity is not used.
- If $\tau_2(s_2) \leq s_1 < \tau_1(s_2)$ (provided $\tau_2 < \tau_1$), then $G_2 \geq 0$ and $G_2 > G_1$, so consuming X_2 is optimal.
- If $s_1 \geq \tau_1(s_2)$, then $G_1 \geq 0$ and $G_1 \geq G_2$, so consuming X_1 is optimal.

When $\tau_2(s_2) \geq \tau_1(s_2)$, the interval for matching from the partner pool is empty and the policy reduces to: not using the opportunity for $s_1 < \tau_1$, and matching from the own pool otherwise. The same structure holds symmetrically for opportunity O_2 arrivals, with indices interchanged.

Proof of Proposition 1. If $a_{i1} \geq \Delta_1 v(s)$ for every s with $s_1 \geq 1$, then $A_1(s_1, s_2) = a_{i1} - \Delta_1 v(s_1, s_2) \geq 0$ at

every such state, so the implication in (SC) holds regardless of the sign of A_2 . The economic content is that a single waiting pair is never worth as much as the reward obtained from consuming one on an opportunity, so the planner is never deterred from matching from the own pool by the option value of the pair it consumes.

Proof of Proposition 2. Between consecutive events, a pool holding $N(s) = s_1 + s_2$ pairs accrues the expected discounted holding cost $\kappa N(s) \mathbb{E} \int_0^\tau e^{-\rho t} dt = \kappa N(s) \frac{1-\delta}{\rho} = \theta N(s)$, so the Bellman operator of the κ -model is $T_\theta f = -\theta N + T_0 f$, where T_0 is the original benchmark operator. We argue by induction on value iteration from $v_0 \equiv 0$. Suppose the iterate v_n satisfies the second-order hypotheses of the corresponding lemma. First, $T_0 v_n$ satisfies them: this is precisely the inductive step proved above for Lemma 2 (respectively Lemma 3), and that step uses only the second-order hypotheses on its argument and generic branch rewards—no step invokes monotonicity or signs of first differences of the argument. Second, the affine function $-\theta N$ has identically zero second differences in every direction used by the hypotheses ($\Delta_i \Delta_j (-\theta N) = 0$ for all i, j , and likewise for the diagonal and anti-diagonal operators). Hence $v_{n+1} = T_\theta v_n = T_0 v_n - \theta N$ satisfies the hypotheses, completing the induction; on $\mathbb{N}_0^2$ the properties pass to the fixed point under the standard weighted-sup-norm conditions for linearly growing costs [Puterman, 1994]. Equivalently, since an affine component of the argument can be absorbed into branch-specific reward constants ($f = g + L$ with L affine gives $\max_k \{r_k + f(s'_k)\} = L(s) + \max_k \{(r_k + \ell_k) + g(s'_k)\}$ with $\ell_k = L(s'_k) - L(s)$ constant), the κ -model is the generic-reward benchmark with shifted rewards, to which Lemmas 2–3 apply directly. The threshold characterizations of Theorems 1–2 are derived from the lemmas’ second-order properties, together with the single-crossing hypothesis (SC) that Theorem 2 carries explicitly, so they hold for the κ -model with its own value function and hence its own threshold functions; (SC) is a first-order condition and must be checked for the κ -model’s value function, as Section 4.11 does at the reported calibrations. $\square$

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used ChatGPT (OpenAI) and Claude (Anthropic) to obtain editorial suggestions, assist in revising the exposition, and perform supplementary checks of the manuscript’s mathematical presentation and internal consistency. The authors independently verified all analytical results, proofs, computations, and citations, reviewed and edited all AI-assisted content, and take full responsibility for the content of the published article.

