Optimal Dynamic Matching under Local Compatibility: An Application to Kidney Exchange

Ömer Faruk Şahin* Duygu Sili † M. Utku Ünver‡ Özgür Yılmaz§

July 2025

Abstract

In the past two decades, the design and implementation of living donor kidney exchange clearinghouses have been a major success story in market design. Instead of batching and optimizing exchanges over a fixed pool of incompatible patient-donor pairs, the busiest programs now operate dynamically, matching pairs as they arrive. This feature has also sparked interest in dynamic matching mechanisms. Yet for general matching problems with high-dimensional state spaces, a full characterization of optimal dynamic mechanisms remains elusive, and only approximate solutions are known.

We develop a new methodology to characterize and compute dynamically optimal mechanisms for bilateral matching over arbitrary state spaces, provided that compatibility between agent types follows a linear spatial structure. This technique applies to optimal dynamic kidney exchange and extends to other spatial matching problems. Our approach leverages second-order properties of the value function, extending recent advances in Markov Decision Processes and queueing systems, which traditionally focus only on substitutable components.

Keywords: Dynamic matching, kidney exchange, dynamic exchange, spatial economics, Poisson arrival, dynamic optimization, Markov Decision Process, discrete convex analysis, $\mathcal{D}$-multimodularity, superconcavity, componentwise concavity, submodularity, supermodularity.

JEL Classification Numbers: C78, D47, C70, D78, C61

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

†University of Messina, Via dei Verdi 75, Messina, 98122 Italy; *e-mail address:* dsili@unime.it.

‡Boston College, Chestnut Hill, MA, 02467 USA; *e-mail address:* unver@bc.edu.

§Koç University, İstanbul, 34450 Turkey; *e-mail address:* ozyilmaz@ku.edu.tr.

1 Introduction

While significant progress has been made over the last 15 years in studying dynamic matching markets—which evolve through the arrival of agents, a common theme in the literature is the difficulty of conducting exact analysis due to the high number of agent types. As a result, most of the literature involving high-dimensional state spaces focuses on approximate or asymptotic analysis, regardless of the presence of other complicating factors such as asymmetric information, so long as arrival uncertainty is present.

This difficulty arises because standard dynamic programming techniques typically become intractable as the dimensionality of the state space increases. Although low-dimensional models can sometimes offer useful economic intuition, most frameworks cannot accommodate more than two types of agents on each side of the market without resorting to approximations or heuristic rules, which are then justified by demonstrating bounded inefficiency—or, more optimistically, asymptotically vanishing inefficiency as market size grows.

In this paper, we introduce a new technique for characterizing optimal mechanisms in dynamic matching problems involving an arbitrary number of agent types, as long as mutually beneficial trades or matches occur *locally*, akin to the differentiated *linear city* model of Hotelling (1929). While an exact characterization remains elusive for arbitrary state spaces, our analysis identifies second-order properties of the optimal value function that hold under such local compatibility structures, which we then use to characterize optimal mechanisms.

We apply this framework to characterize optimal dynamic kidney exchange mechanisms under certain assumptions. While we do not pursue directly, our methodology can also be applied to study other problems such as public housing allocation in which houses and families form two sides of a market and each family is deemed compatible with houses of their size or slightly larger, but not any smaller house or any much larger house—due to waste concerns. Furthermore, our model potentially provides a stylized economic framework for studying spatial economics problems, such as traditional geographically differentiated trade networks and modern on-demand services involving ride-sharing and food delivery.

We start by describing the basics of our model. We assume that there is a linear structure of agent types $1, \ldots, n$ such that an agent of type $i = 2, \ldots, n-1$ can only be matched with an agent of type $j = i-1, i+1$ in a match denoted as ij to generate a matching surplus of a_{ij}, while the type at the ends of the linear order $i = 1, n$ can either be matched by itself with matching surplus a_i or with an agent of its neighbor type $j = 2$ or $j = n-1$ respectively, with matching surplus a_{ij}.

We study the optimal matching mechanisms in a model that assumes Poisson arrivals of agents. Agents incur waiting costs over time, and their preferences depend on the compatibility of their match and these costs. In particular, we analyze the optimal control, aiming to maximize the total discounted match surplus.

This framework essentially embeds the model introduced by Ünver (2010) for dynamic kidney exchange as a special case. Each patient-donor pair can be modeled as a single agent. If one assumes the existence of a sufficiently large number of so-called underdemanded pairs—those with less desirable blood type donors relative to their patients—an assumption motivated by the mathematical structure of blood-type incompatibility and a sufficiently long time lapse since the start of a relatively large kidney exchange clearinghouse, feasible kidney exchanges can be modeled with only $n = 4$ types of agents that are scarce and arrive over time; larger size kidney exchanges reduce to bilateral matching of consecutive types that are not abundant. However, Ünver (2010) only solves a constrained version of the optimal dynamic kidney exchange problem. Specifically, the second main result assumes that overdemanded pairs (the end agent types in our model $i = 1$ and $i = n = 4$)—those with more desirable blood type donors relative to their patients—are matched immediately upon arrival.[1] While this assumption may be plausible in some instances, it restricts the generality of the optimal mechanism considerably.

When the assumption of immediate matching for overdemanded pairs—i.e., agents of type $i = 1$ and $i = n = 4$ is relaxed, the full characterization of the optimal dynamic kidney exchange mechanism becomes more complex. The control state space of the model expands to four dimensions, causing techniques used in Ünver (2010), which analyzed the state space for two dimensions, to be impractical. Via our new techniques and analytical results, we have overcome this issue and do not need to make such assumptions. We discuss the details of the kidney exchange application in our model and how it can be embedded in the next section.

Our methodological contribution. The primary methodological contribution of our paper is the development of a new set of tools for dynamic matching frameworks that can handle optimal matching control in models with multiple state variables.This methodology leverages the second-order properties of the value function in dynamic programming by extending recent advances in queueing theory and Markov Decision Processes (MDPs) to the matching framework. Unlike queueing models, where substitutability is the only decision feature, matching models involve central authorities optimizing trade-offs between decisions that exhibit complementarity and substitutability among different agents.

Our main result demonstrates that an optimal mechanism exhibiting a nuanced multi-threshold structure exists. These thresholds govern the stockpiling of different agent types, al-

[1] We also show in the Online Appendix that an auxiliary result, Proposition 3, used to determine a range of pair-type arrival rates to support this assumption in Ünver (2010), Assumption 2, has an error. Thus, it is not easy to pinpoint underlying fundamentals to guarantee that overdemanded pairs will be immediately matched in an unconstrained optimal mechanism.

lowing the clearinghouse to forgo immediate exchanges in anticipation of better future matching opportunities. However, once the stock of a given agent type exceeds a certain threshold, further stockpiling ceases, and these agents are matched immediately when an opportunity arises, as long as their numbers remain above the threshold. These thresholds are, in general, a function of other current state variables.

We derive this multi-threshold structure by proving that the value function in the optimal control problem satisfies (discrete) concavity in the state variables, where each variable corresponds to the number of agents of a specific type.[2] Concavity explains why it becomes optimal, at some point, to prioritize immediate matches over continued waiting: the marginal value of waiting falls below the immediate gain from conducting an exchange when sufficient stock of a certain agent type exists.

A direct proof of the concavity of an implicitly defined optimal value function is elusive for dynamic matching problems. Instead, we develop a new technique.

The core of our argument uses the fundamental theorem of discounted dynamic programming, which establishes that a unique optimal value function exists and can be computed via value iteration using the optimal MDP operator, a contraction mapping, starting from an arbitrary initial value function.

Our approach decomposes the optimal MDP operator into a sequence of operations involving time *discount*, an *arrival*, the expected value (*uniformization*), and various matching decisions among different agents, depending on the state of the exchange pool. More specifically, a *matching operator* is defined as an optimal operator in deciding whether to conduct a particular exchange or not. For example, take one of the $n+1$ possible basic matching decisions, matching agent type $i = 2, \dots n-1$ with an agent of its neighbor type $j = i-1$ or $j = i+1$. It may also let agent of type i to wait and not be matched at all. Matching transitions the state to a lower one, by decreasing one or more types of agents waiting but provides an immediate matching surplus of a_{ij}. Thus, a basic matching operator for each matching type ij finds the maximum value of two decisions, waiting or immediate matching, incorporating the optimal decision needed. To find the optimal matching decision, we use all basic matching operators in a sequence at a given state, one at a time (we show the order in which we apply them does not matter, and the decomposition is commutative).

We want to show that each event operator propagates the concavity of the value function. A naive approach would involve showing that, starting from an arbitrary concave value function,

[2] A real-valued function f defined on the n-dimensional non-negative integer state space $\mathbb{N}^n$ is concave in its i'th component, if for each $x \in \mathbb{N}^n$, $f(x+2e_i) - f(x+e_i) \leq f(x+e_i) - f(x)$, where $e_i \in \mathbb{N}^n$ is zero in all its components except its i'th component, which is 1. It is (componentwise) concave if it is concave in all its elements.

each value iteration results in a concave function and, thus, in the limit, converges to a concave "optimal" value function. However, concavity alone is too weak a property to be propagated by these event operators. Thus, one cannot simply start with an arbitrary concave value function and propagate this property to the limiting optimal value function, as it does not follow from the definition of our operators that concavity alone will propagate.

Instead, we demonstrate that a combination of multiple second-order properties of the value function, including concavity, is propagated. This propagated property corresponds to an abstract second-order characteristic of discrete functions known as *multimodularity* for a set of vectors (cf. Hajek, 1985, Altman et al., 2003).[3]

The crux of our approach relies on constructing a minimal set of abstract state vectors, known as a *multimodular basis*, which spans all states and through which multimodularity is propagated by the event operators. As a result, the economically meaningful second-order properties of the optimal value function, including concavity, are implied by this property.

Complementarity and substitutability. The decision structure also reveals an important underlying feature of our matching model. Given an agent of type i, agents of types $i-1$ and $i+1$ are *complements*, as the agent can be matched with either of them—whenever these types are well-defined. In contrast, agents of types $i-2$ and $i+2$ are *substitutes* for the agent, as the complementary types $i-1$ and $i+1$ can instead be matched with them, respectively. By iterating this logic, we can say that agents of types $i \pm k$ are complements to an agent of type i when k is odd, and substitutes when k is even.

This structure causes the multimodular matching basis we introduce to differ fundamentally from the queueing basis of Koole (1998), which is almost universally adopted in the queueing literature. The queueing basis assumes only substitutable changes to the state variables, whereas our matching basis incorporates both complements and substitutes, reflecting the richer structure of kidney exchange and matching problems more generally.[4]

One way to appreciate the complexity of such a matching problem is by imagining it as the output of a production function (see, e.g., Agarwal et al., 2019 for this interpretation in kidney exchange). The inputs are agents, and not all agents can produce the output—a match.

In contrast, in the queueing theory literature that utilizes Koole's methods, the analog of the production function features only substitutable inputs.

It turns out that the $n+1$ basic matching decisions, which we defined earlier as $1,\ 12, \ldots,\ (i-1)i,\ i(i+1),\ \ldots,\ (n-1)n,\ n$, form a *multimodular matching basis*. These decisions are denoted

[3] A real-valued function f defined on a discrete state space $\mathcal{S}$ is $\mathcal{D}$-multimodular for a set $\mathcal{D} \subset \mathcal{S}$ if, for any $s \in \mathcal{S}$ and $u, v \in \mathcal{D}$, $f(s+u) + f(s+v) \leq f(s) + f(s+u+v)$ (assuming all arguments remain in $\mathcal{S}$).

[4] In Figure 1 in Section 5.3, we illustrate this structure in the kidney exchange context.

as vectors that transition the state by subtracting them from the current state, and any n of them are linearly independent. Their signs are normalized so that they sum to zero, ensuring that the $(n+1)^{\text{st}}$ operator can be expressed as the negative of the sum of the first n. This matching basis spans (via vector addition and subtraction) all possible state transitions and all states. The existence of such a zero-sum matching basis is central to proving that the second-order properties of the value function are preserved under the event operators.

We then prove that if a function satisfies multimodularity for the matching basis, it continues to satisfy this property after a matching operator is applied to the function. Since all states are spanned by the basis matching vectors, then multimodularity is *closed* under any compounded matching operations. Similarly, other event operators (discount, arrival, and uniformization) also preserve multimodularity in this manner. Therefore, starting with a multimodular initial function in the value iteration guarantees that the limiting value function, which is the unique optimal value function, is also multimodular (Theorem 2). Then our next result, Theorem 3, shows that multimodularity for the matching basis captures all economically relevant second-order properties of the optimal value function, including concavity. This ensures the existence of a nuanced multi-threshold mechanism that generates the optimal value.

The structure of the dynamically optimal mechanism. Concavity is not the only important second-order property of the optimal value function implied by multimodularity. Thanks to additional second-order properties, we can derive further insights into the structure of dynamically optimal mechanisms. These properties allow us to characterize the trade-offs between waiting and matching decisions with greater precision.

We also prove in Theorem 3 that the optimal value function is *supermodular* for complementary types i and $i \pm k$ where k is odd, and *submodular* for substitutable types i and $i \pm k$ where k is even. Furthermore, the optimal value function is *superconcave*, meaning that two agents of the same type are (weakly) better substitutes than two agents where the second agent is of a different type.[5]

The optimal mechanism, on the other hand, has an intuitive structure: when an agent of type i arrives (assuming for ease of exposition that $i \in \{2, \ldots, n-1\}$) at a reachable state s, at most one exchange from each type $1, 12, \ldots, (n-1)n, n$ is conducted. Moreover, the exchanges that can be conducted involve type i or, in a particular way, the complements of this agent (Theorem 4). We also identify how the arrival of two agents of types i and j can trigger exchanges when j alone would not lead to any exchange (Theorem 5). Then, we show that for each agent type i,

[5] A real-valued function f defined on $\mathbb{N}^n$ is ij-supermodular for any two components i and j if, for each $x \in \mathbb{N}^n$, $f(x+e_i) - f(x) \leq f(x+e_i+e_j) - f(x+e_j)$. It is ij-submodular if the previous inequality is reversed. It is ij-superconcave if, for each $x \in \mathbb{N}^n$, $f(x+2e_i) - f(x+e_i) \leq f(x+e_i+e_j) - f(x+e_j)$. A function is superconcave if it is ij-superconcave for all agents of components i and j.

the optimal mechanism follows a threshold structure when the number of other agent types is fixed (Theorem 6): if the number of type-i agents exceeds a threshold, then new exchanges are conducted in the optimal mechanism.

Finally, we focus on the optimal mechanism in kidney exchange—with only four agent types. We give a full characterization of the multi-threshold structure when the A blood-type patient and B blood-type donor pair ($A-B$) arrival rate is substantially higher than the $B-A$ arrival rate (Theorem 7). These types constitute the middle two types—2 and 3—in the four-type state space. The optimal mechanism under general arrival rates is based on a mere extension of this structure, and for the purposes of brevity, we defer its analysis and discussion to the Online Appendix.

Related Literature. We have already mentioned the most relevant literature on queueing theory and MDPs, and how they relate to our paper. It is useful to highlight in more detail what our work contributes to the dynamic matching literature, as this area has been the focus of several papers, beginning with Ünver (2010). Due to the complex nature of the arrival problem, all papers make certain simplifications to achieve tractable results and balance trade-offs.

Ünver (2010) and our paper both assume away tissue-type incompatibility of patients with other donors, which effectively serves as a large-market and limit assumption. Undoubtedly, some patients with a high tissue-incompatibility probability 0.99—known as *very highly sensitized patients*—will accumulate in the pool. In our approach, these pairs are assumed to remain unmatched and are excluded from the analysis. For all other patients, however, our analysis is a useful approximation and a useful approximation in the limit in general.

In contrast, much of the other literature ignores blood types and instead focuses on an abstract notion of "hard-to-match" patients, characterizing compatibility as a probabilistic event to capture highly sensitized patients. This approach significantly expands the dimensions of the ex-post state space, forcing authors to rely on approximate analyses with error bounds.[6]

Our framework, however, allows us to fully characterize the optimal mechanisms, as compatibility in our setup is deterministic rather than probabilistic. Starting with Ünver (2010), the recent consensus in this literature has been that almost greedily maximizing exchanges, rather than engaging in dynamic optimization, results in a minimal loss (for example, see Anderson et al., 2017, Kerimov et al., 2025 and related works).

In contrast, our paper shows that under optimal mechanisms, overdemanded pairs—which always enable immediate matches—are not necessarily matched immediately, as dynamic optimization may require waiting to execute more effective matches—departing from prior litera-

[6] An exception is Sönmez et al. (2020), which models both blood-type and tissue-type incompatibility but uses a high-traffic fluid model (the dynamic analog of continuum models) for tractability that leads to a single steady state.

ture.[7]

A closely related paper to ours is Baccara et al. (2020), which examines a dynamic arrival setting and, much like our work, identifies the best complementary matches, albeit in a two-sided matching market inspired by child adoption. While their work and ours share the ability to derive optimal mechanisms in dynamic settings, the underlying approaches differ substantially. Due to the complexities of higher-dimensional state spaces, their analysis is limited to two types, whereas our framework accommodates multiple state dimensions. Moreover, their focus is primarily on incentive issues, distinguishing their contributions from ours.

Several recent papers in economics have studied dynamic assignment problems with features such as queueing and waiting lists. Bloch and Cantala (2017) analyzes the dynamic assignment of objects to arriving agents, proposing approximately optimal policies based on waiting times. Che and Tercieux (2021) examines how to optimally design queue structures to influence agent behavior, with a focus on incentive alignment in queue selection. Leshno (2022) investigates dynamic rules for overloaded waiting lists, particularly in organ allocation, and shows that dynamic rules outperform static ones under congestion. While these papers share a focus on dynamic assignment mechanisms, they consider the assignment of objects to agents with an emphasis on incentives. In contrast, our paper studies dynamic bilateral matching with spatially structured complementarities, such as those arising in kidney exchange.

To date, no precedent to our methodology allows for modeling and solving optimal mechanisms with multiple state dimensions in economics, except brute-force techniques, which are only practical for state spaces with, at most, a couple of dimensions. Previous approaches that we rely on in operations research primarily addressed substitutable decisions, such as in queueing systems, as discussed earlier. However, matching inherently involves complementarities, making our approach fundamentally distinct.

Finally, it is worth noting that our paper also contributes to the dynamic matching literature more broadly, which includes works such as Kurino (2020) and Doval (2022) modeling dynamic stability—an important tenet in two-sided matching markets. Doval (2025) provides an excellent survey of the economics papers on dynamic matching.

2 Kidney Exchange: Motivating the Dynamic Matching Problem

In this section, we introduce the kidney exchange paradigm (Rapaport, 1986, Roth et al., 2004, 2005, 2007), which accounts for more than one-fifth of living-donor transplants in the U.S. as of 2024 (Sönmez and Ünver, 2025). Kidney exchange serves as the main application of the more general framework that we introduce in Section 3.

[7] If agents expire while waiting and these expiration times are relatively predictable, mechanisms that rely on "patience" may mimic optimal mechanisms in the limit (Akbarpour et al., 2020).

We start by explaining the two medical compatibility requirements for kidney transplantation—the preferred treatment for the most severe forms of kidney disease.

The first requirement is ABO blood-type compatibility. Blood type is determined by the presence of A or B blood proteins (called antigens). A patient can receive a kidney from a donor if the patient has all the ABO antigens present in the donor. Thus, an O blood type patient can receive a kidney only from an O donor; an A patient can receive from an O or A donor; a B patient can receive from an O or B donor; and an AB patient can receive from any blood type donor.

The second requirement is tissue-type compatibility. A patient cannot receive a transplant from a donor if they have high levels of antibodies against the donor's tissue types.[8]

A patient is compatible with her paired living donor—typically a blood relative or close loved one—if they are both ABO- and tissue-type-compatible. In such cases, a direct transplant is pursued, since the two alternatives have poorer outcomes: deceased donation requires long waiting times—due to global shortages—for a typically lower-quality organ, and dialysis entails significantly lower quality of life.[9]

On the other hand, a patient with an incompatible paired living donor becomes a candidate for kidney exchange, in which her donor is swapped with a compatible donor from another pair. We refer to the patient and her incompatible donor as a **pair**. The type of a pair is denoted by $X - Y$, where X is the blood type of the patient and Y is the blood type of her donor.[10]

An **exchange** is a set of (at least two) pairs such that the patient in each pair receives a kidney from the donor of another pair. A **matching** is a collection of such exchanges, where each pair participates in at most one exchange. A matching that includes the maximum possible number of pairs is called **maximal**.

Let $\mathcal{T}$ denote the set of all pair types. We partition $\mathcal{T}$ into four subsets: **underdemanded**, **overde-**

[8] Zenios et al. (2001) report the probability of a patient being tissue-type incompatible with a random donor as 0.11, while Sönmez et al. (2020) note that this rate has declined over time in the U.S. deceased-donor waitlist. However, recent studies show that patients more likely to be tissue-incompatible with random donors tend to accumulate in the U.S. kidney exchange pools (see Agarwal et al., 2019).

[9] Incompatible transplants are performed in some countries after removing the patient's antibodies and administering other immunosuppressive therapies. These procedures carry a higher short-term risk of graft loss in blood-type-incompatible (but tissue-type-compatible) transplants (Massie et al., 2020). Nonetheless, data suggest their long-term survival rates are comparable to those of compatible transplants. In contrast, tissue-type-incompatible transplants do not perform as well (Schinstock et al., 2019). These treatments are also expensive, time-consuming, and offered only at select hospitals due to their complexity. In the U.S. and many other countries, incompatible transplantation is either disfavored relative to kidney exchange or not practiced at all. In this paper, we follow most of the literature in assuming that all feasible transplants must be fully compatible.

