Pool Management in Kidney Exchange Programs

Rafael Azevedo Alves

Master in Network and Information Systems Engineering

Department of Computer Science

Faculty of Sciences of the University of Porto

2026

Pool Management in Kidney Exchange Programs

Rafael Azevedo Alves

Dissertation carried out as part of the Master in Network and Information Systems Engineering

Department of Computer Science

2026

Supervisor

Prof. Dr. João Pedro Pedroso, Associate Professor, Faculty of Sciences of the University of Porto

Co-supervisor

Prof. Dr. Ana Viana, Coordinator Professor, ISEP - School of Engineering, Polytechnic of Porto

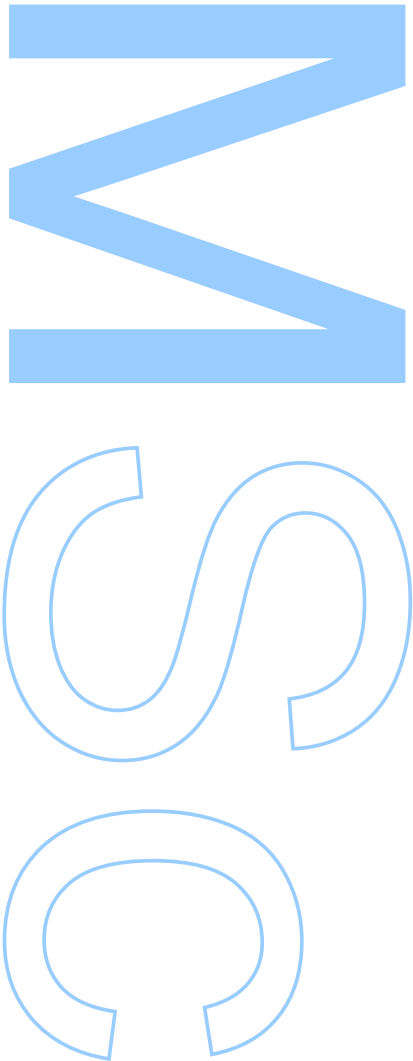

Acknowledgements

My deepest gratitude goes to my supervisor, Prof. João Pedro Pedroso, and my co-supervisor, Prof. Ana Viana. Their continuous availability, support, and guidance, along with their precious feedback, were instrumental in bringing this thesis to a successful conclusion.

I would also like to extend my thanks to Tiago Monteiro for sharing his research data, and to Xenia Klimentova for walking me through the technical groundwork necessary to begin my own implementation. Their initial push provided the crucial momentum needed in the start of this dissertation.

Lastly, I must thank my family and friends for their unwavering support throughout my university years. This achievement would simply not have been the same without them.

Resumo

Os Programas de Doação Renal Cruzada (PDRC) permitem que pares incompatíveis recebam transplantes através de trocas coordenadas de doadores. Apesar dos avanços recentes, a literatura carece de simulações dinâmicas de longa duração que integrem a decisão de aceitação de imunossupressão dos pares parcialmente compatíveis e de análises comparativas de estratégias de priorização baseadas na compatibilidade e no tempo de espera, num contexto de incerteza quanto ao comportamento de aceitação de imunossupressão. Esta dissertação aborda estas lacunas desenvolvendo uma simulação de eventos discretos de um PDRC em horizontes de 6, 12 e 24 anos. Avaliamos quatro políticas de correspondência no espectro eficiência-equidade — duas políticas hierárquicas lexicográficas com priorização baseada na compatibilidade, uma com um critério de desempate por tempo de espera, e outra sem, uma política de maximização de cardinalidade, e uma assente num modelo estático (uma única tentativa de emparelhamento dos pares). A aceitação de imunossupressão é modelada como uma decisão estocástica independente, com taxas de aceitação que variam entre 0% e 100%, ao longo de um total de 5500 simulações.

Os resultados estabelecem cinco conclusões principais: a) a priorização hierárquica melhora substancialmente o número de transplantes de pares incompatíveis, superando o modelo estático; b) o critério de desempate por tempo de espera proporciona proteção adicional, especialmente para os casos com maior tempo de espera; c) a preocupação teórica de que a saída dos pares parcialmente compatíveis para imunossupressão privaria os pares incompatíveis de dadores cruciais não se concretiza: o número de transplantes mantém-se estável em todos os cenários de aceitação de imunossupressão; d) contabilizando conjuntamente transplantes via PDRC e imunossupressão, as políticas de equidade superam a de cardinalidade máxima; e) todas as políticas dinâmicas são computacionalmente viáveis com tempos médios de resolução inferiores a 0,09 segundos por *match run*.