References

- Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton R. Featherstone, and Ömer Karaduman. Market failure in kidney exchange. *American Economic Review*, 109(11):4026–4070, 2019. doi: 10.1257/aer.20180771.
- Leila Agha, Brigham Frandsen, and James B. Rebitzer. Fragmented division of labor and healthcare costs: Evidence from moves across regions. *Journal of Public Economics*, 169:144–159, 2019.
- Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Unpaired kidney exchange: Over-

- coming double coincidence of wants. *Journal of Political Economy*, 128(6):2329–2372, 2020. doi: 10.1086/705736.
- Itai Ashlagi and Alvin E Roth. New challenges in multihospital kidney exchange. *American Economic Review*, 102(3):354–359, 2012.
- Itai Ashlagi and Alvin E. Roth. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863, 2014. doi: 10.3982/TE1357.
- Itai Ashlagi and Alvin E. Roth. Kidney exchange: An operations perspective. *Management Science*, 67(9):5455–5478, 2021. doi: 10.1287/mnsc.2020.3954.
- Itai Ashlagi, Felix Fischer, Ian A. Kash, and Ariel D. Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 91:284–296, 2015.
- Mariagiovanna Baccara, SangMok Lee, and Leeat Yariv. Optimal dynamic matching. *Theoretical Economics*, 15(3):1221–1278, 2020. doi: 10.3982/TE3740.
- David Blackwell. Discounted dynamic programming. *Annals of Mathematical Statistics*, 36(1): 226–235, 1965. doi: 10.1214/aoms/1177700285.
- Danny Blom, Bart Smeulders, and Frits C. R. Spijksma. Rejection-proof mechanisms for multi-agent kidney exchange. *Games and Economic Behavior*, 143:25–50, 2024.
- Randall D. Cebul, James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba. Organizational fragmentation and care quality in the U.S. healthcare system. *Journal of Economic Perspectives*, 22(4):93–113, 2008.
- David M. Cutler. The incidence of adverse medical outcomes under prospective payment. *Econometrica*, 63(1):29–50, 1995.
- Chen Hajaj, John Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Philip J. Held, Frank McCormick, Akinlolu Ojo, and John P. Roberts. A cost-benefit analysis of government compensation of kidney donors. *American Journal of Transplantation*, 16(3):877–885, 2016.
- Krista L. Lentine et al. Managing the costs of kidney paired donation: A survey of contemporary U.S. practice and challenges. *Clinical Transplantation*, 40(2), 2026. doi:10.1111/ctr.70466.
- David A. Mast, William Vaughan, Stephan Busque, et al. Managing finances of shipping living donor kidneys for donor exchanges. *American Journal of Transplantation*, 11(9):1810–1814, 2011.
- NHS Blood and Transplant. Annual report on living donor kidney transplantation 2024/25. Technical report, NHS Blood and Transplant, 2025.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. John Wiley & Sons, New York, 1994. ISBN 978-0-471-72782-8.

- Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.
- Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review*, 97(3):828–851, 2007.
- Susan L Saidman, Alvin E Roth, Tayfun Sönmez, M Utku Ünver, and Francis L Delmonico. Increasing the opportunity of live kidney donation by matching for two-and three-way exchanges. *Transplantation*, 81(5):773–782, 2006.
- Dorry L. Segev, Sommer E. Gentry, Daniel S. Warren, Brigitte Reeb, and Robert A. Montgomery. Kidney paired donation and optimizing the use of live donor organs. *JAMA*, 293(15):1883–1890, 2005. doi: 10.1001/jama.293.15.1883.
- Nancy L. Stokey, Robert E. Lucas, Jr., and Edward C. Prescott. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA, 1989.
- Keith F. Teltser. Do kidney exchanges improve patient outcomes? *American Economic Journal: Applied Economics*, 11(3):427–453, 2019.
- Paul I Terasaki, David W Gjertson, and J Michael Cecka. Paired kidney exchange is not a solution to abo incompatibility. *Transplantation*, 65(2):291, 1998.
- Panos Toulis and David C. Parkes. Design and analysis of multi-hospital kidney exchange mechanisms using random graphs. *Games and Economic Behavior*, 91:360–382, 2015. doi: 10.1016/j.geb.2015.01.001.
- United States Renal Data System. 2018 USRDS annual data report: Epidemiology of kidney disease in the United States. Technical report, National Institutes of Health.
- Stefanos A. Zenios, E. Steve Woodle, and Lainie Friedman Ross. Primum non nocere: Avoiding harm to vulnerable wait list candidates in an indirect kidney exchange. *Transplantation*, 72:648–654, 2001.
- M. Utku Ünver. Dynamic kidney exchange. *The Review of Economic Studies*, 77(1):372–414, 2010.