[10] Most kidney exchange systems match incompatible pairs only in "circular" exchanges, where each donor gives to the next patient in the cycle. Because transplants must be done simultaneously to avoid failed future steps, these systems face logistical limits. In contrast, chains initiated by altruistic donors (a practice used primarily in the U.S.) or deceased donors (a practice used less frequently) can proceed sequentially and reach more patients (see Roth et al., 2006, Rees et al., 2009, Agarwal et al., 2019, Furian et al., 2020). On non-simultaneous exchange proposals, see Ausubel and Morrill (2014), Akbarpour et al. (2024).

manded, **self-demanded**, and **reciprocally demanded** pair types, denoted by $\mathcal{P}^U$, $\mathcal{P}^O$, $\mathcal{P}^S$, and $\mathcal{P}^R$, respectively:

$$\mathcal{P}^U = \{O-A, O-B, O-AB, A-AB, B-AB\},$$
$$\mathcal{P}^O = \{A-O, B-O, AB-O, AB-A, AB-B\},$$
$$\mathcal{P}^S = \{A-A, B-B, O-O, AB-AB\},$$
$$\mathcal{P}^R = \{A-B, B-A\}.$$

The naming of these sets reflects the following characteristics:

- Underdemanded pairs require an overdemanded pair to be matched, as the donor's blood type is less desirable than the patient's.
- Overdemanded pairs, conversely, can be matched with their own type or with other types in various combinations, though some may be wasteful.
- Self-demanded pairs are naturally matched within their own type, as both the donor and patient share the same blood type. Matching them with overdemanded pairs is also possible but may lead to inefficiency.
- Reciprocally demanded pairs ($A-B$ and $B-A$) can be directly matched with each other in a two-way exchange. Using an overdemanded type to facilitate such a match is typically wasteful.

The following two assumptions are standard in the literature and are motivated by long-run pool evolution under asymmetric arrival and matching rates. These are driven by the structure of blood-type compatibility, particularly in the limit when tissue-type incompatibility becomes negligible (see Roth et al., 2007 and Ünver, 2010 for detailed discussion):

Assumption 1. ***(Limit Assumption)*** *No patient is tissue-type incompatible with the donor of another pair.*

Assumption 2. ***(Long-Run Assumption)*** *Under any dynamic matching mechanism, in the long run, there is an arbitrarily large number of underdemanded pairs of each type in the pool at any time.*

We also make the following simplifying assumption with minimal loss of generality. When two or more self-demanded pairs of the same type $X-X$ exist, they can always be matched immediately in a two-way exchange. Moreover, if only one such pair exists, it can be inserted at no cost into an exchange of the form $(\ldots, V-W, Y-Z, \ldots)$ between the $V-W$ and $Y-Z$ pairs, where W can donate to X and X can donate to Y. Thus, self-demanded pairs are not needed for control. For clarity, we assume that such pairs do not arrive (see also Ünver, 2010).

Assumption 3. ***(No Self-Demanded Pairs)*** *There are no self-demanded pair types available for* *exchange.*

We denote the set of overdemanded pairs that can participate in *three-way* or *four-way exchanges* involving an $A-B$ pair by $\mathbb{O}_{A-B}$. When an $A-B$ pair is not used, the maximum feasible exchange size decreases by one, resulting in either two-way or three-way exchanges, respectively. For clarity, we use the notation $\mathbb{O}_{A-B}$ specifically for these overdemanded pairs, which should not be confused with blood type O (the latter does not appear in the following sections). Specifically, we define

$$\mathbb{O}_{A-B} = \{B-O, AB-O, AB-A\}.$$

Similarly, the set of overdemanded pairs associated with $B-A$ is defined as

$$\mathbb{O}_{B-A} = \{A-O, AB-O, AB-B\}.$$

The reasoning is as follows. As shown by Roth et al. (2007), $\mathbb{O}_{A-B}$ pairs (excluding $AB-O$) can form three-way exchanges using an $A-B$ pair and an underdemanded pair, such as $(A-B, B-O, O-A)$ and $(A-B, B-AB, AB-A)$. Without an $A-B$ pair, these exchanges reduce to $(B-O, O-B)$ and $(AB-A, A-AB)$, respectively. Symmetrically, the corresponding three-way exchanges for $\mathbb{O}_{B-A}$ can be formed similarly.

The $AB-O$ pair, however, behaves differently. It can generate a four-way exchange when combined with an $A-B$ or $B-A$ pair, such as $(AB-O, O-A, A-B, B-AB)$ or $(AB-O, O-B, B-A, A-AB)$. Alternatively, it can form a three-way exchange on its own, such as $(AB-O, O-A, A-AB)$ or $(AB-O, O-B, B-AB)$, when paired with underdemanded pairs. Thus, the surplus from an $AB-O$ pair is 4 in a four-way exchange and 3 in a three-way exchange.

Since blood type AB is rare in most countries, controlling $AB-O$ as a separate state is unnecessary and does not affect the results. To simplify, we classify any arriving $AB-O$ pair arbitrarily as either a $\mathbb{O}_{A-B}$ or a $\mathbb{O}_{B-A}$ pair to avoid confusion. Based on this, we make the following assumption:

Assumption 4. ***(Rare** AB **assumption)*** *The maximum exchange size including an $\mathbb{O}_{A-B}$ pair is 3 when also an $A-B$ pair is used and 2 when it is not. Similarly, the maximum exchange size of including an $\mathbb{O}_{B-A}$ pair is 3 when also a $B-A$ pair is used and 2 when it is not.*

Trade-offs and decisions regarding optimal control. When a $B-A$ pair exists in the pool and there is no $A-B$ pair, suppose a composite-type $\mathbb{O}_{B-A}$ pair such as of type $A-O$ arrives: should we match the $A-O$ pair in a two-way exchange (e.g., $(A-O, O-A)$ with an abundant $O-A$ pair) with an immediate surplus 2, and retain the $B-A$ pair to save a future arriving $A-B$ pair (with a surplus of 2 in the future) or should we conduct the three-way exchange $(B-A, A-O, O-B)$ now with an immediate surplus of 3? Thus, if the waiting value is more than 1 at the current state of the pool, we will retain the $B - A$ pair and match it otherwise.

The decision can be much more complicated, as the same decision applies to also an $\mathbb{O}_{A-B}$ pair when it arrives: will we match it immediately in a two-way exchange or wait for a three-way exchange, holding on to it with a future coming $A - B$ pair (and an abundant underdemanded pair).

If the optimal value function is monotonic and concave in the stock of each pair type then a multi-threshold policy will exist: initially, we will hold on to $\mathbb{O}_{A-B}$ or $B-A$ pairs when their stocks are small, and eventually, as their stocks grow, the value of waiting will fall below 1, and we will need to conduct the immediate higher surplus match whenever feasible. The situation may be even more complex if we have to do the symmetric control for $A - B$ pairs depending on the arrival rates of these types.

3 General Model

We next introduce our general dynamic exchange model, which covers kidney exchange as a special case.

3.1 Dynamic matching as a Markov Decision Process

Let $\mathcal{T} = \{1, 2, \ldots, n\}$ be a set of agent types such that an agent of type $i \neq 1, n$ can be matched only with an agent of types $j = i - 1$ or $j = i + 1$, her immediate neighbors, and an agent of type $i = 1$ or $i = n$ can be matched by herself or type $j = 2$ or $j = n - 1$, respectively. We refer to such an ij as an **exchange** such that when $i = j$ we match an agent of type i by herself. We refer to agents (or types) that can be matched with each other as **compatible**.

When compatible agents of types i and j are matched, a surplus of $a_{ij} \in \mathbb{R}_{++}$ is realized. When an agent of type $i = 1$ or $i = n$ is matched by herself a surplus of $a_i \in \mathbb{R}$ is realized.[11] In this case, we sometimes refer to a_i as a_{ii} when convenient.

For each agent type i, we denote the probability that an arriving agent is of type i by p_i, and the associated probability distribution by $p = (p_i)_{i \in \mathcal{T}}$.

We assume that agents arrive over time with a stochastic (discrete) Poisson arrival process in continuous time. Let λ_i denote the arrival rate of agent type $i \in \mathcal{T}$; that is the expected number of agents of type i that arrive per unit time. We formulate the dynamic matching under this arrival structure as a *Markov Decision Process (MDP)*.

Let ρ denote the continuous discount rate. Thus, the surplus generated by realizing K exchanges of types $i_1 j_1, \ldots, i_K j_K$ at time t is $\sum_k a_{i_k j_k}$. Let $M(t)$ denote the surplus of agents matched at time

[11] If agents of types $i = 1$ or $i = n$ can indeed be feasibly matched by themselves then $a_i > 0$ while in some problems this may not be feasible and in that case we allow $a_i < 0$, without loss of generality.

t. We assume that the number of exchanges that can be conducted at each time t is bounded by a capacity and hence, $M(t)$ is also uniformly bounded above (see also Footnote 14). The goal of the central authority is to maximize the *expected surplus*, denoted by ES and defined as

$$ES = \mathbb{E}_{t=0} \sum_{t \in (0,\infty)} M(t) e^{-\rho t}.$$

Note that, although the index set $t \in (0, \infty)$ is uncountable, since $M(t)$ almost surely has countable support and is bounded, this value is finite.

Definition 1. A dynamic exchange mechanism is **dynamically optimal** if it maximizes the expected surplus.

The decision associated with this maximization problem is when to match arriving agents and which agents to match. If an arriving agent is decided not to be matched, she joins the pool of waiting agents. If an ij exchange is executed at time t, the total surplus increases by $a_{ij} e^{-\rho t}$, and if $i \neq j$, then these two matched agents are removed, and if $i = j$, then the matched type i agent is removed.

First, we observe that this is indeed a Markovian process: by the memorylessness property of Poisson arrival, the process starting from any pool with a specified composition is the same regardless of time it takes for the process to achieve this pool. Thus, our state space for this process consists of all possible pools of agents. Therefore, a pool can be denoted by its state, defined as follows:

Definition 2. The set of **states** is the set of non-negative valued integer n-tuples, $\mathcal{S} = \{s = (s_1, s_2, \ldots, s_n) : s_i \in \mathbb{N}\} = \mathbb{N}^n$, where for each state $s = (s_1, s_2, \ldots, s_n)$, the numbers s_i , denote the number of agents of type i.

We illustrate our definitions so far within the kidney exchange setting:

Example 1. *In the kidney exchange setting, an agent is a patient-donor pair. Since we assume there are no self-demanded pairs (Assumption 3) and that there are infinitely many underdemanded pairs (Assumption 2), the only types that are relevant for the process are overdemanded pairs and reciprocal pairs. Since we focus on exchanges of size at most 3, types of overdemanded pairs also do not matter for the process, in so far as they can conduct a three-way exchange with a given reciprocal pair. Thus, there are four types that matter for our formulation of the process: Reciprocal pairs (pairs of type $A-B$ and $B-A$), and overdemanded pairs that can be part of a three-way exchange including a reciprocal pair of type $A-B$ and $B-A$, namely $\mathbb{O}_{A-B}$ and $\mathbb{O}_{B-A}$, respectively. We refer to the following enumeration of pair types: The type space is given as $\mathcal{T} = \{1, 2, 3, 4\}$ and $n = 4$ is the number of pair types, so that for any state $s = (s_1, s_2, s_3, s_4)$,* *1. s_1 is the number of $\mathbb{O}_{A-B}$ overdemanded pairs,*

2. s_2 is the number of $A-B$ reciprocal pairs,

3. s_3 is the number of $B-A$ reciprocal pairs, and

4. s_4 is the number of $\mathbb{O}_{B-A}$ overdemanded pairs.

Since, whenever an overdemanded pair is to be matched, an underdemanded pair is always part of the exchange,[12] *we formulate exchanges as if $\mathbb{O}_{A-B}$ pairs match with themselves, or match with $A-B$ pairs, generating a surplus of 2 and 3, respectively (similarly, for the $\mathbb{O}_{B-A}$ pairs). Thus, we integrate the existence of underdemanded pairs into the surplus values of exchanges and suppress them from the states and actions throughout. Thus, for each state s, if $s_1, s_2 > 0$ (s_3, s_4>0), we can match $\mathbb{O}_{A-B}$ and $A-B$ pairs ($\mathbb{O}_{B-A}$ and $B-A$ pairs) to generate a surplus of 3, if $s_2, s_3 > 0$, we can match $A-B$ and $B-A$ pairs to generate a surplus of 2, if $s_1 > 0$ ($s_4 > 0$), we can match $\mathbb{O}_{A-B}$ ($\mathbb{O}_{B-A}$) pairs with themselves to generate a surplus of 2. Therefore, the surplus a_{ij} is defined as*

$$a_{ij} := \begin{cases} 3 & \textit{if } \{i,j\} = \{1,2\} \textit{ or } \{i,j\} = \{3,4\} \\ 2 & \textit{if } \{i,j\} = \{2,3\} \textit{ or } i = j = 1 \textit{ or } i = j = 4 \end{cases} .$$

Thus, our goal is to maximize the discounted sum of the number of pairs matched.[13]

The total arrival rate of agents is denoted by $\lambda := \sum_{i \in \mathcal{T}} \lambda_i$. Note that, p_i, the probability that an arriving agent is of type i, is equal to $\frac{\lambda_i}{\lambda}$. By standard arguments, the expected discounting until an agent arrives is $\mathbb{E}[e^{-\rho\tau}] = \frac{\lambda}{\lambda+\rho}$, where τ denotes the random variable of the first arrival. We also denote this quantity by $\delta := \frac{\lambda}{\lambda+\rho}$, whenever convenient.

The MDP starting from state s works as follows: A discount of δ is incurred for the time passed until an agent arrives. When an agent of type i arrives, the temporary state becomes $s + e_i$. Then, the central mechanism decides which subset of available exchanges to conduct, where the *surplus* of each exchange is as described above. After conducting a (possibly empty) subset of (possibly empty) available exchanges, the state transitions to $s' \leq s + e_i$, at which the waiting for another agent starts and the process continues. Here $e_i \in \{0,1\}^n$ is the unit vector of agent type i with 1 in component i and 0 in other components.

For two states s, s' such that $s' \leq s$, we write $s \to s'$ if there is a feasible finite exchange sequence $(i_k j_k)$ such that beginning from s, conducting each exchange in order and removing the matched agents transitions the state to s' and the sequence involves at most q exchanges of each type.[14]

[12] Note that there are infinitely many underdemanded pairs in the pool, and an overdemanded pair is always necessary to match them.

[13] The maximization of the discounted expected surplus is equivalent to the cost minimization problem here, where there is unit time cost for waiting in the pool. We skip the derivation of this equivalence, which is already shown (see Ünver, 2010).

[14] Here, q is an arbitrarily large integer quota denoting a finite capacity for each exchange type ij to be conducted per period. It is only utilized for technical convenience to prevent an unbounded number of exchanges being feasible as the

Let $E(s,s')$ be the set of exchange sequences that transitions state from s to s'. We define the **matching value of transition** from s to s' as

$$\mathcal{M}(s,s') := \max_{(i_k j_k)\in E(s,s')} \sum_k a_{i_k j_k}.$$

We set $s \to s$ and $\mathcal{M}(s,s) := 0$ for any $s \in \mathcal{S}$.[15]

We denote the set of real-valued functions on the state space by $\mathcal{V} := \{f | f : \mathcal{S} \to \mathbb{R}\}$. Let $v^*(s)$ denote the maximized expected discounted surplus of the process that starts at state s. Then, we have the following recursive equation:

$$v^*(s) = \delta\left(\sum_i p_i \max_{s' : s+e_i \to s'} \left(\mathcal{M}(s+e_i, s') + v^*(s')\right)\right). \tag{1}$$

This formulation is compact and easy to understand, but, its structure is analytically difficult. To achieve a more tractable formulation, we next define the *event operators* with which we can construct the recursive equation given above.

3.2 Matching and other event operators

Our approach in this work is based on the concept of *event-based dynamic programming (dp)*, which is introduced by Koole (1998) for queueing models and widely studied in that literature.[16] We adopt *event-based dp* and generalize it to our context.

This methodology consists of (i) decomposing the value function (Equation (1) above) in smaller parts, called *(event) operators*, each of which captures an event in the dynamic exchange, and (ii) analyzing the desired properties of the value function through these *operators* such that we can unravel the structure of the optimal mechanism. Here, we establish Part (i), and we return to Part (ii) in Section 5.

An *(event) operator* is a mapping $T : \mathcal{V}^m \to \mathcal{V}$ for any positive integer m. Thus, Equation (1) can be written as a fixed point relation, $T^* v^* = v^*$, where the **optimality operator** $T^* : \mathcal{V} \to \mathcal{V}$ is

state goes to infinity.

[15] We demonstrate transitions within the kidney exchange setting via an example: The transition of $(1,1,0,0) \to (0,0,0,0)$ is obtained by matching $\mathbb{O}_{A-B}$ and $A-B$ pairs, and $(1,1,1,1) \to (1,0,0,1)$ is obtained by matching $A-B$ and $B-A$ pairs. Similarly, $(3,2,1,0) \to (0,0,0,0)$ is obtained by four exchanges: first, matching an $A-B$ with a $B-A$ pair, second, an $A-B$ with an $\mathbb{O}_{A-B}$ and third and fourth, two $\mathbb{O}_{A-B}$'s with themselves. We denote the situation where state s' cannot be reached from state s through exchanges with available pairs, by $s \not\to s'$. For example, $(1,1,1,0) \not\to (1,0,1,0)$ since $A-B$ pairs cannot match with themselves.

For any $s \to s'$, regardless of which exchange sequence is chosen to realize the transition, we have

$$\mathcal{M}(s,s') = 2(s_1 + s_4 - s'_1 - s'_4) + (s_2 + s_3 - s'_2 - s'_3).$$

[16] This approach completely differs from the one used in Ünver (2010) and our analysis is substantially different from Koole's as our *matching operators* do not have an antecedent in his work.

defined for any $f \in \mathcal{V}$ and $s \in \mathcal{S}$ by

$$(T^* f)(s) := \delta\Bigg(\sum_i p_i \max_{s' : s+e_i \to s'} \big(\mathcal{M}(s+e_i, s') + f(s')\big)\Bigg). \tag{2}$$

Moreover, this operator can be written as a composition of smaller operators, reflecting events in dynamic exchange. We start with the characteristic operators of the current context, the *matching operators*.

Definition 3. Matching operators are defined as follows:

- For any two compatible types i and j with $i \neq j$ the **matching operator of $i - j$** $T_{ij} : \mathcal{V} \to \mathcal{V}$ is defined for any $f \in \mathcal{V}$ and $s \in \mathcal{S}$ by

$$(T_{ij} f)(s) := \begin{cases} \max\{f(s), f(s - e_i - e_j) + a_{ij}\} & \text{if } s - e_i - e_j \geq 0 \\ f(s) & \text{otherwise,} \end{cases}$$

- For any type $i \in \{1, n\}$, the **matching operator of i** $T_i : \mathcal{V} \to \mathcal{V}$ is defined for any $f \in \mathcal{V}$ and $s \in \mathcal{S}$ by

$$(T_i f)(s) := \begin{cases} \max\{f(s), f(s - e_i) + a_i\} & \text{if } s - e_i \geq 0 \\ f(s) & \text{otherwise.} \end{cases}.$$

We call $\pm(e_i + e_j)$ and $\pm e_i$ as the **matching vectors** of the operators T_{ij} and T_i, respectively.

We also consider a generic version of *matching operators*:

- For each $w \in \mathbb{Z}_+^n \cup \mathbb{Z}_-^n \setminus \{0\}$, and a surplus $a_w \in \mathbb{R}_+$, the **generalized (matching) operator of w** $T_w : \mathcal{V} \to \mathcal{V}$ is defined for any $f : \mathcal{S} \to \mathbb{R}$ and $s \in \mathcal{S}$ by

$$(T_w f)(s) := \begin{cases} \max\{f(s), f(s - |w|) + a_w\} & \text{if } s - |w| \geq 0 \\ f(s) & \text{otherwise} \end{cases}.$$

Note that for $w = \pm(e_i + e_j)$, $T_{ij} = T_w$ & $a_w = a_{ij}$, and for $w = \pm e_i$, $T_i = T_w$ & $a_w = a_i$. For other w, this operator is more abstract. The generalized operator is utilized in Section 5.1 to motivate our methodology and prove our results in Section 5.2, where it will also be clear why we consider both positive and negative values for w and not fix its sign as a definition.

We next define other operators relevant to the dynamic aspects of the problem:

Definition 4. We define

- the **discount operator** $T_\delta : \mathcal{V} \to \mathcal{V}$ as follows: for any $f \in \mathcal{V}$ and $s \in \mathcal{S}$, $(T_\delta f)(s) := \delta f(s)$,
- for each agent type i, the **arrival operator of i** $T_{\text{arr},i} : \mathcal{V} \to \mathcal{V}$ as follows: for any $f \in \mathcal{V}$ and $s \in \mathcal{S}$, $(T_{\text{arr},i} f)(s) := f(s + e_i)$, and
- the **uniformization operator** $T_p : \mathcal{V}^m \to \mathcal{V}$ as follows: for any $f_1, f_2, \ldots, f_m \in \mathcal{V}$ and $s \in \mathcal{S}$,

$(T_p(f_1, f_2, \ldots, f_m))(s) := \sum_i p_i f_i(s)$.

Typically, $m \leq n$, the number of agent types, $p \in \mathbb{R}^m$, and T_p is an implicit function of m.

We denote the composition of the operators T and T' by $T \circ T'$. Also, for the k consecutive composition of an operator T, we use the notation T^k.