Estas descobertas sustentam uma recomendação clara: adotar políticas de priorização hierárquica baseadas em compatibilidade com desempate por tempo de espera, e encarar o transplante por imunossupressão como uma via complementar para os pares parcialmente compatíveis. Desta forma, potencia-se a que todos os pacientes num PDRC tenham uma oportunidade equitativa de receber um transplante que salva vidas.

Palavras-chave: Programas de Doação Renal Cruzada, Programação Linear Inteira, Otimização, Simulação, Estratégias de Priorização, Imunossupressão, Equidade na Saúde

Abstract

Kidney Exchange Programs (KEPs) allow incompatible patient-donor pairs to receive living-donor transplants through coordinated donor swaps. Despite recent progress, KEP literature lacks long-horizon dynamic simulations embedding the immunosuppressive acceptance decision of half-compatible pairs and a comparative analysis of compatibility-based and waiting-time-based prioritization strategies under immunosuppressive behavioral uncertainty. This thesis addresses these gaps by developing a discrete-event of KEP's operation over horizons of 6, 12, and 24 years. We design and implement four matching policies covering the efficiency-equity spectrum - two hierarchical lexicographic policies with compatibility-based prioritization, with and without a waiting-time tiebreaker, a maximum cardinality efficiency benchmark, and a static upper-bound model - while modeling immunosuppressive acceptance as an independent stochastic decision with acceptance rates varying from 0% to 100% across 5,500 simulation runs.

The results establish five main conclusions. Hierarchical prioritization consistently delivers substantially better outcomes for incompatible pairs, even surpassing the static upper bound. The waiting-time tiebreaker provides additional protection, especially for the longest-waiting cases. The theoretical concern that the withdrawal of half-compatible pairs to pursue direct immunosuppressive transplantation would deprive incompatible pairs of crucial donors does not materialize: the number of transplanted incompatible pairs remain stable across all immunosuppression acceptance rate scenarios. When immunosuppressive and KEP exchange transplants are counted together, equity-constrained policies achieve superior system-level efficiency compared to maximum cardinality. Finally, all dynamic policies remain computationally tractable with mean solve times below 0.09 seconds per match run.

These findings support a clear recommendation: adopt compatibility-based hierarchical prioritization policies with a waiting-time tiebreaker, and embrace immunosuppressive transplantation as a complementary pathway for half-compatible pairs. This ensures that every patient who enters a KEP has an equitable chance to receive a life-saving transplant.

Keywords: Kidney Exchange Programs, Integer Linear Programming, Optimization, Simulation, Prioritization Strategies, Immunosuppression, Healthcare Equity

Table of Contents

List of Tables

List of Figures

List of Abbreviations

ABO A, B, and O Blood Group System. 1, 11– 14, 19, 39

ABOc ABO-Compatible. 12, 22

ABOi ABO-Incompatible. 8, 9, 22, 23

AMK All pairs Matched in the KEP. ix, 40, 47– 50, 52, 53, 55– 59, 61, 68, 69

ATO All pairs Transplanted Overall. ix, 40, 47, 49, 50, 70

DFS Depth-First Search. 17, 34– 36

DSA Donor-Specific Antibodies. 14, 72

ESRD End-Stage Renal Disease. 1

HDTI Half-compatible pairs Directly Transplanted with Immunosuppression. ix, 40, 50

HLA Human Leukocyte Antigen. v, 1, 8, 10– 13, 15, 23, 39, 72, 73

HMK Half-compatible pairs Matched in the KEP. ix, 40, 48, 50

ILP Integer Linear Programming. v, 4, 16– 18, 21, 23, 24, 33, 34, 42, 52, 66, 67, 73, 75

IMK Incompatible pairs Matched in the KEP. ix, x, 40, 52– 61, 63, 64, 68– 70, 85– 87

KEP Kidney Exchange Program. v, vi, ix, 1– 8, 10– 22, 24– 26, 28, 33, 34, 37, 38, 40, 41, 43, 46, 47, 49, 50, 60, 62, 63, 67, 68, 70– 75