Observation 1. *Matching operators T_w commute under composition, so that for any pair of matching vectors w, w', $T_w \circ T_{w'} = T_{w'} \circ T_w$.*

This follows simply by observing that both expressions evaluated by any $f \in \mathcal{V}$ at any $s \in \mathcal{S}$ are equal to $\max\{f(s), f(s-|w|)+a_w, f(s-|w'|)+a_{w'}, f(s-|w|-|w'|)+a_w+a_{w'}\}$, where the terms of the form $f(s-|w|)$ or $f(s-|w'|)$ disappear if $s-|w| \ngeq 0$ or $s-|w'| \ngeq 0$, respectively.

This observation allows us to conclude the following:

Lemma 1. *For any state $s \in \mathcal{S}$ and $f \in \mathcal{V}$,*

$$(T_1^q \circ T_{12}^q \circ \ldots \circ T_{(n-1)n}^q \circ T_n^q f)(s) = \max_{s' : s \to s'} \left(\mathcal{M}(s,s') + f(s')\right).$$

Proof of Lemma 1. The left-hand side consists of a $\max\{\cdot\}$ expression with terms of the form $f(s+d)+a_d$ where d iterates over all possible sums of *matching vectors* (namely, the vectors $-e_1, -e_1-e_2, \ldots, -e_{n-1}-e_n, -e_n$) with each vector used at most q times, and a_d is the sum of values a_w for each *matching vector* w used. Letting $s+d=s'$, and observing $a_d = \mathcal{M}(s, s+d)$ for some a_d (i.e., for the sequence of exchanges achieving the maximum surplus), we see that the left-hand side iterates over all s' such that $s \to s'$. But this is the domain of the $\max\{\cdot\}$ in the right-hand side. So we have the equality. □

We define the **composite matching operator** as

$$T_M := T_1 \circ T_{12} \circ \ldots \circ T_{(n-1)n} \circ T_n.$$

Finally, we obtain the following:

Observation 2. *The optimal value function v^* (Equation (1) in Section 3.1) is equivalent to for any $s \in \mathcal{S}$,*

$$v^*(s) = \left(T_\delta \circ T_p\left(T_M^q \circ T_{\text{arr},1} v^*, T_M^q \circ T_{\text{arr},2} v^*, \ldots, T_M^q \circ T_{\text{arr},n} v^*\right)\right)(s). \tag{3}$$

Thus, the optimality operator T^* (Equation (2)) is a composition of matching operators, arrival operators, discount operator, and uniformization operator.

We return to *event-based dp* and *matching operators* in Section 5 to analyze the desired properties of the value function to solve the optimal control problem. But first, we argue that *concavity* and certain second-order properties of v^* are crucial for this solution.

4 Monotonicity of Optimal Control

The structure of the optimal mechanism is, it turns out, closely related to the second-order properties of the value function v^*. To see this, consider the simpler case of our model and focus on the *matching operator* of the form $(T_i f)(x) = \max\{f(x), f(x-e_i)+1\}$. Then, the optimal mechanism chooses the first action if and only if $f(x) - f(x-e_i) > 1$. Now, the intuition is simply the following: The left-hand side of this inequality is the marginal value of an agent of type i, that is, the opportunity cost of matching, and the right-hand side is the benefit of matching the agent of type i. Now, if there are diminishing returns to pooling type i (*concavity*, see Definition 5), then the left-hand side becomes lower as the pool gets larger. Then, there is a critical level, the *threshold*, such that optimality implies pooling the agent below the level and matching above it. Thus, the optimal mechanism is a *threshold* mechanism (see Koole, 1998 for this argument).

As we analyze later, the structure of the optimal mechanism for the current problem is also determined by the second-order properties of the value function (see Section 6). Here, we focus on these second-order properties and their economic intuitions.

Definition 5. A function $f : \mathbb{N}^n \to \mathbb{R}$ satisfies **concavity in component *i*** if for each $x \in \mathbb{N}^n$,

$$f(x) + f(x+2e_i) \leq 2f(x+e_i). \tag{4}$$

A function is **componentwise concave** if it is *concave in each component* $i \in \{1, 2, \ldots, n\}$.

For each agent type i, Condition (4) is equivalent to

$$f(x+2e_i) - f(x+e_i) \leq f(x+e_i) - f(x),$$

which implies *diminishing marginal return on* i. As we explained above, this connects *componentwise concavity* to a *threshold mechanism* being optimal.

As we argue later in detail, in addition to *concavity*, *substitutability* and *complementarity* properties of agent types are essential in unraveling the structure of the optimal control (see Section 6), and also for the consistency of our theoretical framework (see Section 5). The intuition is clear: For example, in kidney exchange, certain types are *complements* (e.g. $X-Y$ and $Y-X$ pair types) and certain others are *substitutes* (e.g. $X-Y$ and $\mathbb{O}_{Y-X}$ type pairs), and these properties clearly affect which pairs to keep and which others to match at a given state of the MDP.

Definition 6. A function $f : \mathbb{N}^n \to \mathbb{R}$ is ***ij*-submodular** if for each $x \in \mathbb{N}^n$,

$$f(x) + f(x + e_i + e_j) \leq f(x + e_i) + f(x + e_j).$$

A function $f : \mathbb{N}^n \to \mathbb{R}$ is ***ij*-supermodular** if for each $x \in \mathbb{N}^n$,

$$f(x + e_i) + f(x + e_j) \leq f(x) + f(x + e_i + e_j).$$

The connection of these two properties to *substitutability* and *complementarity* is clear: The condition for *ij-submodularity* is equivalent to

$$f(x + e_i + e_j) - f(x + e_j) \leq f(x + e_i) - f(x),$$

which states that marginal value of an additional agent of type i decreases with each additional agent of type j. Thus, *ij-submodularity* states that i and j agent types are **substitutes**. Similarly, *ij-supermodularity* states that i and j agent types are **complements**.

Definition 7. A function $f : \mathbb{N}^n \to \mathbb{R}$ is ***ij*-superconcave** if for each $x \in \mathbb{N}^n$,

$$f(x + e_i) + f(x + e_i + e_j) \geq f(x + e_j) + f(x + 2e_i).$$

A function is **superconcave** if for each i, j, it is ij-superconcave.

Superconcavity states that the marginal value of i is lower when we have an additional i instead of an additional j. Economically, this states that no other type is a closer *substitute* to i than i itself.

5 How to Solve the Optimal Control Problem

We have explained so far that the second-order properties of the value function are crucial in understanding the structure of the optimal policy. In particular, we argued that *concavity* of the value function implies monotonicity of the optimal control, i.e., existence of a *threshold-type* dynamic mechanism. We now present our methodology for proving the second-order properties of the value function. We first state the existence of the optimal value function which follows directly from the fundamental theorem of MDPs:[17]

Theorem 1. *The function $v^* : \mathcal{S} \to \mathbb{R}$ defined by the fixed point relation*

$$v^*(s) = \delta\left(\sum_i p_i \max_{s' : s+e_i \to s'} \left(\mathcal{M}(s + e_i, s') + v^*(s')\right)\right) \tag{5}$$

[17] See Chapter 6 in Puterman (1994), for this result and the related discussion.

exists and is unique. Moreover, for any $v_0 : \mathcal{S} \to \mathbb{R}$*, the sequence* (v_k) *defined by the relation* $v_k = T^* v_{k-1}$ *converges uniformly to* v^**, where the operator* T^* *is defined by the relation*

$$(T^* v)(s) := \delta\left(\sum_i p_i \max_{s' : s+e_i \to s'} \left(\mathcal{M}(s+e_i, s') + v(s')\right)\right). \tag{6}$$

The second part is particularly important for the methodology of *event-based dp*. Notice that, the choice of v_0 in the sequence (v_k) is arbitrary, thus, we can start the sequence with a function v_0 satisfying a given property, as long as the property is not logically inconsistent with itself.[18] This observation allows us to focus on the operator T^* itself. We first define the following notion:

Definition 8. We say an operator $T : \mathcal{V} \to \mathcal{V}$ **propagates** a property P, if for each function $v \in \mathcal{V}$ satisfying P, $Tv \in \mathcal{V}$ also satisfies P.

We observe the following: if a property P is *propagated* by the operator T^*, and P is preserved under *uniform convergence*,[19] then, by the second part of Theorem 1, the fixed point of the operator T^* also satisfies P. Thus, our proof strategy is to show that the operator T^* *propagates* the desired second-order properties of the value function.

The difficulty with this approach is that, as it could be that properties P and P' are not, separately and alone, *propagated* by some operator T, their intersection $P \cap P'$ is *propagated* by T. Indeed, it is often the case that, desired properties are too weak to be *propagated* by T, and we look for additional properties so that the intersection of these properties are *propagated* by the operator as a stronger property. We next observe (in Section 5.1) that the current model of dynamic exchange suffers from this difficulty: The second-order properties defined in Section 4 are not *propagated* separately. Furthermore, if we consider all these properties together (instead of considering whether they are *propagated* one-by-one), they are not *propagated* either. This motivates the definition of *multimodularity*, which, it turns out, is just strong enough to be *propagated* by the operator T^*.

5.1 Motivation for multimodularity in exchange context

First, we utilize the *generalized matching operator* defined in Definition 3. Additionally, we formulate a generic version for the second-order properties given in Section 4.

Definition 9. Let $u, v \in \mathbb{Z}^n$. A function $f : \mathbb{N}^n \to \mathbb{R}$ satisfies **property** $\boldsymbol{P(u, v)}$ if for each $x \in \mathbb{N}^n$,

[18] In particular, since all of the properties we define have weak inequalities, the constant zero function $v_0(s) := 0$ satisfies every property we present in this paper.

[19] The definition of *uniform convergence* is standard: for any sequence (x_k) that satisfy the property and converges uniformly to x, the limit x also satisfies the property.

such that $x+u, x+v, x+u+v \in \mathbb{N}^n$, we have

$$f(x+u)+f(x+v) \leq f(x)+f(x+u+v). \tag{7}$$

Note that $P(u,v)$ states that the vectors u and v are *complements* (that is, there is *supermodularity* between u and v). Also, note that $P(u,v)$ is defined such that u and v are not restricted to be positive vectors and this provides a general formulation for the properties defined in Section 4: For example, taking $u=e_i$ and $v=-e_i$, $P(u,v)$ is equivalent to *componentwise concavity* in component i, or taking $u=e_i$ and $v=-e_j$, it is equivalent to ij-submodularity.

Now, the problem of characterizing the set of second-order properties such that they are *propagated* together (as explained above) can now be defined as exploring the set of these properties in the form of $P(u,v)$. As it will be clear in Section 5.4, the crux of the *propagation* problem is the *propagation* by the *matching operator* (defined in Definition 3 in Section 3.2). Thus, a necessary condition for the solution to this problem of characterizing the set of properties is the condition on how $P(u,v)$ is *propagated* by T_w. Now, by definition of the *matching operator* T_w, showing that $P(u,v)$ is *propagated* by T_w requires to show the following for any $x \in \mathbb{N}^n$:

$$\begin{aligned}
&\max\{f(x+u), f(x+u+w)+a_w\}+\\
&\max\{f(x+v), f(x+v+w)+a_w\}\\
\leq\ &\max\{f(x), f(x+w)+a_w\}+\\
&\max\{f(x+u+v), f(x+u+v+w)+a_w\}.
\end{aligned}$$

Here we introduce two definitions that will be helpful throughout the paper. The left-hand side of the inequality equals to one of the four expressions: (i) $f(x+u)+f(x+v)$, (ii) $f(x+u+w)+a_w+f(x+v+w)+a_w$, (iii) $f(x+u)+f(x+v+w)+a_w$, and (iv) $f(x+u+w)+a_w+f(x+v)$. We call the first two cases the *symmetric* cases, as these cases occur when both $\max\{\cdot\}$ expressions are equal to their first (respectively second) argument. We call the next two possible expressions as the *first case of asymmetry* and the *second case of asymmetry* respectively. In the first case of asymmetry the first $\max\{\cdot\}$ expression equals to its first argument and second expression equals to its second argument, and in the second case of asymmetry it is vice versa.

Now, we can further observe that, in order to show the inequality, we can arbitrarily replace $\max\{\cdot\}$ expressions on the right-hand side of the inequality by one of their arguments. To see this, note that showing that $x<a$ or $x<b$ suffices for showing $x<\max\{a,b\}$. Thus, depending on the cases on the left-hand side, we can pick the arguments on the right-hand side arbitrarily.

Using this observation, our inequality follows easily in the symmetric cases: We can pick the arguments on the right-hand side so that they match the symmetry on the left-hand side. Then, the inequality will either become $f(x+u)+f(x+v) \leq f(x)+f(x+u+v)$ or $f(x+u+w)+f(x+v+w) \leq f(x+w)+f(x+u+v+w)$, both of which follows from the fact that f satisfies $P(u,v)$. Thus, what remains to show is the asymmetric cases.

Now suppose the left-hand side of the equation has first case of asymmetry, so that it is equal to $f(x+u)+f(x+v+w)+a_w$. Then, it would suffice to show that this expression is less than $f(x)+f(x+u+v+w)+a_w$, since the right-hand side of the inequality is greater than this. This is equivalent to showing $P(u, v+w)$. The same can be observed in the second case of asymmetry as well. Thus, a sufficient condition for the *propagation* of a set of properties is the following: if a set of properties in the form of $P(u,v)$ is *propagated*, then $P(u, v+w)$ should also be in this set. We refer to the set of properties satisfying this condition as **closed**.

We illustrate the concept of *closedness* in *propagation* in a two-dimensional setting. Suppose $f : \mathbb{Z}^2 \to \mathbb{R}$ is a *concave* function in each of the two components. Also, T_1 and T_2 are two operators, defined by the equations $(T_1 f)(x) = \max\{f(x), f(x-e_1)+a_1\}$ and $(T_2 f)(x) = \max\{f(x), f(x-e_2)+a_2\}$, respectively. We show that, in this setting, even though *concavity* does not *propagate* alone, it *propagates* together with *superconcavity*. Thus the set of properties $\{$*concavity, superconcavity*$\}$ is closed with respect to the set of operators $\{T_1, T_2\}$ (for the proof, see the Online Appendix).

Unfortunately, these results do not generalize to multiple dimensions, for the simple reason that, when we have ij*-superconcavity*, and we have an operator T_k defined similar to T_1 above for some $k \neq i, j$, the necessary second-order property we obtain does not follow from *superconcavity*.[20] For this reason, the set of properties we have defined so far are also not *closed*.

To make this point precise in the kidney exchange domain, we note that, to show that *concavity in component* $i = 2$ is *propagated* by the operator T_{-e_1} (which is the *matching operator* deciding whether an overdemanded pair should be matched in a two-way exchange), we also need to show that the function is ij*-superconcave* in components $i = 2$ and $j = 1$, thus, we need to show $P(e_2, e_1 - e_2)$. But, to show that *superconcavity* is propagated by the *matching operator* T_{-e_4}, the required property is of the form $P(e_2 - e_4, e_1 - e_2)$, which is not implied by our existing properties. Moreover, to show *submodularity* and *supermodularity* properties, we need additional properties of the form $P(u,v)$, which are not implied by the set of our properties defined so far. Thus, this set of four second-order properties $\{$*concavity, superconcavity, submodularity, supermodularity*$\}$ is not *closed*: they cannot be *propagated* by the *matching operator*.

The puzzle becomes the following: what is the set of restrictions on u and v such that the set of properties of the form $P(u,v)$ is closed, that is, when is it the case that whenever we have a

[20] Ease of two dimensions here comes from the fact that, when we have an operator for a component k and a property for some components i, j, k must coincide with one of i or j, which makes the proofs trivial.

matching operator T_w, $P(u, v+w)$ is also contained in this set? It turns out that the concept of $\mathcal{D}$-*multimodularity*, introduced by Hajek (1985) (see also Altman et al., 2003), helps us develop a methodology to solve the *closedness* puzzle and captures this general form.

5.2 $\mathcal{D}$-Multimodularity

We have seen that all of the second-order properties we are interested in is of the form $P(u, v)$, and their difference lies in the domain from which vectors u and v are chosen. The approach of $\mathcal{D}$–*multimodularity* is based on the idea of constructing a set of vectors $\mathcal{D}$, such that *for any* two vectors $u, v \in \mathcal{D}$, our set of properties includes $P(u, v)$.

By choosing both vectors u, v arbitrarily from the same domain, the obtained set of properties is strong enough to satisfy *closedness* property we discussed in Section 5.1 (see Lemma 2 below).

Definition 10. A set of integer-valued vectors $\mathcal{D} \subseteq \mathbb{Z}^n$ is called a **multimodular basis** of $\mathbb{Z}^n$ if

i. $\sum_{v\in\mathcal{D}} v = 0$, and
ii. $\mathrm{span}(\mathcal{D}) = \mathbb{Z}^n$.[21]

We are now ready to define $\mathcal{D}$-*multimodularity*.

Definition 11. (Hajek, 1985) A function $f : \mathbb{N}^n \to \mathbb{R}$ is **$\mathcal{D}$-multimodular** if for each $x \in \mathbb{N}^n$, $u, v \in \mathcal{D}$, such that $x+u, x+v, x+u+v \in \mathbb{N}^n$, we have

$$f(x+u) + f(x+v) \leq f(x) + f(x+u+v).^{22} \tag{8}$$

Note that $\mathcal{D}$-*multimodularity* is equivalent to property $P(u, v)$ being satisfied for each $u, v \in \mathcal{D}$. The usefulness of this definition lies in the following result:

Lemma 2. *A function $f : \mathbb{N}^n \to \mathbb{R}$ is $\mathcal{D}$-multimodular if and only if for any two disjoint subsets $U, V \subset \mathcal{D}$ we have*

$$f\left(x + \textstyle\sum_{u\in U} u\right) + f\left(x + \textstyle\sum_{v\in V} v\right) \leq f(x) + f\left(x + \textstyle\sum_{u\in U} u + \textstyle\sum_{v\in V} v\right).$$

See Appendix A for the proof of this result.

By Lemma 2, instead of taking two distinct vectors $u, v \in \mathcal{D}$ and stating that they are *complements*, we can alternatively state that any two sums of distinct vectors in disjoint subsets of $\mathcal{D}$ are *complements*. The proof of the statement is intuitive: If u and v are *complements*, and u

[21] Note that in our context $\mathrm{span}(\mathcal{D}) = \{z \in \mathbb{Z}^n : z = \sum_{v\in\mathcal{D}} \alpha_v v \;\; \exists\, (\alpha_v)_{v\in\mathcal{D}} \in \mathbb{Z}^{|\mathcal{D}|}\}$.

[22] In Hajek (1985), as well as in other prominent works on *multimodularity*, this inequality is of reversed form. The reason we define it differently here is that we focus on a *reward-maximization* problem, as opposed to a *cost-minimization* problem that is frequent in this literature. Thus, we search for *concavity*-related properties, as opposed to *convexity*-related properties.

and w are *complements*, then u and $v+w$ should also be *complements*, since the marginal value of u is increasing in both v and w. Thus, we can prove the statement by induction on the size of the sets U, V. If the statement holds for sets of size at most k, by applying the definition of $\mathcal{D}$*-multimodularity* again we can see that the statement holds for sets of size $k+1$ as well.

Lemma 2 explains the reason why the concept of $\mathcal{D}$*-multimodularity* solves the *closedness* problem discussed in Section 5.1. Notice that as a special case of Lemma 2, we can take three vectors $u, v, w \in \mathcal{D}$, and conclude that a $\mathcal{D}$*-multimodular* function f also satisfies property $P(u, v+w)$. Thus, the *closedness* property defined in Section 5.1 is satisfied for each *matching operator* T_w such that $w \in \mathcal{D}$. Thus, as long as the *matching operators* are of the form T_w for some $w \in \mathcal{D}$, sufficient propagation results are obtained. Although there are some minor nuances, this is essentially the main idea behind the proof of the core result (Proposition 2) we utilize for our main theorem (Theorem 3).

Another motivation revealed by Lemma 2 is the reason behind property (i) in Definition 10. Notice that, since $\sum_{u\in\mathcal{D}} u = 0$, we have that for $u \in \mathcal{D}$, the sum of elements excluding u is $-u$. Thus, u and $-u$ are always sums of distinct vectors in disjoint subsets of $\mathcal{D}$, and by Lemma 2, property $P(u, -u)$ is satisfied. For example, when $u = \pm e_i$, $P(u, -u)$ is componentwise concave in i. This is also true for all vectors u that can be written as sums of distinct vectors in disjoint subsets of $\mathcal{D}$. This observation will become useful in later sections when we discuss the implications of $\mathcal{D}$*-multimodularity*.

5.3 Multimodular basis of matching

We construct a *multimodular basis*, capturing both the decisions in dynamic exchange and the necessary second-order properties.

Definition 12. The *multimodular basis* $\mathcal{D}^M = \{e_1, -e_2-e_1, e_3+e_2, \ldots, (-1)^{n-1}(e_n+e_{n-1}), (-1)^n e_n\}$ of $\mathbb{Z}^n$ is called the **matching basis**.

It is straightforward to check the *matching basis* is indeed a *multimodular basis*. Also, this basis contains *matching vectors* relevant for the exchange context: Our *matching operators* are of the form $(T_w f)(x) = \max\{f(x), f(x-|w|)+a_w\}$ (see Definition 3), where $w \in \{\mp e_1, \mp(e_2+e_1), \ldots, \mp(e_n+e_{n-1}), \mp e_n\}$. We separate these operators into two groups: We write $(T_w^+ f)(x) = \max\{f(x), f(x+w)+a_w\}$ for $w \in \mathcal{D}^M \cap \mathbb{Z}_-^n$ and $(T_w^- f)(x) = \max\{f(x), f(x-w)+a_w\}$ for $w \in \mathcal{D}^M \cap \mathbb{Z}_+^n$. Thus, all of our operators are of the form $T_w^\pm$ for some $w \in \mathcal{D}^M$, where the sign of $T^\pm$ depends on whether $w < 0$ or $w > 0$. *Matching vectors* being members of the *matching basis* will be crucial in our main theorem below.

In Figure 1, we present the *matching vectors* and *matching basis* reflecting complementarity and substitutability relations in the context of kidney exchange: There are two sides of the market, Sides A and B, for A and B type patients, respectively. Each side contains a reciprocal pair and an overdemanded pair such that each reciprocal pair can be matched with pairs from the other side, with the cross-type pair via a two-way exchange and with the overdemanded pair via a three-way exchange (including an underdemanded pair). This representation provides a clear view of complementarity and substitutability structure: pairs on the same side of the market are substitutes of each other, and each reciprocal pair is a complement with the cross-type pair and also with the overdemanded pair from the other side. This figure also depicts the insight for how the *matching basis* is constructed.

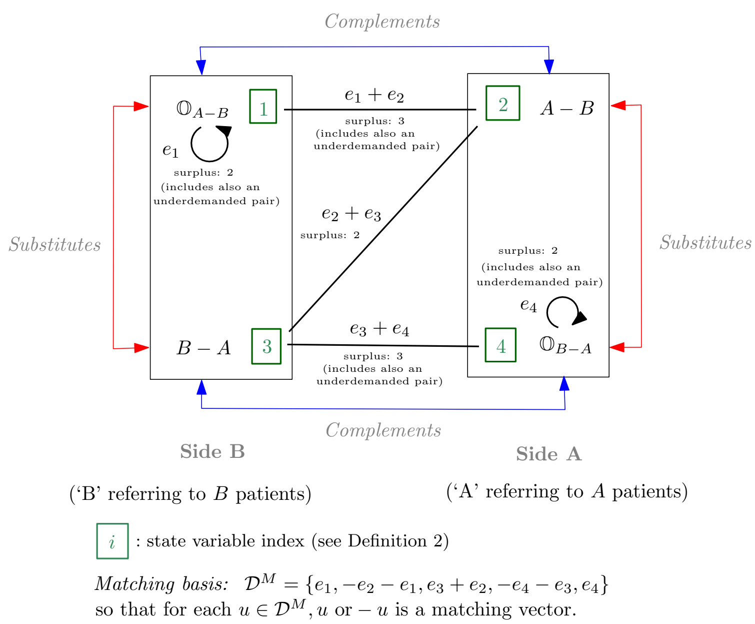

Figure 1: Matching vectors and multimodular basis in kidney exchange

5.4 Main result

As we have previously explained in Section 3.2, *optimality operator* T^* itself can be written as a composition of *matching operators* and other operators capturing *discounting*, *arrival* and *linear combination* (see Observation 2 in Section 3.2 for this argument).

First, we state the following result on *propagation*:

Proposition 1. *Let $f, f_1, f_2, \ldots, f_m$ be $\mathcal{D}^M$-multimodular functions from $\mathbb{N}^n$ to $\mathbb{R}$ with $m \leq n$. Then, $T_\delta f$, for any agent type i, $T_{\mathrm{arr},i} f$, and $T_p(f_1, f_2, \ldots, f_m)$ are also $\mathcal{D}^M$-multimodular.*

Proof of Proposition 1. $\mathcal{D}^M$*-multimodularity* of the *discount operator* and *uniformization operators* follows from multiplying each inequality with the discount factor (respectively the weights p_i) and summing the resulting inequalities in the latter case. $\mathcal{D}^M$*-multimodularity* of the *arrival operator* follows from letting $s' = s + e_i$. □

Since by Proposition 1, *arrival*, *uniformization* and *discount operators propagate* $\mathcal{D}^M$*-multimodularity*, to prove that v^* satisfies $\mathcal{D}^M$*-multimodularity*, by Observation 2 (Section 3.2), it is sufficient to prove that each *matching operator propagates* $\mathcal{D}^M$*-multimodularity*.

The result below states that $\mathcal{D}^M$*-multimodularity* is *propagated* by *generalized matching operator* T_w for each $w \in \mathcal{D}^M$ defined in Definition 3:

Proposition 2. *Let* $f : \mathbb{N}^n \to \mathbb{R}$ *be a* $\mathcal{D}^M$*-multimodular function. Then, for each* $w \in \mathcal{D}^M$*, the function* $T_w f$ *is also* $\mathcal{D}^M$*-multimodular.*

We only present the proof for $w < 0$ here. The proof for the case $w > 0$ as well as the rigorous treatment of the boundary conditions (i.e., if one or more of the vectors that go into $f(\cdot)$ as argument are not non-negative) are in Appendix B.

Proof of Proposition 2 for $w < 0$. Let $u, v \in \mathcal{D}^M$. Fix $x \in \mathbb{N}^n$. Since $w < 0$, $(T_w f)(x) = \max\{f(x), f(x+w) + a_w\}$. In what follows, assume that all arguments of f are non-negative, and thus, $f(\cdot)$ is well-defined. First, we assume $w \neq u, v$. We need to show the following:

$$\begin{aligned} &\max\{f(x+u), f(x+u+w) + a_w\} + \\ &\max\{f(x+v), f(x+v+w) + a_w\} \\ \leq\ &\max\{f(x), f(x+w) + a_w\} + \\ &\max\{f(x+u+v), f(x+u+v+w) + a_w\}. \end{aligned}$$

If the maximizing actions are the same on the left-hand side,[23] the statement follows immediately from $\mathcal{D}^M$*-multimodularity* of f. Suppose they are different. Then, the left-hand side is equal to either $f(x+u) + f(x+v+w) + a_w$ or to $f(x+v) + f(x+u+w) + a_w$. Since u, v, w are distinct and f is $\mathcal{D}^M$*-multimodular*, by Lemma 2, both of these expressions are smaller than $f(x) + f(x+u+v+w) + a_w$, which is weakly smaller than the right-hand side of the inequality.[24]

Now, suppose $u = w$ and $a = a_w = a_u$. Then, we need to show

$$\begin{aligned} &\max\{f(x+u), f(x+u+u) + a\} + \\ &\max\{f(x+v), f(x+v+u) + a\} \\ \leq\ &\max\{f(x), f(x+u) + a\} + \\ &\max\{f(x+u+v), f(x+u+v+u) + a\}. \end{aligned}$$

[23] By maximizing action we mean which of its arguments equals to the value of $\max\{\cdot\}$.

[24] $\max\{f(x), f(x+w) + a_w\} + \max\{f(x+u+v), f(x+u+v+w) + a_w\} \geq f(x) + f(x+u+v+w) + a_w$.

Symmetric actions again follow easily. By $\mathcal{D}^M$- *multimodularity*, we have

$$f(x) + f(x+2u) \leq 2f(x+u) \text{ (concavity), and}$$
$$f(x+u) + f(x+v) \leq f(x) + f(x+u+v).$$

Summing these inequalities, and rearranging it, we obtain

$$f(x+v) - f(x+u+v) \leq f(x+u) - f(x+2u).$$

Thus, if $f(x+u) - f(x+2u) \leq a$, then we have $f(x+v) - f(x+u+v) < a$ as well. Thus, if $\max\{f(x+u), f(x+u+u)+a\} = f(x+2u)+a$, then $\max\{f(x+v), f(x+v+u)+a\} = f(x+u+v)+a$. Thus, the only case we have to check for different actions is the case with the left-hand side being equal to $f(x+u)+f(x+u+v)+a$. But, choosing the second and first arguments in the $\max\{\cdot\}$'s respectively, this expression is less than the right-hand side. The case for $v = w$ is symmetric. □

The next result is the core result for the rest of the paper.

Theorem 2. *The optimal value function v^* is $\mathcal{D}^M$-multimodular.*

Proof of Theorem 2. The result follows immediately from Observation 2 in Section 3.2, and Propositions 1 and 2 above. □

Using this result, we prove the crucial second-order properties of the value function of the dynamically optimal mechanism.

Theorem 3. *The optimal value function v^* is componentwise concave, superconcave, and for each pair of types i, j where j and i are both even or both odd, i.e., have the same parity, ij-submodular and for each pair of types i, j where i and j have different parities, ij-supermodular.*

See Appendix C for the proof of this result.

Theorem 3 states that the optimal value function is concave in each component. Thus, the marginal optimal value of stocking an additional agent of any given type is a monotonically decreasing function of the stock of that type. Moreover, the optimal value function is supermodular for types that can be utilized in exchange together, and thus these types are indeed complements. In the kidney exchange setting, complementary pair types are $\mathbb{O}_{A-B}$ & $A-B$, $A-B$ & $B-A$, $B-A$ & $\mathbb{O}_{B-A}$, and $\mathbb{O}_{A-B}$ & $\mathbb{O}_{B-A}$. Additionally, it is submodular for types that can be utilized instead of each other in an exchange, and thus these types are substitutes. The substitute types are $\mathbb{O}_{A-B}$ & $B-A$ and $A-B$ & $\mathbb{O}_{B-A}$. The optimal value function is also superconcave, which implies that two agents of the same type are better substitutes than two agents, one of a different type and one from the same type.

6 The Structure of the Optimal Mechanism

We have argued that the monotonicity properties of the value function are crucial for the structure of the optimal mechanism and explained the intuition behind it (see Section 4). In addition, we have shown that the value function satisfies a certain set of these properties (Theorem 3). Next, we characterize the optimal mechanism. We first define the following.

Definition 13. For a value function v^*, a state $s \in \mathcal{S}$ is **reachable** if for each s' with $s \rightarrow s'$ we have

$$v^*(s) > \mathcal{M}(s, s') + v^*(s').$$

We denote the set of *reachable* states by $\mathcal{R} \subseteq \mathcal{S}$. We call a state **unreachable** if it is not reachable.

A *reachable* state is such that there is no sequence of exchanges (including only -some of- the existing pairs) which increases the surplus. An *unreachable* state s, on the other hand, implies the existence of another state s' with $s \rightarrow s'$ such that transitioning from s to s' (by means of a sequence of exchanges) is not only feasible but also increases surplus. The motivation is simple: since *unreachable* states imply transitions with additional surplus, the optimal mechanism never reaches these states.

When the state is s, and the incoming pair is of type i, the state first transitions to $s + e_i$, before the relevant decision is made. Then, the optimal mechanism maximizes $\sum_k \#_k a_{w_k} + v(s + e_i - \sum_k \#_k |w_k|)$ subject to $x + e_i - \sum_k \#_k |w_k| \geq 0$, where $\#_k$ denotes the number of exchanges of type w_k used. Defining $w_k = \pm e_k$ for any $k \in \{1, n+1\}$ and $w_k = \pm(e_{k-1} + e_k)$ for any $k \in \{2, \ldots, n\}$, we have

$$\mathcal{D}^M = \{|w_1|, -|w_2|, |w_3|, ..., (-1)^n |w_{n+1}|\}. \tag{9}$$

Note that for each agent type i odd, e_i is the sum of first i elements in this set, and for each agent type i even, e_i is the sum of last $n + 1 - i$ elements in this set.

Now we define a special set of matching vector indices corresponding to each agent type i:

Definition 14. For each agent type (or state component index) $i \in \tau = \{1, \ldots, n\}$, define an index set of matching vectors in $\mathcal{D}^M$, $\boldsymbol{K_i} \subset \{1, \ldots, n+1\}$, as

$$K_i := \{i, i-2, i-4, \ldots\} \cup \{i+1, i+3, i+5, \ldots\}.$$

Observe that for i odd, we have the following two equalities:

$$\{w_k : k \in K_i\} = \{|w_1|, |w_3|, \dots, |w_i|\} \cup \{-|w_{i+1}|, -|w_{i+3}|, \dots\} \text{ and}$$

$$\mathcal{D}^M = \{\underbrace{|w_1|, -|w_2|, |w_3|, \dots, |w_i|}_{\text{the sum is } e_i}, \overbrace{-|w_{i+1}|, \dots, (-1)^n |w_{n+1}|}^{\text{the sum is } -e_i}\} \tag{10}$$

where (10) follows as explained in Section 5.3.

Moreover, for i even, we have

$$\{w_k : k \in K_i\} = \{-|w_2|, -|w_4|, \dots, -|w_i|\} \cup \{|w_{i+1}|, |w_{i+3}|, \dots\} \text{ and}$$

$$\mathcal{D}^M = \{\underbrace{|w_1|, -|w_2|, |w_3|, \dots, -|w_i|}_{\text{the sum is } -e_i}, \overbrace{|w_{i+1}|, \dots, (-1)^n |w_{n+1}|}^{\text{the sum is } e_i}\}. \tag{11}$$

We refer to the first i vectors in (10) if i is odd, and the last $n - i + 1$ vectors in (11) if i is even, as the *set of matching vectors summing up to* e_i. This structure enables us to express the sets of indices K_i in a more compact way.

Observation 3. *For each i, the set K_i is characterized by the set of all indices ℓ for which ℓ is odd and w_ℓ is included in the set of matching vectors summing up to e_i, or ℓ is even and w_ℓ is not included in the set of matching vectors summing up to e_i.*

We use the set K_i and its complement $K_i^c := \{1, \dots, n+1\} \setminus K_i$ to classify how the marginal value of each exchange of type k changes when an agent of type i arrives (i.e., $\Delta_{|w_k|} v^*(s + e_i) := v^*(s + e_i) - v^*(s + e_i - |w_k|)$ provided that $s + e_i - |w_k| \geq 0$). First note that, by Lemma 2, if $|w_k|$ and e_i are sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$, $P(e_i, |w_k|)$ holds. Then, by definition of K_i, for v^*, for each $k \in K_i$, $P(e_i, -|w_k|)$ holds and, for each index $k \in K_i^c$, $P(e_i, |w_k|)$ holds. Now, by definition of Property P, for v^*, the marginal value decreases whenever $P(e_i, -|w_k|)$ holds, and it increases whenever $P(e_i, |w_k|)$ holds. Thus, when an agent of type i arrives, for each type $k \in K_i$, the marginal value of $|w_k|$ decreases, and for each type K_i^c, the marginal value of $|w_k|$ increases.

Our structural theorems and their proofs use this important conclusion of Lemma 2. Thus, we have the following observation:

Observation 4. *For each agent type i and matching vector index $k \in K_i$, $P(e_i, -|w_k|)$ holds for v^*, thus, the the marginal value of $|w_k|$ decreases; moreover, for each index $\ell \in K_i^c$, $P(e_i, |w_\ell|)$ holds for v^*, thus, the marginal value of $|w_k|$ increases.*

We now state our next theorem regarding the structure of exchanges used in the optimal mechanism:

Theorem 4. *When the state is reachable and the incoming agent is of type i, the optimal mechanism conducts only exchanges of type w_k for $k \in K_i$. Moreover, it conducts each such exchange at most once, so that all decisions can be identified with a set of indices $A(s,i) \subseteq K_i$ for each reachable state s and agent type i such that when the state is s and the incoming agent is of type i, optimal mechanism transitions to $s + e_i - \sum_{k \in A(s,i)} |w_k|$ in exchange for an immediate reward of $\sum_{k \in A(s,i)} a_{w_k}$.*

See Appendix D for the proof of this theorem.

By Theorem 4, the maximum number of exchanges of each type that can be conducted in a period, quota q, in the problem definition is never binding, since regardless of how high the quota is, starting from any reachable state, at the optimal mechanism, each exchange vector is used at most once.

Theorem 4 shows that the only exchanges that can be conducted when a type i agent arrives at a reachable state s are captured by indices in a subset $A(s,i)$ of K_i. This is because, when an agent of type i arrives, K_i is the set of matching vector indices whose marginal value decreases, and K_i^c is the set of indices whose marginal value increases.

We next show that after a type i agent arrives, there is a particular structure in how conducted exchanges change at two consecutively reachable states that only differ in whether there is an additional type j agent or not.

Theorem 5. *Let s be a reachable state such that $s + e_j$ is also reachable for some agent type j. Then, we have for any agent type i,*

$$A(s,i) \cap K_j \subseteq A(s+e_j, i) \subseteq A(s,i) \cup K_j.$$

See Appendix D for the proof of this theorem.

Theorem 5 states that the decision of the optimal mechanism after an type i agent arrives, i.e., the sets $A(s,i)$ and $A(s+e_j,i)$ (recall that they are both subsets of K_i), can only potentially add vectors from K_j and remove vectors from K_j^c when this arrival occurs at state s versus $s + e_j$, respectively. This is because the sets K_j and K_j^c characterize the indices of matching vectors whose marginal values respectively decrease and increase when an additional type j agent exists in the pool. Thus, in the latter state, the optimal mechanism can only additionally conduct exchanges whose marginal value decreases and only stop conducting exchanges whose marginal value increases.

Moreover, we establish that marginal value changes are monotonic.

Theorem 6. *Let s be a reachable state such that $s+e_j$ is not reachable for some agent type j. Then, $s+ke_j$ is not reachable for any positive integer k.*

Proof of Theorem 6. Suppose s is reachable and $s+e_j$ is not. Then, $A(s,j) \neq \emptyset$. Let $\ell \in A(s,j) \subseteq K_j$. Then, the property $P(-|w_\ell|, e_j)$ holds for v^* by Observation 4. Thus, we also have the property $P(-|w_\ell|, (k-1)e_j)$ hold for v^* by the claim in the proof of Lemma 2. Thus, setting $s' = s+e_j$, $u = -|w_\ell|$, $v = (k-1)e_j$ in the definition of property P we obtain,

$$v^*(s+ke_j) + v^*(s+e_j-|w_\ell|) \leq v^*(s+e_j) + v^*(s+ke_j-|w_\ell|).$$

Thus, by rearranging the terms, $\Delta_{|w_\ell|}v^*(s+ke_j) \leq \Delta_{|w_\ell|}v^*(s+e_j) \leq a_{|w_\ell|}$. Since $s+ke_j-|w_\ell| \geq s+e_j-|w_\ell| \geq 0$, it is feasible for the optimal mechanism to transition to state $s+e_j-|w_\ell|$. Hence, optimal mechanism would (at least) conduct the exchange of vector $|w_\ell|$ at state $s+ke_j$, implying that $s+ke_j$ is not reachable. □

In the following section, we utilize these insights to derive the optimal mechanism for the kidney exchange application, i.e., when $n = 4$. In this case, depending on the arrival probabilities of agent types, we can obtain more structured threshold mechanisms.

7 The Structure of the Optimal Kidney Exchange Mechanism

Recall that kidney exchange is a special case of the general problem with 4 state variables. When the state is $s = (s_1, s_2, s_3, s_4)$, s_1 refers to the number of type $\mathbb{O}_{A-B}$ pairs, s_2 is the number of type $A-B$ pairs, s_3 is the number of type $B-A$ pairs, and s_4 is the number of type $\mathbb{O}_{B-A}$ pairs. In this case,

1. $K_1 = \{1,2,4\}$, i.e., for $\mathbb{O}_{A-B}$ type pairs, marginal value decreasing exchanges are a two-way exchange of an $\mathbb{O}_{A-B}$ pair with an underdemanded pair (represented by w_1), a three-way exchange of an $\mathbb{O}_{A-B}$ with an $A-B$ and an underdemanded pair (w_2), and a three-way exchange of an $\mathbb{O}_{B-A}$ pair with a $B-A$ and an underdemanded pair (w_4).
2. $K_2 = \{2,3,5\}$, i.e., for $A-B$ type pairs, marginal value decreasing exchanges are a three-way exchange of an $\mathbb{O}_{A-B}$ with an $A-B$ and an underdemanded pair (w_2), a two-way exchange of an $A-B$ and a $B-A$ pair (w_3), and a two-way exchange of an $\mathbb{O}_{B-A}$ pair with an underdemanded pair (w_5).
3. $K_3 = \{1,3,4\}$, i.e. for $B-A$ type pairs, marginal value decreasing exchanges are a three-way exchange of an $\mathbb{O}_{B-A}$ with a $B-A$ and an underdemanded pair (w_1), a two-way exchange of an $A-B$ and a $B-A$ pair (w_3), and a two-way exchange of an $\mathbb{O}_{A-B}$ pair with an underdemanded pair (w_4).

4. $K_4 = \{2, 4, 5\}$, i.e., for $\mathbb{O}_{B-A}$ type pairs, marginal value decreasing exchanges are a three-way exchange of an $\mathbb{O}_{A-B}$ with an $A - B$ and underdemanded pair (w_2), a three-way exchange of an $\mathbb{O}_{B-A}$ pair with a $B - A$ and an underdemanded pair (w_4), and a two-way exchange of an $\mathbb{O}_{B-A}$ pair with an underdemanded pair (w_5).

We begin analyzing kidney exchange with a simple observation regarding reachable states:

Observation 5. *For each s,*

i. $s_2 > 0$ *and* $s_3 > 0$ *imply* $s \notin \mathcal{R}$,

ii. $s_1 = s_3 = s_4 = 0$ *or* $s_1 = s_2 = s_4 = 0$ *implies* $s \in \mathcal{R}$.

The first observation follows from the fact that, pair types excluding $A - B$ and $B - A$ pairs are overdemanded, and thus can be matched in arbitrary times. Thus, matching existing $A - B$ and $B - A$ pairs earlier implies no cost, but a benefit of earlier match. The second part follows from the fact that these are the states without overdemanded pairs and with at most one type of reciprocal pair, and they are always *reachable* since they do not admit any feasible exchange.