NDD Non-Directed Donor. ix, 1, 6, 15, 17, 28– 30, 33– 39, 61– 63

NEAD Non-simultaneous Extended Altruistic Donor. 15, 16

PICEF Position-Indexed Chain-Edge Formulation. 16, 23

PRA Panel Reactive Antibody. ix, 2, 8, 13, 19, 25, 38, 40, 52, 54, 55, 72

1. Introduction

Chronic kidney disease is a major public health problem, affecting millions of individuals worldwide and placing a significant burden on health care systems. For patients suffering from End-Stage Renal Disease (ESRD), kidney transplantation is universally recognized as the most effective treatment compared to dialysis, offering a better quality of life and higher survival rates [1]. However, every year, thousands of individuals with this disease face a critical lack of suitable kidney donors, with a severe disparity between the demand for and the supply of organs from deceased and living donors.

Living donor kidney transplantation leads to better clinical outcomes for recipients when compared to deceased donation. However, a significant proportion of willing patient-donor pairs are unable to proceed with transplantation due to medical incompatibility such as blood type (A, B, and O Blood Group System (ABO)) incompatibility or tissue incompatibility (Human Leukocyte Antigen (HLA) incompatibility/positive crossmatch).

To solve this problem, Kidney Exchange Programs (KEPs) were established. In a KEP, patients who are incompatible with their intended donors can swap donors with other pairs, so that each patient receives a compatible kidney. Figure 1.1 illustrates the two fundamental exchange structures in a KEP: a **cycle** involves two or more pairs exchanging donors in a closed loop—for example, the donor of Pair 1 donates to the patient of Pair 2, while the donor of Pair 2 donates to the patient of Pair 1; a **chain** is initiated by a Non-Directed Donor (NDD), who has no intended recipient, donating to the patient of Pair 1 whose donor then donates to the patient of Pair 2, and so on. The donor in the last pair of the chain either donates to a patient in the deceased donors waiting list, or remains in the pool and initiates a chain in the following matching run.

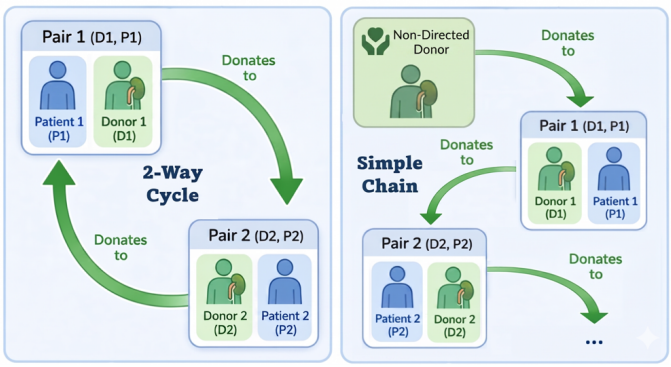

Figure 1.1: Illustration of a 2-way cycle (left) and a simple chain initiated by a NDD (right) in a KEP.

At its core, the matching of pairs in a KEP is a combinatorial optimization problem: given a pool of patient-donor pairs, represented as vertices in a directed compatibility graph where arcs represent which cross-pair transplants are medically feasible, the task is to identify the set of feasible cycles and chains that maximize a certain objective—typically the total number of transplants. Spanning operations research, market design, and clinical medicine, this problem has attracted substantial academic attention since its formal introduction by Roth et al. [2]. In the two decades since, KEPs have been successfully implemented across multiple countries and have significantly increased the number of living donor kidney transplants [3].

1.1. Kidney Exchange Programs' Challenges

Despite their strong theoretical foundation, the real-world implementation of KEPs faces a number of challenges that motivate the present work.

- **Dynamic Pool Management:** Real-world KEPs do not operate on a static pool of pairs. Pairs arrive continuously and depart the pool after transplantation or due to other external circumstances (death, clinical decisions, alternative treatments found). As a result, matching decisions must be made repeatedly under uncertainty concerning future arrivals. As highlighted by Biró et al. [3], different countries adopt distinct prioritization rules and operational strategies, underscoring the importance of simulation-based approaches to evaluate alternative matching policies.