Given the state space is $\mathcal{S} = \{s = (s_1, s_2, s_3, s_4) : s_i \in \mathbb{N}\} = \mathbb{N}^4$, the set $\mathcal{R}$ lies, in general, in the four-dimensional Euclidian space.[25] We consider a special case of the problem where it is optimal not to keep any pair of a certain overdemanded type in the pool. As we see next, this special case corresponds essentially to an analysis in the usual two-dimensional space and is very useful in understanding the structure of the optimal mechanism and the intuition behind it. The general case is a mere extension of this structure and intuition.[26]

7.1 Unbalanced dynamic exchange

We call the special case of the dynamic kidney exchange problem, where it is optimal not to keep any pair of a certain overdemanded type in the pool, as *unbalanced*. The set of conditions for a problem being *unbalanced* depends on the problem parameters and we cannot provide these exact analytical conditions. But, numerically, a problem is *unbalanced* if the difference between the arrival rates of the reciprocal pairs is above a certain threshold.[27] We suppose, without loss of generality, that $A - B$ types arrive more frequently than $B - A$ pairs by a sufficiently high margin. This is the case we consider in this section (and the other case is essentially the same and symmetric to this case). The intuition for this problem being *unbalanced* is the following: It is

[25] Note that, Observation 5 implies a set of restrictions on the *reachable* states, but not on the set of dimensions.

[26] For the purposes of brevity, we defer the analysis of this general case to the Online Appendix on balanced dynamic exchange.

[27] It is important to note here that it could be that the conditions for *unbalancedness* are very weak. In fact, using the reported simulation values in Ünver (2010), the problem turns out to be *unbalanced* for all values p_{A-B}, p_{B-A} such that $p_{A-B} \neq p_{B-A}$. Thus, even minor deviations from equal probabilities can imply *unbalanced* problem. Such different probabilities that trigger the unbalanced case were also recorded in the field by Terasaki et al. (1998) in their exchange-pool sample as $p_{A-B}/p_{B-A} = 5/3$.

unlikely that, there will be many $B - A$ pairs waiting in the pool for incoming $A - B$ pairs. Thus, keeping an $\mathbb{O}_{B-A}$ pair in the pool has low value, and it is optimal not to keep this pair in the pool. Similarly, since $A - B$ type arrive more frequently and thus, there will be more $A - B$ pair arrivals to the pool expectedly, it is optimal to match these pairs with $\mathbb{O}_{A-B}$ pairs, which is implied by their marginal value being smaller than one.

Definition 15. A problem $(\mathcal{T}, (\lambda_i)_{i\in\mathcal{T}}, \rho)$ with the value function v^* is called **unbalanced**, if for each $s \in \mathcal{R}$,

- there are no $\mathbb{O}_{B-A}$ pairs available in the pool at any reachable state,
- whenever $s - e_2 \geq 0$, $v^*(s) - v^*(s - e_2) < 1$.[28]

Suppose there are $A - B$ pairs in the pool. This is the trivial case since there is no decision to make: An incoming $\mathbb{O}_{A-B}$ pairs is matched with an existing $A - B$ pair, an incoming $\mathbb{O}_{B-A}$ pair is matched with an underdemanded pair, an incoming $B-A$ pair is matched with an existing $A-B$ pair, and an incoming $A - B$ pair is pooled (since there are no exchanges available for them). Thus, we focus on the nontrivial case of reachable states: when there are no $A - B$ pairs in the pool, which means, by Observation 5 (i), and Definition 15, that, at any *reachable* state, there are (potentially) $B - A$ and $\mathbb{O}_{A-B}$ pairs in the pool.

7.1.1 Multi-dimensional threshold mechanism

We have shown that the set of *reachable* states is such that either there are only $A-B$ pairs, or $B-A$ and $\mathbb{O}_{A-B}$ pairs. Thus, for an *unbalanced* problem, we have $\mathcal{R} \subseteq \{s: s_2 = s_4 = 0\} \cup \{s: s_1 = s_3 = s_4 = 0\}$. This facilitates the use of two-dimensional space such that we can visualize the structure of the optimal mechanism. Since $A - B$ and $B - A$ pairs are never present together in the pool at any time, we can label the x-axis as $s_2 - s_3$ such that whenever it is positive, there are $A - B$ pairs (and no $B - A$ pairs) in the pool, and otherwise, there are $B - A$ pairs (and no $A - B$ pairs). Moreover, since, by definition of an *unbalanced* problem, there are no $\mathbb{O}_{B-A}$ pairs in the pool, we refer the y-axis as s_1, the number of $\mathbb{O}_{A-B}$ pairs in the pool. In Figure 2, we depict the set of all potentially *reachable* states. As we argued above, the case $s_2 > 0$ is trivial and in this section, we focus on the case $s_2 = s_4 = 0$, thus, the left-hand-side (LHS) quadrant in Figure 2.

Now, we analyze the optimal mechanism for the *unbalanced* problem. It turns out that the optimal mechanism is a generalized version of the simple threshold mechanism, where, for each arriving pair, a threshold function determines whether to keep the pair in the pool or to match it with an existing pair. These thresholds are utilized in the mechanism described in Table 1, which we refer to as the **multi-dimensional threshold mechanism**. There are three threshold functions: $t^{1,3}$ is utilized when the arriving pair is type 1 ($\mathbb{O}_{A-B}$) or 3 ($B-A$), t^2 is utilized when the

[28] This means that that is, the marginal value of $A - B$ pairs is always less than one.

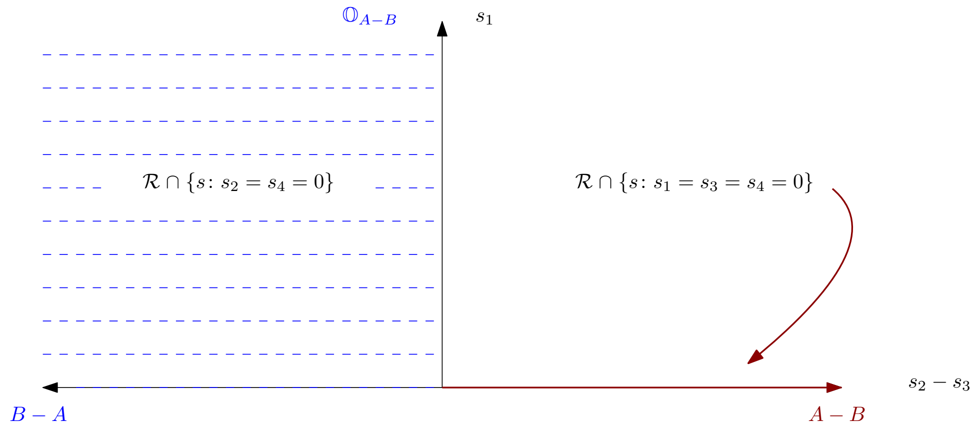

Figure 2: Depiction of the potentially *reachable* states

arriving pair is type 2 ($A-B$), and t^4 is utilized when the arriving pair is type 4 ($\mathbb{O}_{B-A}$). All three functions determine thresholds relating the number of $\mathbb{O}_{A-B}$ and $B-A$ pairs with each other. While $t^{1,3}$ determines a threshold number of $\mathbb{O}_{A-B}$ pairs as a function of s_3 (the number of $B-A$ pairs), t^2 and t^3 determine a threshold number of $B-A$ pairs as a function of s_1 (the number of $\mathbb{O}_{B-A}$ pairs). The decision of which exchange to conduct at a state s after a new pair arrives depends on whether the respective numbers s_1 and s_3 exceed the threshold. The left-hand column of Table 1 corresponds to the non-trivial case $s_2 = s_4 = 0$ (also to the LHS of Figure 2), where we explain how these thresholds are utilized. The right-hand column corresponds to the trivial case $s_1 = s_3 = s_4 = 0$ (also to the RHS of Figure 2), and the optimal decision is trivial or there is no decision to make.

The broad idea here is to prove the existence and the properties of these *threshold functions* by exploiting $\mathcal{D}^M$*-multimodularity* (Theorem 2) and the resulting second-order properties (Theorem 3) of the value function v^*. Our next result characterizes a *multi-dimensional threshold mechanism* and states its optimality.

Theorem 7. *Let $(\mathcal{T}, (\lambda_i)_{i\in\mathcal{T}}, \rho)$ be an unbalanced kidney exchange problem with the value function v^*. Then, there exist three threshold functions $t^{1,3}$, t^2, and t^4 such that the induced multi-dimensional threshold mechanism is optimal. Moreover, these threshold functions satisfy the following properties:*

- *i.* $t^{1,3}:\mathbb{N}\to\mathbb{N}$ *is a* ***non-increasing*** *function of the number of type 3 ($B-A$) pairs such that*
  - *i.1.* *for a state s with $s_1 \geq 1$, $v^*(s) - v^*(s-e_1) \leq 2$ if and only if $s_1 > t^{1,3}(s_3)$,*
  - *i.2.* *for each $k \in \mathbb{N}$, $t^{1,3}(k+1) \geq t^{1,3}(k) - 1$.*
- *ii.* $t^2:\mathbb{N}\to\mathbb{N}$ *is a* ***non-decreasing*** *function of the number of type 1 ($\mathbb{O}_{A-B}$) pairs such that*
  - *ii.1.* *for a state s with $s_3 \geq 1$ and $s_1 \geq 1$, $v^*(s-e_3) - v^*(s-e_1) \leq 1$ if and only if $s_3 < t^2(s_1)$,*
- *iii.* $t^4:\mathbb{N}\to\mathbb{N}$ *is a* ***non-increasing*** *function of the number of type 1 ($\mathbb{O}_{A-B}$) pairs such that*
  - *iii.1.* *for a state s with $s_3 \geq 1$, $v^*(s) - v^*(s-e_3) \leq 1$ if and only if $s_3 > t^4(s_1)$,*

MULTI-DIMENSIONAL THRESHOLD MECHANISM

Case 1: Arriving pair is type $\mathbb{O}_{A-B}$ (type 1)				
	$s_2 = 0$:	1.1 $s_1 + 1 \leq t^{1,3}(s_3)$: Keep in the pool. 1.2 $s_1 + 1 > t^{1,3}(s_3)$: Match with an underdemanded pair.	$s_2 > 0$:	• Match with an existing $A-B$ pair and an underdemanded pair.
Case 2: Arriving pair is type $A-B$ (type 2)				
	$s_1 > 0$ or $s_3 > 0$:	2.1 $s_1 = 0$: Match with an existing $B-A$ pair. 2.2 $s_3 = 0$: Match with an existing $\mathbb{O}_{A-B}$ pair and an underdemanded pair. 2.3 $s_1 > 0$ and $s_3 > 0$: 2.3.1 $s_3 \leq t^2(s_1)$: Match with an existing $\mathbb{O}_{A-B}$ pair and underdemanded pair. 2.3.2 $s_3 > t^2(s_1)$: Match with an existing $B-A$ pair.	$s_1 = s_3 = 0$:	• Keep in the pool.
Case 3: Arriving pair is type $B-A$ (type 3)				
	$s_2 = 0$:	3.1 $s_1 \leq t^{1,3}(s_3 + 1)$: Keep in the pool. 3.2 $s_1 > t^{1,3}(s_3 + 1)$: Keep the $B-A$ pair in the pool, match an existing $\mathbb{O}_{A-B}$ pair with an underdemanded pair.	$s_2 > 0$:	• Match with an existing $A-B$ pair.
Case 4: Arriving pair is type $\mathbb{O}_{B-A}$ (type 4)				
	$s_3 > 0$:	4.1 $s_3 \leq t^4(s_1)$: Match with an underdemanded pair. 4.2 $s_3 > t^4(s_1)$: Match with an existing $B-A$ pair and an underdemanded pair.	$s_3 = 0$:	• Match with an underdemanded pair.

Table 1: The description of the multi-threshold mechanism with threshold functions $t^{1,3}$, t^2, t^4.

iii.2. for each $k \in \mathbb{N}$, $t^4(k+1) \geq t^4(k) - 1$.

See Appendix E for its proof.

The existence of the threshold function $t^{1,3}$ and the properties *(i.1.), (i.2.)* characterizes the optimal decision on the arriving pairs of types $\mathbb{O}_{A-B}$ and $B-A$.

Property *(i.1.)* states that the marginal value of keeping a pair of type $\mathbb{O}_{A-B}$ in the pool is greater than 2 whenever the number of $\mathbb{O}_{A-B}$ pairs in the pool is less than a certain threshold that depends on the number of available $B-A$ pairs, and it is less than 2 after this threshold. The intuition simply follows from the *componentwise concavity* of v^* (see Theorem 3).

We next explain the intuition for non-increasingness of the threshold function, $t^{1,3}$: First, note that, $B-A$ pairs and $\mathbb{O}_{A-B}$ pairs are *substitutes* since both can be used for matching future excess $A-B$ pairs in the pool (this is by *ij-submodularity* for $i=1$ and $j=3$ of the value function, which follows from *$\mathcal{D}^M$-multimodularity* (Theorem 2). Thus, as the number of $B-A$ pairs in the pool increases, the marginal value of keeping an arriving pair of type $\mathbb{O}_{A-B}$ in the pool decreases, and thus, incentives to pool these pairs weaken.

Property *(i.2.)* states that this threshold is not only non-increasing but also does not decrease faster than the rate of a linear function with slope one. The intuition is *superconcavity*: Even though $B-A$ pairs are substitutes for $\mathbb{O}_{A-B}$ pairs, they are not as close of a substitute as $\mathbb{O}_{A-B}$ pairs themselves. Thus, when we have an additional $\mathbb{O}_{A-B}$ pair but one less $B-A$ pair, the effect of the additional $\mathbb{O}_{A-B}$ pair dominates, and marginal value of $\mathbb{O}_{A-B}$ pairs decreases. Thus, when we have an additional $B-A$ pair, the number of $\mathbb{O}_{A-B}$ pairs to be removed from the pool (for an immediate exchange surplus of 2) can not exceed 1.

These results regarding the threshold functions t^2 and t^4 and their interpretation are analytically symmetric (and the intuition is similar) to properties *(i.1.)* and *(i.2.)*. We skip the explanation of the intuition for brevity purposes.

We next depict the dynamically optimal mechanism on a graph. The set of reachable states $\mathcal{R}$ is defined by the function $t^{1,3}$, since it determines whether the existing $\mathbb{O}_{A-B}$ pair has a marginal value greater than 2. The other two functions t^4 and t^2 define two new regions that correspond to two different decisions. We next explain that these three regions interact in a way that allows us to obtain a simple illustration of the optimal mechanism.

To see this, we use the following observation.

Observation 6. *For a state s with* $s_1, s_3 > 0$, *if* $v^*(s) - v^*(s-e_1) < 2$ *and* $v^*(s) - v^*(s-e_3) > 1$ *then* $v^*(s-e_3) - v^*(s-e_1) < 1$; *if* $v^*(s) - v^*(s-e_1) > 2$ *and* $v^*(s) - v^*(s-e_3) < 1$ *then* $v^*(s-e_3) - v^*(s-e_1) > 1$.

First, we depict the regions induced by threshold functions in Figure 3.

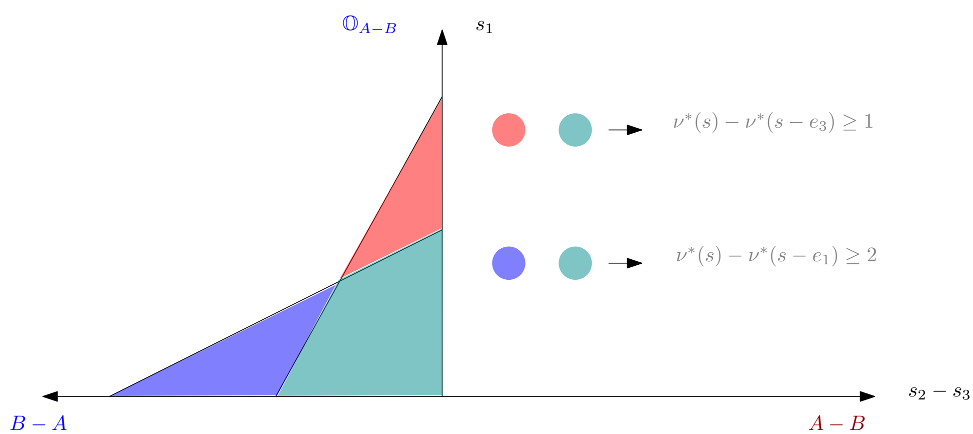

Figure 3: Regions induced by $t^{1,3}$ and t^4

Now, by Observation 6, in the red region in Figure 3, we have $v^*(s-e_3) - v^*(s-e_1) < 1$ and in the blue region we have $v^*(s-e_3) - v^*(s-e_1) > 1$. Thus, $s_3 < t^2(s_1)$ in the red region and $s_3 > t^2(s_1)$ in the blue region. This means that the function t^2, as depicted in Figure 4, should pass through the intersection point of blue, red and dark green regions in Figure 3.

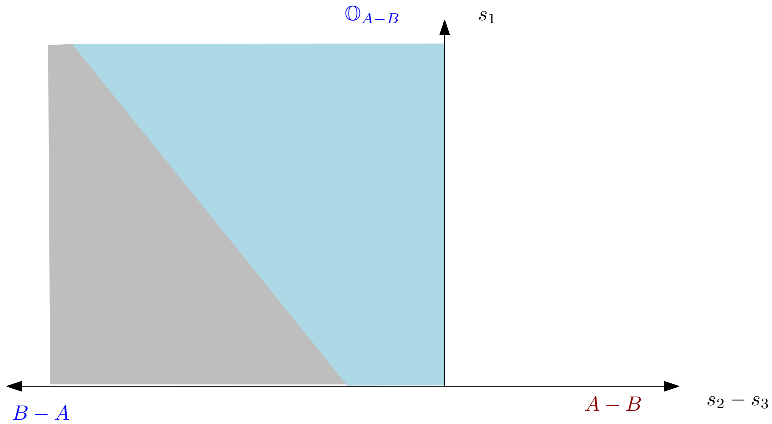

Figure 4: Regions induced by t^2

Moreover, the red region in Figure 3 is not *reachable*, and thus, irrelevant for the optimal mechanism. This leaves us with the regions in Figure 5 below, that together describe the optimal mechanism:

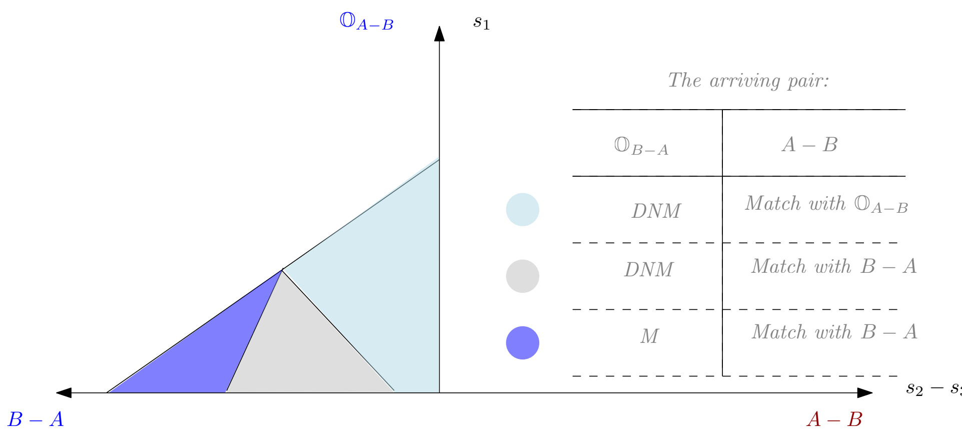

Figure 5: The multi-dimensional threshold mechanism

The three regions together with the x-*axis* characterize the set of *reachable* states. The decision of whether to keep or match an arriving $\mathbb{O}_{A-B}$ pair, as well as the decision to match an existing $\mathbb{O}_{A-B}$ pair in a *two-way* exchange when a $B-A$ pair arrives, is determined by the set of *reachable* states $\mathcal{R}$. Other decisions, namely, decision of *Match* (Decision 4.2) vs *Do-Not-Match* (DNM) (Decision 4.1) and decision of *Match with* $\mathbb{O}_{A-B}$ (Decision 2.3.1) vs *Match with* $B-A$ (Decision 2.3.2) are determined by these regions.[29] [30]

[29]Note that we do not color the axis, since they correspond to trivial decisions. When there is no $B-A$ pair, *Match* and *Match with* $B-A$ decisions are irrelevant. Similarly, when there is no $\mathbb{O}_{A-B}$ pair, *Match with* $\mathbb{O}_{A-B}$ is irrelevant.

[30]Moreover, although the structure depicted above is a theoretical necessity, what we observe numerically is much simpler. Numerically, the blue and gray regions in Figure 5 are empty, and set of *reachable* states constitute the light blue region: Whenever there exists an $\mathbb{O}_{A-B}$ pair, incoming $A-B$ pairs are matched with these overdemanded pairs in *three-way* exchanges. Moreover, the DNM region contains all *reachable* states with some $\mathbb{O}_{A-B}$ pairs in the pool. Thus, the numerical results imply that the optimal mechanism is described *(i)* by the triangular region of reachable states $\mathcal{R}$, and *(ii)* a number $\bar{s}$ denoting the threshold for the M or DNM decision. But, we are unable to show that this special case we observe numeri-

8 Conclusion

Our contributions in this paper are twofold: one pertains to the methodology of dynamic matching theory, and the other concerns the practice of market design for kidney exchange.

First, through our methodological contribution, we introduce novel tools for analyzing dynamic matching frameworks in which substitutes and complements in matching are well defined—as on a spatial linear mutual compatibility graph over types. We extend tools recently developed in queueing theory—specifically, within event-based dynamic programming using Markov Decision Processes—to dynamic matching. Standard techniques in queueing rely almost entirely on substitutable inputs. Our methodology provides an elegant and tractable framework for optimally controlling multi-dimensional state spaces by characterizing the second-order properties of the optimal value function. The optimal mechanism is a complex multi-threshold policy that prescribes conducting a certain set of exchanges when the number of agents of a given type exceeds a threshold, and otherwise taking no action, as a function of the other state variables. We also characterize the structure of these optimal exchanges. This technique is applicable to a range of practical problems and high-level models, from on-demand ride-sharing to spatial bilateral trade economies.

As our second contribution, we use this methodology to characterize the optimal dynamic kidney exchange mechanisms under certain large market assumptions, extending the work of Ünver (2010) by removing the assumption that overdemanded types are matched immediately and pointing out an erroneous interim result motivating this assumption. This application reduces to a four-state-variable case of our more general model with specific exchange surpluses. As a result, we demonstrate that a multi-threshold mechanism is optimal, controlling at most three types of pairs simultaneously: overdemanded pairs complementing $A-B$, overdemanded pairs complementing $B-A$, and either $A-B$ or $B-A$ (but not both).

When the arrival rates of patient-donor pair types $A-B$ and $B-A$—the two central adjacent types—are balanced, our general characterization applies; however, near-greedy optimization performs well in such cases (as explained in the Online Appendix). Greedy optimization matches each pair in the largest feasible exchange immediately upon arrival. This intuition aligns with several findings in the literature such as Anderson et al. (2017) and

In contrast, when arrival rates are unbalanced—e.g., when $A-B$ arrives significantly more frequently than $B-A$, as observed in past data—retaining both $B-A$ pairs and overdemanded pairs complementing $A-B$ (such as $B-O$) may be optimal. Interestingly, this policy is *anti-greedy*: while a $B-O$ pair can be immediately matched with an underdemanded pair (e.g., $O-B$), it

cally is also a theoretical necessity. Thus, whenever we mention the optimal unbalanced mechanism, we refer to the three regions depicted above. We provide a numerical example in the Online Appendix for the unbalanced case.

is optimal to retain $B-O$ pairs until their stock reaches a threshold. Such policies are dynamic and underscore the importance of explicitly modeling blood types, a practice often overlooked in more recent dynamic matching literature. Indeed, under realistic arrival rates where $A-B$ arrives more frequently than $B-A$, the thresholds can exceed 30 pairs for both $B-A$ and the overdemanded pairs that complement $A-B$.

References

Agarwal, Nikhil, Itai Ashlagi, Eduardo Azevedo, Clayton R. Featherstone, and Ömer Karaduman. 2019. "Market failure in kidney exchange." *American Economic Review* 109 (11): 4026–70.

Akbarpour, Mohammad, Julien Combe, YingHua He, Victor Hiller, Robert Shimer, and Olivier Tercieux. 2024. "Unpaired Kidney Exchange: Overcoming Double Coincidence of Wants without Money." *The Review of Economic Studies* rdae081. 10.1093/restud/rdae081.

Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan. 2020. "Thickness and information in dynamic matching markets." *Journal of Political Economy* 128 (3): 783–815.

Altman, Eitan, Bruno Gaujal, and Arie Hordijk. 2003. *Discrete-event control of stochastic networks: Multimodularity and Regularity*. Springer, Berlin Heidelberg New York Barcelona Hong Kong London Milan Paris Tokyo.

Anderson, Ross, Itai Ashlagi, David Gamarnik, and Yash Kanoria. 2017. "Efficient dynamic barter exchange." *Operations Research* 65 (6): 1446–1459.

Ausubel, Lawrence M, and Thayer Morrill. 2014. "Sequential kidney exchange." *American Economic Journal: Microeconomics* 6 (3): 265–285.

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv. 2020. "Optimal dynamic matching." *Theoretical Economics* 15 (3): 1221–1278.

Bloch, Francis, and David Cantala. 2017. "Dynamic assignment of objects to queuing agents." *American Economic Journal: Microeconomics* 9 (1): 88–122.

Che, Yeon-Koo, and Olivier Tercieux. 2021. "Optimal queue design." In *Proceedings of the 22nd ACM Conference on Economics and Computation,* 312–313.

Doval, Laura. 2022. "Dynamically stable matching." *Theoretical Economics* 17 (2): 687–724.

Doval, Laura. 2025. "Dynamic Matching." In *Handbook of the Economics of Matching, forthcoming,* edited by Yeon-Koo Che, Bernard Salanié, Pierre-André Chiappori Volume 2. Elsevier.

Furian, Lucrezia, Antonio Nicolò, Caterina Di Bella, Massimo Cardillo, Emanuele Cozzi, and Paolo Rigotti. 2020. "Kidney exchange strategies: new aspects and applications with a focus on deceased donor-initiated chains." *Transplant International* 33 (10): 1177–1184.

Hajek, Bruce. 1985. "Extremal Splittings of Point Processes." *Mathematics of Operations Research* 10(4) 543–556.

Hotelling, Harold. 1929. "Stability in competition." *Economic Journal* 39 (153): 41–57.

Kerimov, Süleyman, Itai Ashlagi, and Itai Gurvich. 2025. "On the optimality of greedy policies in dynamic matching." *Operations Research* 73 560–582.

Koole, Ger. 1998. "Structural results for the control of queueing systems using event-based dynamic programming." *Queueing Systems* 30 323–339.

Kurino, Morimitsu. 2020. "Credibility, efficiency, and stability: A theory of dynamic matching markets." *The Japanese Economic Review* 71 (1): 135–165.

Leshno, Jacob D. 2022. "Dynamic matching in overloaded waiting lists." *American Economic Review* 112 (12): 3876–3910.

Massie, Allan B, Babak J Orandi, Madeleine M Waldram, Xun Luo, Anh Q Nguyen, Robert A Montgomery, Krista L Lentine, and Dorry L Segev. 2020. "Impact of ABO-incompatible living donor kidney transplantation on patient survival." *American Journal of Kidney Diseases* 76 (5): 616–623.

Puterman, Martin L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* John Wiley & Sons, Inc., Hoboken, New Jersey.

Rapaport, Felix T. 1986. "The case for a living emotionally related international kidney donor exchange registry." *Transplantation Proceedings* 18 5–9.

Rees, Michael A., Jonathan E. Kopke, Ronald P. Pelletier et al. 2009. "A Non-simultaneous Extended Altruistic Donor Chain." *The New England Journal of Medicine* 360 1096–1101.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2004. "Kidney Exchange." *Quarterly Journal of Economics* 119 457–488.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2005. "Pairwise Kidney Exchange." *Journal of Economic Theory* 125 151–188.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2007. "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences." *American Economic Review* 97 828–851.

Roth, Alvin E., Tayfun Sönmez, M. Utku Ünver, Francis L. Delmonico, and Susan L. Saidman. 2006. "Utilizing List Exchange and Nondirected Donation Through Chain Paired Kidney Donations." *American Journal of Transportation* 6 2694–2705.

Schinstock, Carrie A, Andrew J Bentall, Byron H Smith, Lynn D Cornell, Matthew Everly, Manish J Gandhi, and Mark D Stegall. 2019. "Long-term outcomes of eculizumab-treated positive crossmatch recipients: allograft survival, histologic findings, and natural history of the donor-specific antibodies." *American Journal of Transplantation* 19 (6): 1671–1683.

Sönmez, Tayfun, and M. Utku Ünver. 2025. "Matching under Non-transferable Utility: Applications." In *Handbook of the Economics of Matching, forthcoming*, edited by Yeon-Koo Che, Bernard Salanié, Pierre-André Chiappori Volume 2. Chap. 3, Elsevier.

Sönmez, Tayfun, M. Utku Ünver, and M. Bumin Yenmez. 2020. "Incentivized Kidney Exchange." *American Economic Review* 110(7) 2198–2224.

Terasaki, Paul I, David W Gjertson, and J Michael Cecka. 1998. "Paired kidney exchange is not a solution to ABO incompatibility." *Transplantation* 65 (2): 291.

Ünver, M. Utku. 2010. "Dynamic Kidney Exchange." *The Review of Economic Studies* 77 (1) 372–414.

Zenios, Stefanos A., E. Steve Woodle, and Lainie Friedman Ross. 2001. "Primum Non Nocere: Avoiding Increased Waiting Times for Individual Racial and Blood-type Subsets of Kidney Wait List Candidates in a Living Donor/cadaveric Donor Exchange Program." *Transplantation* 72 648–654.

Appendix A Proof of Lemma 2

Claim 1. *If the properties* $P(u,v)$ *and* $P(u,w)$ *hold for some function* f, *then* $P(u,v+w)$ *also holds.*

Proof. We first rewrite the definitions of $P(u,v)$ and $P(u,w)$: for each x, and x',

$$f(x+u)+f(x+v) \le f(x)+f(x+u+v),$$
$$f(x'+u)+f(x'+w) \le f(x')+f(x'+u+w).$$

Substituting $x' = x+v$ into the second inequality, we obtain

$$f(x+u+v)+f(x+w+v) \le f(x+v)+f(x+u+v+w).$$

Lastly, summing the last and first inequalities, we obtain for each x,

$$f(x+u)+f(x+v+w) \le f(x)+f(x+u+v+w),$$

which is equivalent to $P(u,v+w)$. □

Let $U = \{u_1,u_2,...u_k\}$ and $V = \{v_1,v_2...,v_l\}$. Since each of these vectors is a distinct member of $\mathcal{D}^M$, we have the property $P(u_i,v_j)$ for any two indices i,j. Then, using Claim 1, we have $P(u_1,v_1+v_2)$, and repeatedly using the claim we have $P(u_1,\sum_{v\in V} v)$. Again, using the same argument for u_2 we have $P(u_2,\sum_{v\in V} v)$, and combining with the previous conclusion we have $P(u_1+u_2,\sum_{v\in V} v)$. We iterate using the claim—this time on u_i's—to obtain $P(\sum_{u\in U} u,\sum_{v\in V} v)$, completing the proof.

Appendix B Complete Proof of Proposition 2

To complete the proof of Proposition 2, we first present the following lemma, which deals with the boundary cases in the proof.

Lemma 3. *Let* $x \in \mathbb{N}^4$, *and let* $u,v,w \in \mathcal{D}^M$ *be three distinct vectors in* $\mathcal{D}^M$ *such that* $x+u, x+v \ge 0$. *If* $w \le 0$ *and* $x+u+w \ge 0$ *or* $x+v+w \ge 0$, *then* $x+u+v+w \ge 0$. *If* $w \ge 0$ *and* $x+u-w \ge 0$ *or* $x+v-w \ge 0$, *then* $x-w \ge 0$

Proof of Lemma 3. First, suppose $w \le 0$, $x+u+w \ge 0$ or $x+v+w \ge 0$. Now, suppose $x_i+u_i+v_i+w_i < 0$ for some i. Since $x_i \ge 0$, we must have $u_i+v_i+w_i < 0$. By construction of $\mathcal{D}^M$, exactly one of this numbers is -1 whereas the other two are zero.

Suppose $w_i = -1$ and $u_i = v_i = 0$. Then, we must have $x_i = 0$ as well, since if $x_i > 0$, we would have $x_i+u_i+v_i+w_i = 0$. Then, we have $x_i+u_i+w_i = x_i+v_i+w_i = -1$. This means $x+u+w \ngeq 0$ and $x+v+w \ngeq 0$ and contradicts our assumption.

Now suppose $w_i = 0$, then, $x_i+u_i+v_i+w_i = x_i+u_i$ or $x_i+u_i+v_i+w_i = x_i+v_i$, since one of u_i and v_i equals to zero by the above observation. But by assumption, we have $x_i+v_i \ge 0$ and $x_i+u_i \ge 0$, contradicting $x_i+u_i+v_i+w_i < 0$

Now we consider with the symmetric case of $w \geq 0$. Suppose $x_i - w_i < 0$. Then, we must have $w_i = 1$ and $x_i = 0$. Then, by the observation above, exactly one of u_i and v_i equals to zero while the other equals to -1. Thus, $u_i, v_i \leq 0$, and $x_i + u_i - w_i \leq x_i - w_i < 0$ and $x_i + v_i - w_i \leq x_i - w_i < 0$, contradicting our assumption. □

Finally, we present the proof for the case $w \geq 0$, in which case the *matching operator* takes the form $Tf = \max\{f(x), f(x-w) + a_w\}$.

We have to show

$$\begin{aligned} &\max\{f(x+u), f(x+u-w) + a_w\} \;+\; \max\{f(x+v), f(x+v-w) + a_w\} \\ \leq\; &\max\{f(x), f(x-w) + a_w\} \;+\; \max\{f(x+u+v), f(x+u+v-w) + a_w\}. \end{aligned}$$

First suppose u, v, w are distinct elements from $\mathcal{D}^M$.

Again, symmetric actions are trivial. If the actions are different, then by Lemma 3, we have $x - w \geq 0$, thus, the right-hand side of the inequality is weakly greater than $f(x-w) + f(x+u+v) + a_w$. In different action cases, the left-hand side is equal to either $f(x+u) + f(x+v-w) + a_w$ or $f(x+v) + f(x+u-w) + a_w$. Since $x - w \geq 0$, we can let $x' = x - w \geq 0$ to write the left-hand side as $f(x' + u + w) + f(x' + v) + a_w$ or $f(x' + u) + f(x + v + w) + a_w$. By *multimodularity*, both of these expressions are weakly less then $f(x') + f(x' + u + v + w) + a_w$, which by substituting $x' = x - w$ equals to $f(x-w) + f(x+u+v) + a_w$. We showed this expression is weakly less than the right-hand side, so the proof ends here.

Now suppose wlog that $w = u$. Then, the expression becomes

$$\begin{aligned} &\max\{f(x+u), f(x) + a_u\} \;+\; \max\{f(x+v), f(x+v-u) + a_u\} \\ \leq\; &\max\{f(x), f(x-u) + a_u\} \;+\; \max\{f(x+u+v), f(x+v) + a_u\}. \end{aligned}$$

After noting that symmetric actions are trivial and $x+u, x+v \geq 0$, we show that the left-hand side can only be equal to $f(x) + f(x+v) + a_u$, if the actions are different. First, note that if $x + v - u \ngeq 0$, then this is trivial. Then, assuming $x + v - u \geq 0$, we can let $x' = x - u \geq 0$, and by combining *concavity* and *multimodularity* as above, we can write

$$f(x' + v) - f(x' + u + v) \leq f(x' + u) - f(x' + 2u)$$

and substituting again we have

$$f(x + v - u) - f(x + v) \leq f(x) - f(x + u).$$

Thus, if $f(x+u) \geq f(x) + a_u$ then $f(x+v-u) - f(x+v) \leq f(x) - f(x+u) \leq -a_u$, so $f(x+v-u) + a_u \leq f(x+v)$. This shows the only possible case for actions being different is if the left-hand side is equal to $f(x) + f(x+v) + a_u$. Finally, we see that picking the first argument in the first max and the second argument in the second max on the right-hand side, the right-hand side becomes $f(x) + f(x+v) + a_u$, completing the proof.

Appendix C Proof of Theorem 3

As we mentioned earlier, all the properties of the theorem are of the form $P(u,v)$ for some vectors $u,v \in \mathbb{Z}^n$. By Theorem 2, the value function v^* is $\mathcal{D}^M$*-multimodular*. Thus, by Lemma 2, $P(u,v)$ holds for any u,v such that $u = \sum_{u'\in U} u'$, $v = \sum_{v'\in V} v'$ and U,V are disjoint subsets of $\mathcal{D}^M$. It suffices to show that each of the properties is equivalent to $P(u,v)$ for some u,v that are sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$.

First, we make the simple observation that the first i terms in $\mathcal{D}^M$ sum up to e_i if i is odd and $-e_i$ if i is even. Similarly, since elements of $\mathcal{D}^M$ sum to 0, the remaining $n+1-i$ elements sum up to $-e_i$ and e_i if i is odd or even, respectively.

From this observation, it follows that e_i and $-e_i$ are always sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$. Since concavity in component i is equivalent to $P(e_i,-e_i)$, componentwise concavity follows.

Suppose $i,j \in \tau$ are distinct types with the same parity, i.e., both are odd or both are even, and suppose without loss of generality that $i < j$. First, suppose i and j are odd. Then, the first i terms sum up to e_i and last $n+1-j$ terms sum up to $-e_j$, and since $i < j$, these sums are disjoint. If i and j are even, the first i terms sum up to $-e_i$, and the last $n+1-j$ terms sum up to e_j. Since $i < j$, these sums are disjoint. Thus in either case, $P(e_i,-e_j)$ and $P(-e_i,e_j)$ hold, respectively. They are both equivalent to each other and to ij*-submodularity*.[31]

Next, suppose $i,j \in \tau$ have different parities. By the same argument as before, we can take the first i terms and last $n+1-j$ terms for $i < j$ to obtain conclude $P(e_i,e_j)$ or $P(-e_i,-e_j)$ holds. These are both equivalent to ij*-supermodularity*.

For *superconcavity*, take any two distinct $i,j \in \tau$. First, we observe that if i and j have different parities, so that ij*-supermodularity* is satisfied, ij*-superconcavity* is trivial. We can see this by writing

$$v^*(x+e_i) + v^*(x+e_i+e_j) \;\geq\; v^*(x+e_i) + v^*(x+e_i) + v^*(x+e_j) - v^*(x) \;\geq\; v^*(x+e_j) + v^*(x+2e_i),$$

where the first inequality follows from ij*-supermodularity* and the second inequality follows from *concavity in component* i.

Suppose i and j have the same parity. Then, *superconcavity* in these components is equivalent to $P(e_i, e_j - e_i)$, which is equivalent to $P(-e_i, e_i - e_j)$. Assume i and j are odd. If $i < j$, then the first i terms sum up to e_i and the first j terms sum up to e_j. Thus, terms from $i+1$ to j sum up to $e_j - e_i$. Thus, e_i and $e_j - e_i$ are sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$, which proves $P(e_i, e_j - e_i)$. If $i > j$, the last $n+1-i$ terms sum to $-e_i$, and the first j terms sum to e_j. Thus, terms from $j+1$ to i sum to $e_i - e_j$, and $-e_i$ and $e_i - e_j$ are sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$, which proves $P(-e_i, e_i - e_j)$. In either case, we have ij*-superconcavity*. When i and j are even the symmeric proof holds. We showed that v^* is *superconcave*.

[31] $P(u,v)$ is equivalent to $P(-u,-v)$. This can be observed by the simple change of variables $x' = x - u - v$.

$$
\begin{aligned}
w_k \in \{w_\ell : \ell \in K_i\} = \{|w_1|, |w_3|, \ldots, |w_i|\} \cup \{-|w_{i+1}|, -|w_{i+3}|, \ldots\} \\
\text{if } k > i: \\
\mathcal{D}^M = \{\underbrace{|w_1|, -|w_2|, |w_3|, \ldots, |w_i|}_{e_i=\sum_{\ell=1}^{i} w_\ell}, -|w_{i+1}|, \ldots \underbrace{-|w_k|}_{=w_k} \ldots, (-1)^n|w_{n+1}| \\
e_i - |w_k| = \sum_{\ell \in \{1,2,\ldots,i\}\cup\{k\}} w_\ell \quad \text{is a sum of distinct vectors in } \mathcal{D}^M \\
\text{if } k \leq i: \\
\mathcal{D}^M = \{\underbrace{|w_1|, -|w_2| \ldots, \overbrace{|w_k|}^{=w_k}, \ldots, |w_i|}_{e_i=\sum_{\ell=1}^{i} w_\ell}, -|w_{i+1}|, \ldots, (-1)^n|w_{n+1}| \\
e_i - |w_k| = \sum_{\ell \in \{1,2,\ldots,i\}\setminus\{k\}} w_\ell \quad \text{is a sum of distinct vectors in } \mathcal{D}^M
\end{aligned}
$$

Figure 6: Illustration of Observation 7 when i is odd and $B = \{k\} \subseteq K_i$.

Appendix D Proofs of Theorems 4 and 5

Before proving Theorem 4, we will make a couple of observations.

First, note that if the optimal mechanism transitions to $s + e_i - \sum_k \#_k|w_k|$ after a type i agent arrives at state s, then for each ℓ such that $\#_\ell > 0$, we must have that $\Delta_{|w_\ell|}v^*(s + e_i - \sum_k \#_k|w_k| + |w_\ell|) = v^*(s + e_i - \sum_k \#_k|w_k| + |w_\ell|) - v^*(s + e_i - \sum_k \#_k|w_k|) < a_{w_\ell}$, since otherwise optimal mechanism would have kept an additional vector w_ℓ as the marginal value of the vector would be greater than a_{w_ℓ}.

Second, if s is reachable, and $s - |w_\ell| \geq 0$, we must have that $\Delta_{|w_\ell|}v^*(s) > a_{w_\ell}$, since otherwise s would not be reachable. Thus, we can further conclude that $\Delta_{|w_\ell|}v^*(s + e_i - \sum_k \#_k|w_k| + |w_\ell|) < a_{w_\ell} < \Delta_{|w_\ell|}v^*(s)$. Thus, if we show that v^* satisfies property $P(|w_\ell|, e_i - \sum_k \#_k|w_k| + |w_\ell|)$ and that $s - |w_\ell| \geq 0$, we will arrive at a contradiction, since $P(|w_\ell|, e_i - \sum_k \#_k|w_k| + |w_\ell|)$ implies $\Delta_{|w_\ell|}v^*(s + e_i - \sum_k \#_k|w_k| + |w_\ell|) > \Delta_{|w_\ell|}v^*(s)$.