- **Efficiency and Fairness:** Not all patients face equal matching chances, making the trade-off between efficiency and fairness a central challenge in KEP design. Most of the traditional optimization models for KEPs focus on maximizing cardinality, trying to obtain the highest number of transplants possible at a given time [3]. While this approach is efficient at the system level, it raises significant ethical and practical concerns, particularly regarding fairness, because it tends to overlook pairs that are hard-to-match—pairs with blood type O patients [4] or with patients exhibiting high Panel Reactive Antibody (PRA) levels [5, 6]—, which makes these pairs often experience prolonged waiting times and reduced chances of transplantation, often ending up leaving the program due to severe clinical deterioration or death before a match is secured. Quantifying this trade-off in a realistic dynamic setting is essential for policy decisions in operational KEPs.

- **Inclusion of Half-compatible Pairs:** Another relevant challenge concerns half-compatible pairs, for whom transplantation is medically possible but requires intensified immunosuppressive treatment [7], forcing patients who opt for this alternative to commit to a strict, lifelong routine of daily medications that prevent organ rejection. This creates a behavioral dynamic with important consequences for the pool: half-compatible pairs face a strategic decision between remaining in the exchange pool in the hope of a fully compatible match, or accepting direct immunosuppressive transplantation. The early departure of these pairs may remove highly desired donors from the compatibility graph, changing the long-term matching possibilities for the remaining, strictly incompatible pairs, which makes accurately modeling this behavior essential for realistic simulation of KEP outcomes.

1.2. Objectives

The main objective of this thesis is to evaluate the performance of KEPs under various priority rules and different immunosuppressive acceptance rates. To this end, we employ a discrete-event simulation-based framework that enables the systematic analysis of different matching policies under realistic KEP conditions—capturing the dynamic arrival and departure of patient-donor pairs, along with the periodic execution of the matching algorithm at discrete time intervals.

To achieve this goal, the following specific objectives are defined:

- Develop a dynamic simulator framework that models the stochastic arrival of patient-donor pairs and the evolution of the exchange pool over long time horizons of 6, 12, and 24 years.
- Design and evaluate different matching policies that balance efficiency and fairness: two hierarchical lexicographic policies, a maximum cardinality policy, and a static policy. The latter considers all pairs entering the pool simultaneously within a single optimization problem, thereby yielding an expected upper bound on the number of possible transplants.
- Model the immunosuppressive acceptance decision of half-compatible pairs as a stochastic process with varying acceptance rates, and quantify how it affects pool composition and matching outcomes over time.
- Assess the efficiency-equity trade-off across all policies' configurations, with particular focus on the outcomes for the hardest-to-match subpopulations.

1.3. Methodology and Contributions

To achieve the objectives of Section 1.2, this thesis makes the following methodological and scientific contributions:

- **Dynamic Simulation Framework:** We implement a discrete-event simulation of KEP operation, adapted from the methodology of Monteiro et al. [8], with extended time horizons of up to 24 years and a clear immunosuppressive decision-making mechanism for half-compatible pairs. The simulation framework models pair arrivals via a Poisson process, generates compatibility graphs using distributional assumptions with a clinical foundation, and executes match runs at 90-day intervals using Integer Linear Programming (ILP).
- **Hierarchical Policies:** We design two novel hierarchical policies that divide the pool into compatibility-based priority levels and solve a sequence of ILP problems lexicographically. The first policy incorporates accumulated waiting time as a tiebreaker within each compatibility level, while the second applies a pure compatibility-based approach without temporal consideration. This design assesses whether transplantation opportunities for the most vulnerable pairs are maximized and directly addresses the fairness limitations of standard maximum cardinality matching.
- **Immunosuppression Integration:** We incorporate the immunosuppressive acceptance decision into the dynamic simulation, modeling half-compatible pair behavior, across five acceptance rate scenarios ranging from 0% to 100%, adapted from Andersson and Kratz [7]. This allows us to quantify the long-term pool impact when half-compatible pairs withdraw from the pool to pursue direct immunosuppressive transplantation.
- **Comprehensive Empirical Evaluation:** We conduct a total of 5,500 simulation runs across a full factorial experimental design covering four policies, five immunosuppression acceptance rates, and three time horizons. This generates robust statistical evidence on the key factors under study.

1.4. Thesis Outline

The remainder of this thesis is organized as follows:

- **Chapter 2 - Literature Review** provides a comprehensive review of the academic literature supporting this work, covering the historical development and global adoption of KEP, the mathematical optimization approaches used to solve the matching problem, dynamic and static simulation frameworks, fairness and equity considerations, and the emerging literature on immunosuppressive therapy in KEP.