Third, e_i is the sum of the first i or last $n+1-i$ vectors in $\mathcal{D}^M$, depending on whether i is odd or even. Moreover, if w_k is a vector not included in this sum for any $k \in K_i$ then k is even and, hence, as $w_k < 0$ and $|w_k| = -w_k$, we obtain $e_i - |w_k| = e_i + w_k$ is also a sum of distinct vectors in $\mathcal{D}^M$ (i.e., the sum of matching vectors included in the sum equating e_i and w_k). Similarly, if w_k is a vector included in this sum for any $k \in K_i$, then k is odd, and hence, as $w_k > 0$ and $w_k = |w_k|$, we obtain $e_i - |w_k| = e_i - w_k$ is also a sum of distinct vectors in $\mathcal{D}^M$ (i.e., the sum of matching vectors included in the sum equating e_i except w_k). Therefore, $e_i - |w_k|$ for any $k \in K_i$ is a sum of distinct vectors in $\mathcal{D}^M$ (e.g., see Figure 6 when i is odd). We observe that this argument can be iterated, and thus, we obtain the following:

Observation 7. *For any $B \subseteq K_i$, $e_i - \sum_{k\in B}|w_k|$ is a sum of distinct vectors in $\mathcal{D}^M$. This sum includes all vectors w_k that are included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i (first i or last $n+1-i$ vectors depending on the parity of i) except those whose indices are included in B, and also includes all vectors w_k for $k \in B$ such that w_k is not included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i.*

Similarly, for any $A \subseteq K_j^c$, $e_i + \sum_{k\in A}|w_k|$ is a sum of distinct vectors in $\mathcal{D}^M$. This sum includes all vectors w_k that are included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i (first i or last $n+1-i$ vectors depending on the parity of i) except those whose indices are included in A, and also includes all vectors

w_k for $k \in A$ such that w_k is not included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i.

Moreover, if we have $A \subseteq K_j^c$ and $B \subseteq K_j$, we can merge the two cases, and conclude that $e_j + \sum_{k \in A} |w_k| - \sum_{k \in B} |w_k|$ is a sum of distinct set of vectors in $\mathcal{D}^M$.

We are ready to state the proof of the theorem:

Proof of Theorem 4. Suppose s is a reachable state and the optimal mechanism uses $\#_\ell \geq 0$ exchanges of each exchange type w_ℓ after an agent of type i arrives. We will show that $\#_\ell \leq 1$ for each ℓ and that $\#_\ell = 0$ for any $\ell \notin K_i$. Let $A = K_i \cap \{\ell : \#_\ell > 0\}$, and $\#'_\ell = \begin{cases} \#_\ell - 1 & \text{if } \ell \in A \\ 0 & \text{if } \ell \notin A \end{cases}$. Then, $e_i - \sum_\ell \#_\ell |w_\ell| = e_i - \sum_{\ell \in A} |w_\ell| - \sum_\ell \#'_\ell |w_\ell|$. By Observation 7, $e_i - \sum_{\ell \in A} |w_\ell|$ is a sum of distinct vectors in $\mathcal{D}^M$. We will show that $\#'_\ell = 0$ for any ℓ. We will start with showing that if $\#'_\ell > 0$, then v^* satisfies $P(|w_\ell|, e_i - \sum_k |w_k|)$.

Suppose $\#'_\ell > 0$. First, assume $\ell \notin K_i$: two cases exist depending on whether ℓ is odd or even.

If ℓ is odd: Then $w_\ell = |w_\ell|$, and by definition of K_i, as $\ell \notin K_i$, w_ℓ is not included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i. Moreover, it is not included in the sum of matching vectors in $\mathcal{D}^M$ equating $e_i - \sum_{k \in A} |w_k|$, since we only potentially add or remove elements inside $A \subseteq K_i$ when we add the term $-\sum_{k \in A} |w_k|$. Thus, $w_\ell = |w_\ell|$ and $e_i - \sum_{\ell \in A} |w_\ell|$ are sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$, which means $P(|w_\ell|, e_i - \sum_{\ell \in A} |w_\ell|)$ holds for v^*.

If ℓ is even: Then $w_\ell = -|w_\ell|$, and since $\ell \notin K_i$, $w_\ell = -|w_\ell|$ is included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i. Moreover, it is also included in the sum of matching vectors in $\mathcal{D}^M$ equating $e_i - \sum_{\ell \in A} |w_\ell|$, since we can only add or remove elements in $A \subseteq K_i$ when we add the term $-\sum_{k \in A} |w_k|$, and $\ell \notin K_i$. Thus, w_ℓ is included in the sum $e_i - \sum_{k \in A} |w_k|$. Therefore, it is not included in the sum $-(e_i - \sum_{k \in A} |w_k|)$, which consists of the sum of the remaining vectors in $\mathcal{D}^M$. Hence w_ℓ and $-(e_i - \sum_{k \in A} |w_k|)$ are sums of distinct vectors in disjoint subsets of $\mathcal{D}^M$, and thus we have $P(w_\ell, -(e_i - \sum_{k \in A} |w_k|))$ holds for v^*. Since $-w_\ell = |w_\ell|$, equivalently we have $P(|w_\ell|, e_i - \sum_{k \in A} |w_k|)$ holds for v^*, which is what we wanted to show.

Now assume that $\ell \in K_i$. By definition of $\#'_\ell$, we have that $\#_\ell > 1$ and $\ell \in A$. Thus, either ℓ is odd and $w_\ell = |w_\ell|$ is included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i but it is not included in the sum of matching vectors in $\mathcal{D}^M$ equating $e_i - \sum_{k \in A} |w_k|$ since it is removed with the term $-\sum_{k \in A} |w_k|$. This case yields $P(|w_\ell|, e_i - \sum_{k \in A} |w_k|)$ holds for v^*. Or, ℓ is even and $w_\ell = -|w_\ell|$, and thus, it is not included in the sum of matching vectors in $\mathcal{D}^M$ equating e_i but is included in $e_i - \sum_{\ell \in A} |w_\ell|$ since it is added with the term $-\sum_{\ell \in A} |w_\ell|$. This yields $P(w_\ell, -(e_i - \sum_{\ell \in A} |w_\ell|))$ holds for v^* which means equivalently $P(|w_\ell|, e_i - \sum_{\ell \in A} |w_\ell|)$ holds for v^*.

We thus showed that for each $\#'_\ell > 0$, we have $P(|w_\ell|, e_i - \sum_{\ell \in A} |w_\ell|)$ holds for v^*. Summing all of these inequalities, we conclude that $P(\sum_\ell \#'_\ell |w_\ell|, e_i - \sum_{\ell \in A} |w_\ell|)$ holds for v^*.

We will now show that if $s + e_i - \sum_{\ell \in A} |w_\ell| - \sum_\ell \#'_\ell |w_\ell| \geq 0$, then we have $s - \sum_\ell \#'_\ell |w_\ell| \geq 0$ as well. To see this, note that each entry in $-\sum_{\ell \in A} |w_\ell|$ is non-positive, thus, only positive contribution to $s + e_i - \sum_{\ell \in A} |w_\ell| - \sum_\ell \#'_\ell |w_\ell|$ comes from e_i. This means that only index of $s - \sum_\ell \#'_\ell |w_\ell|$ that can be negative is the i^{th} index. But only two vectors w_ℓ that has a non-zero entry in the i^{th} index are w_i and w_{i+1}, both of which are included in K_i. Thus, if i^{th} index of $s - \sum_\ell \#'_\ell |w_\ell|$ is negative, it must be due to some $k \in K_i$ (more specifically, some $k \in \{i, i+1\}$), such that $\#'_k > 0$. By definition of $\#'_k$, we have that $k \in A$, and thus, $|w_k|$ is also included in the sum $-\sum_{\ell \in A} |w_\ell|$. But since $s + e_i - \sum_{\ell \in A} |w_\ell| - \sum_\ell \#'_\ell |w_\ell| \geq 0$, and the $+e_i$ term can only compensate for one term of the form $-|w_k|$, we must also have that $s - \sum_\ell \#'_\ell |w_\ell| \geq 0$.

Finally, using our observations, we can see that property $P(\sum_\ell \#'_\ell |w_\ell|, e_i - \sum_{\ell \in A} |w_\ell|)$ for v^* together with $s - \sum_\ell \#'_\ell |w_\ell| \geq 0$ and the fact that optimal mechanism transitions to $s + e_i - \sum_{\ell \in A} |w_\ell| - \sum_\ell \#'_\ell |w_\ell|$, contradicts with the state s being reachable, since the optimal mechanism would instead transition to $s - \sum_\ell \#'_\ell |w_\ell|$. Thus, we must have that $\#'_\ell = 0$ for all ℓ, which is what we wanted to show. □

We will now make a generalization to Observation 7, which will allow us to prove Theorem 5.

Lemma 4. *Let $A, A', B, B' \subseteq \{1, ..., n+1\}$ be pairwise disjoint sets of indices such that $A, A' \subseteq K_j^c$ and $B, B' \subseteq K_j$. Then, $P(\sum_{k \in A} |w_k| - \sum_{k \in B} |w_k|, e_j + \sum_{k \in A'} |w_k| - \sum_{k \in B'} |w_k|)$ holds for v^*.*

Proof of Lemma 4. By Observation 7, $e_j + \sum_{k \in A'} |w_k| - \sum_{k \in B'} |w_k|$ is a sum of distinct vectors in $\mathcal{D}^M$, moreover, since $A \cap A' = \emptyset$, for any $k \in A$, $|w_k|$ is not part of this sum of distinct vectors. Thus, we have that for any $k \in A$, property $P(|w_k|, e_j + \sum_{\ell \in A'} |w_\ell| - \sum_{\ell \in B'} |w_\ell|)$ holds for v^* by Lemma 2. Similarly, since $B \cap B' = \emptyset$, for any $k \in B$, $-|w_k|$ is not part of this sum of distinct vectors ,thus we have $P(-|w_k|, e_j + \sum_{\ell \in A'} |w_\ell| - \sum_{\ell \in B'} |w_\ell|)$ holding for v^*. Summing the relevant inequalities, we observe that property $P(\sum_{k \in B} |w_k| - \sum_{k \in A} |w_k|, e_j + \sum_{k \in A'} |w_k| - \sum_{k \in B'} |w_k|)$ holds for v^* by Claim in the proof of Lemma 2. □

We also use the following simple property of the optimal mechanism in proving Theorem 5.

Lemma 5. *Let s be a reachable state such that optimal mechanism transitions to $s + e_i - \sum_{k \in A(s,i)} |w_k|$ after an agent of type i arrives. Let $u = \sum_{k \in A} |w_k| - \sum_{k \in B} |w_k| \neq 0$ where $A \subseteq \{1, \ldots, n+1\}$, $B \subseteq A(s,i)$, and $s + e_i - \sum_{k \in A(s,i)} |w_k| - u \geq 0$. Then,*

$$\Delta_u v^*(s + e_i - \textstyle\sum_{k \in A(s,i)} |w_k|) = v^*(s + e_i - \sum_{k \in A(s,i)} |w_k|) - v^*(s + e_i - \sum_{k \in A(s,i)} |w_k| - u) > a_u$$

where a_u is the sum of rewards of the matching vectors with indices included in A minus the sum of rewards of the matching vectors with indices included in B.

Proof of Lemma 5. Since $B \subseteq A(s,i)$ and $s + e_i - \sum_{k \in A(s,i)} |w_k| - u \geq 0$, it is feasible for the optimal mechanism to transition to state $s + e_i - \sum_{k \in A(s,i)} |w_k| - u$, by using vectors from A and abstain from using vectors from B. In doing so, the mechanism would get the additional rewards of vectors from A and lose rewards of vectors from B, thus having the additional reward of a_u. However, this would result in transitioning to the state $s + e_i - \sum_{k \in A(s,i)} |w_k| - u$ instead of $s + e_i - \sum_{k \in A(s,i)} |w_k|$, which means losing the value $\Delta_u v^*(s + e_i - \sum_{k \in A(s,i)} |w_k|)$. Since optimal mechanism decides to not use these additional sum of vectors u (even though it is feasible to do so), it must be the case that $\Delta_u v^*(s + e_i - \sum_{k \in A(s,i)} |w_k|) > a_u$. □

Proof of Theorem 5. For simplicity, given a set of indices A of matching vectors in $\mathcal{D}^M$, as denoted in Equation (9), we will write $s - \Sigma[A]$ instead of $s - \sum_{k\in A}|w_k|$.

Suppose s is a reachable state and that j is an agent type such that $s+e_j$ is also reachable. Then, when the arriving pair is of some type i, the mechanism transitions to $s+e_i-\Sigma[A(s,i)]$ from state s and to $s+e_j+e_i-\Sigma[A(s+e_j,i)]$ from $s+e_j$.

Next, by rearranging the terms in $s+e_j+e_i-\Sigma[A(s+e_j,i)]$, we can obtain the following decomposition:

$$\begin{aligned}
&s+e_j+e_i-\Sigma\Big[A(s+e_j,i)\Big]\\
&\quad= s+e_i+e_j-\Sigma\Big[A(s,i)\Big]-\Sigma\Big[\big(A(s+e_j,i)\setminus A(s,i)\big)\Big]+\Sigma\Big[\big(A(s,i)\setminus A(s+e_j,i)\big)\Big]\\
&\quad= s+e_i-\Sigma\Big[A(s,i)\Big]+e_j-\Sigma\Big[\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j\Big]-\Sigma\Big[\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j^c\Big]\\
&\qquad\quad+\Sigma\Big[\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j\Big]+\Sigma\Big[\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j^c\Big]\\
&\quad= s+e_i-\Sigma\Big[A(s,i)\Big]+\overbrace{e_j-\Sigma\Big[\underbrace{\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j}_{\subseteq K_j}\Big]+\Sigma\Big[\underbrace{\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j^c}_{\subseteq K_j^c}\Big]}^{\text{Define } v:=}\\
&\qquad\quad-\Big(\overbrace{\Sigma\Big[\underbrace{\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j^c}_{\subseteq K_j^c}\Big]-\Sigma\Big[\underbrace{\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j}_{\subseteq K_j}\Big]}^{\text{Define } u:=}\Big).
\end{aligned} \tag{12}$$

Our proof strategy here is to show that vector u defined in (12) satisfies $u=0$, and in particular, each of the sums in the last line of (12) are taken over empty sets of vectors, i.e., both sets $\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j^c$ and $\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j$ are empty. This will imply that $A(s+e_j,i)\subseteq A(s,i)\cup K_j$ and $A(s,i)\cap K_j\subseteq A(s+e_j,i)$, respectively, completing the proof of the theorem.

Suppose $u\neq 0$. We will first use Lemma 5, with $s' := s+e_j$ and $u' := -u$.

Using the above decomposition, the vector $u' = -u$ is of the form $\Sigma[A]-\Sigma[B]$ where $B=(A(s+e_j,i)\setminus A(s,i))\cap K_j^c \subseteq A(s+e_j,i)$. We will show that $s+e_i-\Sigma[A(s,i)]+v \geq 0$. Suppose the ℓ^{th} index of $s+e_i-\Sigma[A(s,i)]+v$ is negative. Since $s+e_i-\Sigma[A(s,i)]+v-u\geq 0$, it must be the case that ℓ^{th} index of $\Sigma[(A(s,i)\setminus A(s+e_j,i))\cap K_j]$ is positive, since the other non-negative summation term in u enters (12) with a minus sign and is overall non-positive. Moreover, since $s+e_i-\Sigma[A(s,i)]$ is also a reachable state by definition, $s+e_i-\Sigma[A(s,i)] \geq 0$. Hence, ℓ^{th} index of v must be negative, and since other terms in v are non-negative, ℓ^{th} index of $-\Sigma[(A(s+e_j,i)\setminus A(s,i))\cap K_j]$ must be negative.

Notice that each expression of the form $\Sigma[\cdot]$ is a sum of absolute values of distinct vectors in $\mathcal{D}^M$, and for each index ℓ there are exactly two vectors in $\mathcal{D}^M$ that have non-zero entries in their ℓ^{th} index, w_ℓ and $w_{\ell+1}$. However, the sets $(A(s,i)\setminus A(s+e_j,i))\cap K_j$ and $(A(s+e_j,i)\setminus A(s,i))\cap K_j$ are disjoint, thus, it cannot be the case that w_ℓ or $w_{\ell+1}$ belong to both of them.

Thus, it must be the case that one of these sets contain ℓ whereas the other contains $\ell+1$. However, both of these sets are subsets of K_j, which does not contain any consecutive elements except for j and

$j+1$. Thus, we must have that $\ell = j$. However, since there is an e_j term in v, this term can compensate for the -1 in the $-\Sigma[(A(s+e_j,i)\setminus A(s,i))\cap K_j]$, which means $s+e_i-\Sigma[A(s,i)]+v \geq 0$. Thus, using Lemma 5, we conclude that

$$\Delta_{u'}v^*(s'+e_i-\Sigma[A(s',i)]) = v^*(s'+e_i-\Sigma[A(s',i)]) - v^*(s'+e_i-\Sigma[A(s',i)]-u') > a_{u'}.$$

Substituting s' and u',

$$v^*(s+e_j+e_i-\Sigma[A(s+e_j,i)]) - v^*(s+e_j+e_i-\Sigma[A(s+e_j,i)]+u) > -a_u.$$

Further substituting v, this can be written as

$$v^*(s+e_i-\Sigma[A(s,i)]+v-u) - v^*(s+e_i-\Sigma[A(s,i)]+v) > -a_u,$$

and finally multiplying both sides by -1, we have

$$\Delta_u v^*(s+e_i-\Sigma[A(s,i)]+v) < a_u. \tag{13}$$

We will now use Lemma 5 with s and u. Notice that u is of the form $\Sigma[A]-\Sigma[B]$ where $B=(A(s,i)\setminus A(s+e_j,i))\cap K_j \subseteq A(s,i)$. We will show that $s+e_i-\Sigma[A(s,i)]-u \geq 0$. Using the same reasoning as before, if ℓ^{th} index of $s+e_i-\Sigma[A(s,i)]-u$ is negative, it must be the case that ℓ^{th} index of $\Sigma[(A(s+e_j,i)\setminus A(s,i))\cap K_j^c]$ and $\Sigma[(A(s,i)\setminus A(s+e_j,i))\cap K_j^c]$ are both non-zero. Again with the same reasoning, this can only happen when both ℓ and $\ell+1$ belong to these sets. However, both of these sets are subsets of K_j^c, which does not contain any two consecutive elements. Thus, we have that $s+e_i-\Sigma[A(s,i)]-u \geq 0$ and we can use Lemma 5 to conclude

$$\Delta_u v^*(s+e_i-\Sigma[A(s,i)]) > a_u.$$

Combining this with inequality (13) we obtain

$$\Delta_u v^*(s+e_i-A(s,i)) \;>\; a_u \;>\; \Delta_u v^*(s+e_i-A(s,i)+v).$$

Finally, we will use Lemma 4. Let $A=\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j^c$, $A'=\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j^c$, $B=\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j$, and $B'=\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j$, and using Lemma 4 we arrive at property $P(\sum_{k\in A}|w_k|-\sum_{k\in B}|w_k|, e_j+\sum_{k\in A'}|w_k|-\sum_{k\in B'}|w_k|)$ for v^*, where $A=\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j^c$, $A'=\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j^c$, $B=\big(A(s,i)\setminus A(s+e_j,i)\big)\cap K_j$, and $B'=\big(A(s+e_j,i)\setminus A(s,i)\big)\cap K_j$. Substituting the definitions of u and v, this property is equivalent to $P(u,v)$.

However, this property is equivalent to the inequality $\Delta_u v^*(s) \leq \Delta_u v^*(s+v)$ for any s. Letting $s' = s+e_i-A(s,i)$, this implies $\Delta_u v^*(s+e_i-A(s,i)) < \Delta_u v^*(s+e_i-A(s,i)+v)$ which contradicts the inequality we obtained above. Thus, we must have that $u=0$, which proves the desired inclusion relations. □

Appendix E Proof of Theorem 7

We first state and prove the following result:

Lemma 6. *Let v^* be the optimal value function, and s be a state. Then:*

i. *If $s_1 > 0$:*
   - *If $v^*(s) - v^*(s-e_1) \le 2$, then $v^*(s+e_1) - v^*(s) \le 2$.*
   - *If $v^*(s) - v^*(s-e_1) \le 2$, then $v^*(s+e_3) - v^*(s+e_3-e_1) \le 2$.*
   - *If $v^*(s+e_3) - v^*(s+e_3-e_1) \le 2$, then $v^*(s+e_1) - v^*(s) \le 2$.*

ii. *If $s_3 > 0$:*
   - *If $v^*(s) - v^*(s-e_3) \le 1$, then $v^*(s+e_3) - v^*(s) \le 1$.*
   - *If $v^*(s) - v^*(s-e_3) \le 1$, then $v^*(s+e_1) - v^*(s+e_1-e_3) \le 1$.*
   - *If $v^*(s+e_1) - v^*(s+e_1-e_2) \le 1$, then $v^*(s+e_3) - v^*(s) \le 1$.*

iii. *If $s_1 > 0$ and $s_3 > 0$:*
   - *If $v^*(s-e_3) - v^*(s-e_1) > 1$, then $v^*(s) - v^*(s+e_3-e_1) > 1$.*
   - *If $v^*(s+e_1-e_3) - v^*(s) > 1$, then $v^*(s-e_3) - v^*(s-e_1) > 1$.*

Proof of Lemma 6. All of these second-order properties are implied by Theorem 3: By *concavity*, $v^*(s+e_1)-v^*(s) \le v^*(s)-v^*(s-e_1) \le 2$, and by *superconcavity*, $v^*(s+e_1-e_3)-v^*(s-e_3) \le v^*(s)-v^*(s-e_1) \le 2$, and by *submodularity* in components 1 and 3, $v^*(s+e_3)-v^*(s+e_3-e_1) \le v^*(s)-v^*(s-e_1) \le 2$. This concludes Lemma 6(i). The proof of Lemma 6(ii) follows by the symmetric argument. The same argument using *superconcavity* in components 1 and 3 proves Lemma 6(iii). □

Proof of Theorem 7. By Lemma 6(i), we can let $t(s)$ to be the supremum of s' such that $v^*\big((s',0,s,0)\big) - v^*\big((s'-1,0,s,0)\big) > 2$. Such a supremum is finite since if we have infinitely many overdemanded pairs, an additional overdemanded pair has no future value and must be immediately used, and thus has a marginal value of is at most 2. By *concavity*, this function satisfies property *(i.1.)* and by *superconcavity* it satisfies property *(i.2.)* in Theorem 7. By a symmetric argument, properties *(ii.1.)*, and *(iii.1.)*, *(iii.2.)* in Theorem 7 follow from Lemma 6(ii) and 6(iii).

To show the optimality of the multi-dimensional threshold mechanism, we consider each case (we skip all the trivial decisions in all cases).

Case 1: The arriving pair is $\mathbb{O}_{A-B}$ type. In this case, the mechanism can decide to switch to state $s+e_1$, or stay in state s for an immediate surplus of 2. The former is chosen if and only if $v^*(s+e_1) > v^*(s)+2$. By property *(i.1.)*, this is equivalent to $s_1 + e_1 > t^{1,3}$.

Case 2: The arriving pair is $A-B$ type. In this case, the mechanism can decide to match this $A-B$ pair with an existing $\mathbb{O}_{A-B}$ pair, gain a surplus of 3 and transition to state $s-e_1$. Alternatively, it can decide to match this $A-B$ pair with an existing $B-A$ pair, gain a surplus of 2 and transition to state $s-e_3$. Thus, the decision depends on whether $v^*(s-e_3)+2 < v^*(s-e_1)+1$. By property *(ii.1.)*, this is equivalent to determining whether $s_3 < t^2(s_1)$.

Case 3: The arriving pair is $B-A$ type. Again, the mechanism can transition to states $s+e_3$ or $s+e_3-e_1$ to gain a surplus of 2. This means we check whether $v^*(s+e_3)-v^*(s+e_3-e_1) < 2$ or not. By property *(i.1.)*, this is equivalent to whether $s_1 \leq t^{1,3}(s_3+1)$ or not.

Case 4: The arriving pair is $\mathbb{O}_{B-A}$. The mechanism can transition to state $s-e_3$ and obtain a surplus of 3 or stay at state s to obtain a surplus of 3. Thus, we check whether $v^*(s)-v^*(s-e_3)$, which, by property *(iii.1.)*, is equivalent to checking the threshold t^4. $\square$

Online Appendix Not Intended for Print

I. About Proposition 3 in Ünver (2010)

Assumption 2 in Ünver (2010) regarding "overdemanded pairs are matched as soon as they arrive" is motivated by the following proposition.

Proposition 3 (Ünver, 2010). *If arrival probabilities p_{A-B} and p_{B-A} are sufficiently close to each other, then under any dynamically efficient multi-way matching mechanism, overdemanded type pairs are matched as soon as they arrive at the exchange pool.*

There is an analytical flaw in the proof. The proof of Proposition 3 works as follows: We consider two cases, the first where an overdemanded pair is pooled for at least one period, and the second where an overdemanded pair is matched immediately as they arrive. Then, we calculate the *upper bound* for the surplus of the first case, and the *lower bound* for the surplus of the second case. Then, we show that when $|p_{A-B} - p_{B-A}|$ is small enough, the lower bound is greater than the upper bound, thus, it is never optimal to accumulate overdemanded pairs.

The upper bound for the first case is calculated as

$$\frac{\lambda(p_{A-B} + p_{B-A})}{\lambda(p_{A-B} + p_{B-A}) + \rho}\left(3 + \frac{\lambda(p_{A-B} + p_{B-A})}{\rho}\right). \tag{14}$$

The coefficient on the left is the expected discounting until a reciprocal pair arrives; this is when a *three-way* exchange is conducted (thus, the term 3 inside the parenthesis). The second expression the parenthesis is the *upper bound* for the surplus of reciprocal pairs. It is an upper bound because it is argued that it assumes all incoming reciprocal pairs are matched as soon as they arrive.

The lower bound for the second case is calculated as

$$2 + 2\frac{\lambda(\min\{p_{B-A}, p_{A-B}\})}{\rho}. \tag{15}$$

The first item represents the surplus obtained by matching the overdemanded pair immediately with an underdemanded pair (thus, 2). The second term is argued to be a lower bound for the future surplus of all reciprocal pairs. The reasoning is as follows: If reciprocal pairs are matched exclusively with each other, the amount of exchanges would be bounded by the less frequently arriving type, and thus, this lower bound would be equal to $\sum 2(\delta \min\{p_{A-B}, p_{B-A}\})^k$.

We argue and claim that both of these bounds are not proper.

Expression (14) is not an upper bound. If there are sufficiently many $B-A$ pairs in the pool, then the upper bound for the future surplus would be $\frac{2\lambda(p_{A-B})}{\rho}$, which is bigger than $\frac{\lambda(p_{A-B}+p_{B-A})}{\rho}$. Thus, Expression (14) is an upper bound only if we assume that there are no pairs in the pool as the next reciprocal pair arrives.

Expression (15) is not a lower bound. Suppose, *without loss of generality*, that $B-A$ pairs arrive less frequently then $A-B$ pairs. Then, this approach assumes that whenever a $B-A$ pair arrives, there is an $A-B$ pair in the pool. But, this would be equivalent to the future surplus in the extreme case where there are infinitely many $A-B$ pairs at the beginning, which is not a lower bound assumption.[32]

To argue further, suppose $p_{A-B} = p_{B-A}$. Then, the upper bound for the future surplus ($\frac{\lambda(p_{A-B}+p_{B-A})}{\rho}$) would be equal to the lower bound for the future surplus ($2\frac{\lambda \min\{p_{B-A}, p_{A-B}\}}{\rho}$). This means that if arrival probabilities are equal to each other, then the future surplus is (exactly) equal to $\frac{2\lambda p_{A-B}}{\rho}$ and it is *independent from the state* of the Markov decision process (MDP). But, this is impossible, since the future surplus certainly depends on the state.

Remark 1. *We have argued so far that the proof of Proposition 3 is incorrect. But, this does not actually prove that Proposition 3 is incorrect. To prove that it is incorrect, we have to show that there exists some λ, ρ and arrival probabilities (that satisfy the other assumptions in Ünver, 2010) with $p_{A-B} = p_{B-A}$, such that in the optimal mechanism, it is optimal to accumulate at least one overdemanded pair type for at least one period. The numerical example in Figure 8 in Section IV of this Online Appendix proves this.*

II. Illustration of propagation and closedness: A two-dimensional example.

We illustrate the concept of *closedness* in *propagation* in a two-dimensional setting. Suppose we have a function $f : \mathbb{Z}^2 \to \mathbb{R}$ that is *concave* in each of the two components. Suppose we have two operators, T_1 defined by the equation $(T_1 f)(x) = \max\{f(x), f(x-e_1)+a_1\}$ and T_2 defined by the equation $(T_2 f)(x) = \max\{f(x), f(x-e_2)+a_2\}$. Suppose we want to show that *concavity* is *propagated* by the operators so that $T_1 f$ and $T_2 f$ are also *concave* in each component. Take the *concavity in* e_2 property and the operator $T_1 f$. We need to show for any $x \in \mathbb{N}^n$,

$$\begin{aligned}
&\max\{f(x), f(x-e_1)+a_1\} + \\
&\max\{f(x+2e_2), f(x+2e_2-e_1)+a_1\} \\
\le\ & 2\max\{f(x+e_2), f(x+e_2-e_1)+a_1\}
\end{aligned}$$

Applying the argument we developed above, symmetric cases easily follow, and the first case of asymmetry requires showing that for any $x \in \mathbb{N}^n$,

$$f(x) + f(x+2e_2-e_1) \le f(x+e_2) + f(x+e_2-e_1)$$

which does not generally hold, unless f is also *superconcave*. Thus, we are not able to propagate *concavity* alone, and the set of properties that consists of *concavity* in the two components are not closed for the operators T_1 and T_2.

Now suppose that, in addition to being *directionally concave*, f is also *superconcave*, and we want to show that $T_1 f$ and $T_2 f$ are also *superconcave*, in addition to being *concave*. Then, by the above

[32] This calculated bound would be closer to the lower bound if $|p_{A-B} - p_{B-A}|$ is large enough. In this case, it would be more likely that there are $A-B$ pairs in the pool whenever a $B-A$ pair arrives. But, the proof is only for the cases where $|p_{A-B} - p_{B-A}|$ is *small* enough.

argument, we are able to *propagate concavity*.[33] What remains to show is that $T_1 f$ and $T_2 f$ are also *superconcave*. By symmetry, it suffices to show that ij-*superconcavity* for components $i = 1$ and $j = 2$ is *propagated* by the two operators. T_1 propagating ij-*superconcavity* for components $i = 1$ and $j = 2$ is equivalent to for any $x \in \mathbb{N}^n$

$$\begin{aligned} &\max\{f(x+e_2), f(x+e_2-e_1)+a_1\} + \\ &\max\{f(x+2e_1), f(x+e_1)+a_1\}. \\ \leq\ &\max\{f(x+e_1), f(x)+a_1\} + \\ &\max\{f(x+e_1+e_2), f(x+e_2)+a_1\} \end{aligned}$$

Again, the symmetric cases easily follow. The first case of asymmetry also easily follows since by choosing the first and second arguments on the right-hand side, the two sides of the inequality become equal. Suppose we have the second case of asymmetry so that the left-hand side equals $f(x+e_2-e_1)+a_1+f(x+2e_1)$. Again choosing the first and second arguments respectively on the right-hand side, it suffices to show that $f(x+2e_1)-f(x+e_1) \leq f(x+e_2)-f(x+e_2-e_1)$. By *concavity*, $f(x+2e_1)-f(x+e_1) \leq f(x+e_1)-f(x)$ and by *superconcavity* $f(x+e_1)-f(x) \leq f(x+e_2)-f(x+e_2-e_1)$. By combining the two expressions, inequality follows.

To show $T_2 f$ is *superconcave*, we need to show for any $x \in \mathbb{N}^n$,

$$\begin{aligned} &\max\{f(x+e_2), f(x)+a_2\} + \\ &\max\{f(x+2e_1), f(x+2e_1-e_2)+a_2\}. \\ \leq\ &\max\{f(x+e_1), f(x+e_1-e_2)+a_2\} + \\ &\max\{f(x+e_1+e_2), f(x+e_1)+a_2\} \end{aligned}$$

Noting that symmetric cases follow, we first look at the first case of asymmetry, where the left-hand side equals $f(x+e_2)+f(x+2e_1-e_2)+a_2$. Choosing the first and second arguments on the right-hand side and letting $x' = x+e_1$, it suffices to show $f(x'+e_2-e_1)+f(x'+e_1-e_2) \leq 2f(x')$. By *superconcavity*, we have $f(x'+e_1)-f(x') \leq f(x'+e_2)-f(x'+e_2-e_1)$, and $f(x'+e_2)-f(x') \leq f(x'+e_1)-f(x'+e_1-e_2)$. Multiplying the first inequality by -1 and summing the inequalities, we obtain $f(x'+e_2-e_1)-f(x) \leq f(x)-f(x'+e_1-e_2)$, from which the desired inequality follows. In the second case of asymmetry, we can choose the first and second arguments to obtain the definition of *concavity in* e_1. This finishes the proof of our desired propagation results.

Notice that, even though we could not prove $T_1 f$ and $T_2 f$ are *concave* by just assuming that f is concave, we are able to prove $T_i f$ is *concave* and *superconcave* by assuming f is *concave* and *superconcave*. Thus, even though *concavity* did not *propagate* alone, it *propagates* together with *superconcavity*. This makes the set of properties {*concavity*, *superconcavity*} closed with respect to the set of operators $\{T_1, T_2\}$. With the same methodology, we can further show that an operator of the form $(T_{12}f)(x) = \max\{f(x), f(x-e_1-e_2)+a_{12}\}$ will also *propagate* the desired properties, if we further

[33] Second case of asymmetry similarly follows from *superconcavity*.

restrict the set of properties by including *supermodularity*. □

III. Numerical example for the unbalanced dynamic exchange

We provide a numerical example in Figure 7 for the unbalanced case. The figure depicts the set of *reachable* states for a *unbalanced* problem for $p_{A-B} = 0.0635$, $p_{B-A} = 0.0381$, $p_{\mathbb{O}_{A-B}} = 0.00626$, $p_{\mathbb{O}_{B-A}} = 0.021$, $\delta = 0.999$. The horizontal axis is $s_2 - s_3$ and the vertical axis is s_1. Note that $s_4 = 0$ for all reachable states.

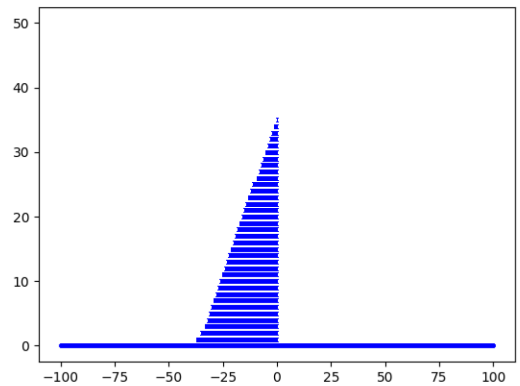

Figure 7: A numerical example

IV. Balanced dynamic exchange

We have described a special case of the optimal mechanism, the *unbalanced dynamic exchange*, in the main text. Here, we refer to the other case as a *balanced dynamic exchange*. Numerically, a problem is a **balanced dynamic exchange** when the arrival probabilities p_{A-B} and p_{B-A} are sufficiently close to each other. In a *balanced dynamic exchange* problem, the only restriction on the set of *reachable* states is given by Observation 5 (i) and Theorem 6. This latter theorem implies that the set of reachable states is connected in each dimension, inducing an interval structure.

However, intuitively, one would expect the set of *reachable* states to not contain a state s with $s_1 > 0$ and $s_2 > 0$, or $s_3 > 0$ and $s_4 > 0$, since existing $A-B$ and $\mathbb{O}_{A-B}$ pairs can make a *three-way* exchange. Moreover, if the mechanism implements a policy such that the $A-B$ pair is reserved for future incoming $B-A$ pair, then it is also expected intuitively that the $\mathbb{O}_{A-B}$ pair is not reserved, and matched immediately in a *two-way* exchange.

But, it turns out that this intuition is incorrect. In fact, numerically, we observe *reachable* states s with $s_1, s_2, s_4 > 0$ at the same time. We explore the intuition behind reserving $A-B$ pairs together with $\mathbb{O}_{A-B}$ pairs for future exchanges. First, we analyze the optimal mechanism for this most general case. The state space is in general four dimensional. But, by Observation 5 (i) and letting $s' = s_2 - s_3$ to denote the *x-axis* as in Section 7.1, we obtain a three dimensional set of *reachable* states. We first illustrate the numerically calculated set of *reachable* states for certain parameter values (see Figure 8).[34]

[34] As always, any state s with $s_1 = s_4 = 0$ is reachable, so we have the linear set of points at the bottom.

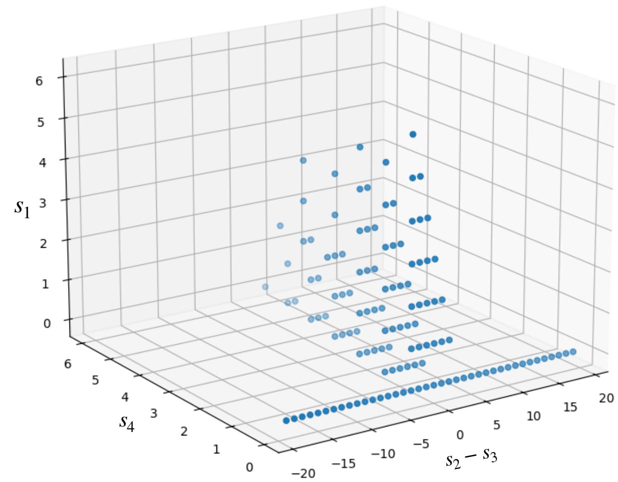

Figure 8: The set of *reachable* states for a *balanced* problem for $p_{A-B} = p_{B-A} = 0.0508$, $p_{\mathbb{O}_{A-B}} = p_{\mathbb{O}_{B-A}} = 0.001346$, $\delta = 0.999$. Note that the set of reachable states is symmetric with respect to the plane $s_2 - s_3 = 0$.

This graph points to the similarity between *unbalanced* and *balanced* problems. Recall that the set of *reachable* states in an *unbalanced* problem is a triangular region (Figure 5). Here, this is instead a polyhedral shape. The intuition for this shape is the same as before: The $\mathcal{D}^M$*-multimodularity* of the value function (Theorem 2) implies *substitutability* and *complementarity* relation between different types, which in turn implies trade-offs generating triangular regions of *reachable* states.

Another complication with the *balanced* problem is that the number of decisions is much bigger. For example, suppose that a $B-A$ pair arrives. Then, the mechanism can: *(i)* match the incoming pair with an existing $A-B$ pair, *(ii)* match the incoming pair with an $\mathbb{O}_{B-A}$ pair, *(iii)* match the incoming pair with an existing $A-B$ pair and match an existing $\mathbb{O}_{A-B}$ pair with an underdemanded pair, *(iv)* match the incoming pair with an existing $A-B$ pair and match an existing $\mathbb{O}_{B-A}$ pair with an underdemanded pair. Although we can eliminate some of these actions depending on the state, at least three actions remain for each state. This implies that the threshold mechanism for the *balanced* problem is a multiple thresholds mechanism, such that for t_1, t_2 with $t_1 < t_2$, the actions depend on whether $s < t_1$, $t_1 < s < t_2$, or $t_2 < s$. Moreover, thresholds themselves are multidimensional functions of the form $t : \mathbb{N}^2 \to \mathbb{N}$, and their properties as functions are again determined by $\mathcal{D}^M$*-multimodularity* and Theorems 4-6.

We now explain the only counter-intuitive result which the *balanced* problem entails, the fact that there are *reachable* states s with $s_1 > 0$ and $s_2 > 0$, i.e. it is possible for optimal mechanism to hold $\mathbb{O}_{A-B}$ pairs together with $A-B$ pairs. The reason this result is counter-intuitive is because, in Section 7.1, the optimal mechanism pools $\mathbb{O}_{A-B}$ to save future incoming $A-B$ pairs, as they will likely be excessive. Thus, why would the optimal mechanism keep reserving $\mathbb{O}_{A-B}$ pairs when an $A-B$ pair arrives? We illustrate the answer with an example.

Example 2. *Suppose that each arriving pair is an $A-B$ pair with $1/2$ probability and a $B-A$ pair with $1/2$ probability. Thus, we assume that no underdemanded or overdemanded pair arrives after the* *process starts. Further assume that we have an $\mathbb{O}_{A-B}$ pair at the start of the process (but no such pair will arrive later). Thus, we can (i) match the existing $\mathbb{O}_{A-B}$ pair with an underdemanded pair to gain a surplus of 2 at time 0, (ii) wait for an $A-B$ pair to arrive and match the existing $\mathbb{O}_{A-B}$ in a three-way exchange. Suppose we opted for the first option and gained a surplus of 2. Now suppose that next three arriving pairs are all $A-B$. Since the probabilities are equal, it would be expected that this excess of $A-B$ pairs persists for some time. Thus, if we waited for a couple of periods, and matched the $\mathbb{O}_{A-B}$ pair in a three-way exchange when the excess occurs, we could have saved an $A-B$ pair for waiting an excessive period of time. Now suppose that we decided to wait for an incoming $A-B$ pair, and matched the $\mathbb{O}_{A-B}$ pair with the first $A-B$ pair. Now suppose that after we do this, three $B-A$ pairs arrived in a row. Same as before, this imbalance will imply a loss of surplus, since excess $B-A$ pairs will likely wait for long. But if we did not match the $\mathbb{O}_{A-B}$ pair with an $A-B$ pair, we could have used this $A-B$ pair to match one of the excess $B-A$ pairs, and matched the $\mathbb{O}_{A-B}$ pair with an underdemanded pair, and used one excess pair without losing any surplus.*

We see that, both strategies fail to achieve the maximum surplus for some sequence of incoming pairs. Now consider the following strategy: We set some threshold for the imbalance, say, three excess pairs of $A-B$ and two excess pairs of $B-A$, and the $\mathbb{O}_{A-B}$ pairs are reserved, until one of these forms of imbalances occur. If the first case of imbalance occurs, the $\mathbb{O}_{A-B}$ pair is matched with one of the $A-B$ pairs. If the second case of imbalance occurs, the $\mathbb{O}_{A-B}$ pair is matched in a two-way exchange. This strategy is not dominated by either of the strategies we have discussed above, it is better than at least one of these strategies for some sequence of incoming pairs. This is precisely the strategy we numerically observe (although the exact numbers for thresholds depend on the problem).

Optimal mechanism uses overdemanded pairs as a *buffer* for future imbalances. It pools the overdemanded pairs of both types up to certain numbers,[35] until the reciprocal pairs reach a state of imbalance, in which case, it uses one of the overdemanded pairs it has been reserving in the pool to *mitigate* this imbalance. Thus, we call the set of reachable states we observe in the numerically calculated figure above as the *buffer zone*.

Although the optimal mechanism for a *balanced* problem has a more complicated set of decisions and states, it is closer to being a greedy algorithm (that always conducts the exchange of maximum size) than the one for an *unbalanced* problem. This is because, when arrival probabilities are close to each other, the case of future imbalance of reciprocal pairs is less likely. Thus, the incentives to not conduct the maximal exchange diminish. For this reason, for a *balanced* problem, always conducting the maximal exchange provides a good approximation to the optimal mechanism.

[35] However, it does so in a much smaller scale than an *unbalanced* mechanism.