- **Chapter 3 - Methodology** describes the formal mathematical framework of the implemented KEP policies, the four matching policies, the dynamic and static simulation designs, the data generation process, the performance metrics evaluated, and the experimental design and computational implementation of this work.

- **Chapter 4 - Results and Analysis** presents and analyzes the experimental findings, examining pool dynamics, transplantation outcomes, the efficiency-equity trade-off, timeout behavior, waiting time distributions, computational performance, and how the policies perform in different temporal scenarios.

- **Chapter 5 - Conclusion** summarizes the main contributions and results of this thesis including their significance for the design of fair and efficient Kidney Exchange Programs, acknowledges the limitations of the work, and identifies directions for future research.

2. Literature Review

This chapter provides a comprehensive review of the literature supporting this thesis. We begin by tracing the historical development and global adoption of Kidney Exchange Programs, followed by an examination of the medical compatibility constraints that govern kidney transplantation. From there, we explore the mathematical optimization approaches and dynamic simulation frameworks developed to solve the KEP matching problem, before discussing fairness and equity considerations that motivate the hierarchical prioritization strategies implemented in this study. Subsequently, we review the emerging clinical and computational literature on immunosuppressive therapy, whose integration into KEP represents a central focus of this work. The chapter concludes by identifying research gaps that this thesis directly addresses.

2.1. Kidney Exchange Programs

2.1.1 Historical Development and Evolution

The concept of kidney exchange—in which two or more incompatible patient-donor pairs swap their donors so that each patient receives a compatible kidney—was first described in the medical literature during the 1980s. Rapaport [9] proposed the first formal framework for paired kidney donation, acknowledging that mutually incompatible pairs could potentially benefit one another through coordinated exchanges. Early implementations were limited and bilateral (involving exactly two pairs), due to the challenge of identifying compatible matches manually.

The first formalization of KEP as an optimization problem was proposed by Roth et al. [2], where the incompatible patient-donor pairs could be viewed as participants in a matching market with compatibility-based restrictions. Subsequent work in Roth et al. [4, 10] developed the core algorithmic and economic principles underlying modern KEPs, demonstrating that centralized systems relying on cycle-based matching could substantially increase transplant rates compared to decentralized bilateral exchanges. These early contributions established the now-standard computational framework of KEPs: represent the pool as a directed compatibility graph, identify feasible cycles (and later, chains initiated by NDDs), and solve the resulting problem as a maximum weighted matching optimization problem.

Since then, multiple operational KEPs worldwide adopted this framework [3, 11], including national programs in the United Kingdom, the Netherlands, Portugal, Spain, USA, and

Canada. At the transnational level, the Scandiatransplant Exchange Program (STEP) demonstrated that international pool merging substantially increases matching opportunities for the hardest-to-match patients [12]. This is further supported by recent multi-year simulation studies, such as Druzsin et al. [13], which confirmed that pooling isolated national registries into a single, massive joint pool yields significant gains in total number of transplants compared to fragmented national matching runs. The push for centralized national registries remains an active area of policy design, as seen in the efforts to establish a unified national KEP in Germany to prevent hospital fragmentation and maximize the effective pool size [14], which was only recently approved by the German government. Today, KEP represents one of the most successful applications of market design and operations research to a critical life-saving healthcare challenge.

2.1.2 National and Transnational KEP Models

Though the fundamental goal of all KEPs is to facilitate exchanges between incompatible pairs, their operational designs, matching policies, legal frameworks, and institutional contexts differ considerably [3]. Below we provide a non-exhaustive list of national and transnational KEPs, with a short description on the matching rules and logistics involved.

1. **United Kingdom**

   - **Program:** *United Kingdom Living Kidney Sharing Scheme* (UKLKSS) was established in 2007 and includes all four nations (England, Wales, Scotland, Ireland).
   - **Matching:** The program operates through a centralized algorithm, with match runs typically executed every three months. Crucially, rather than relying on a simple cardinality algorithm. The algorithm focuses on (i) maximizing the number of two-way exchanges, (ii) maximizing the total number of transplants, (iii) minimizing the number of three-way exchanges, (iv) maximize the number of cross arcs, and (v) maximize the sum of match quality scores, which includes factors such as sensitization level, accumulated waiting time, and donor age [3, 15, 16].
   - **Logistics:** Kidneys are transported from the donor's local center to the recipient's center [3].

2. **Netherlands**

   - **Program:** *Dutch National Kidney Exchange Program*, which was the first KEP to be implemented in Europe, was established in 2004.
   - **Matching:** Match runs occur typically every three months, with algorithms capable of finding cycles with sizes up to four. The algorithm focuses on (i) maximizing the number of transplants, and the tiebreakers include (ii) prioritize identical blood groups over simply compatible ones, (iii) matches for hard-to-match patients, (iv) minimize the cycles/chains lengths, (v) maximize the geographical spread of matches across participating transplant centers (for logistic reasons), and (vi) opt for matches for patients with the longest dialysis vintage [15]. More recently, a novel allocation algorithm (CIAT) was introduced that allows for ABO-Incompatible (ABOi) and "low-level" HLA-incompatible matching to maximize the chances of transplant for highly sensitized patients who have been on the wait-list for long periods [11, 15].
   - **Logistics:** Donors travel to the recipient's transplant center for the procedure [3].

3. **Portugal**

   - **Program:** *Programa Nacional de Doação Renal Cruzada* (PNDRC) was established in 2010. Internationally, it cooperates with countries such as Spain and Italy.
   - **Matching:** Match runs occur typically every three months, and the algorithm aims to (i) maximize the number of transplants in cycles (typically of size two or three) and chains, while (ii) prioritizing biological compatibility — such as the same blood type and PRA levels — and (iii) age [17].
   - **Logistics:** Kidneys are transported from the donor's local center to the recipient's center [3].

4. **Spain**

   - **Program:** *Programa Nacional de Donación Renal Cruzada* was established in 2009.
   - **Matching:** Match runs occur typically every four months. The matching process relies heavily on a "greedy heuristic" procedure based on assigned weights. This approach (i) maximizes the number of transplants achieved using two-way

cycles of maximum weight, and only when there are no more two-way cycles, it (ii) starts to choose three-way cycles of maximum weight [18].

- **Logistics:** Kidneys are transported from the donor's local center to the recipient's center [3].

5. **United States**

   - **Program:** The United States system is highly decentralized, consisting of three primary national programs: the *National Kidney Registry* (NKR, 2008), the *United Network for Organ Sharing* (UNOS, 2010) KPD, and the *Alliance for Paired Donation* (APD, 2006).
   - **Matching:** Match runs occur very frequently, typically on a weekly basis, and, in the UNOS's program, the matching process is based on an optimization algorithm that uses assigned weights to (i) maximize the total number of transplants by identifying all possible cycles (usually limited to a maximum length of three) and chains, and selecting the solution with the highest total priority score, while prioritizing factors such as the (ii) patient sensitization and (iii) pediatric status [19].
   - **Logistics:** In these programs, kidneys are transported from the donor's local center to the recipient's center, but donors also travel to the recipient's transplant center for the procedure [20].

6. **South Korea**

   - **Program:** South Korea was a global pioneer in paired kidney exchanges, having performed the first procedures in 1991. Today, its national system is managed by the *Korean Network for Organ Sharing* (KONOS).
   - **Matching:** The matching process in South Korea tries to maximize the total number of transplants. However, it has been disappearing from recent Korean practice, with the rapid increase of ABOi living-donor transplants (from 4.7% in 2009 to 26.3% in 2018) becoming a concerning clinical trend [21]. This decline is driven by the fact that many patients with incompatible donors undergo direct transplants using desensitization protocols with immunosuppression.
   - **Logistics:** Donors travel to the recipient's transplant center for the procedure [21].

7. **ScandiaTransplant Exchange Program (Transnational)**

   - **Program:** *Scandiatransplant Exchange Program* (STEP) was launched in 2019 and involves Finland, Sweden, Denmark, Norway, and Iceland [12].
   - **Matching:** Match runs occur every three or four months. The algorithm allows for two-way exchanges, three-way exchanges, and chains with no upper limit, and strictly prioritizes matches in the following order: (i) number of transplants, (ii) short cycles, (iii) low match probability pairs (determined by their calculated Panel-Reactive Antibody, or cPRA, a metric evaluating the recipient's immunological sensitization based on both current and historical HLA antibodies), (iv) compatible blood groups [12].
   - **Logistics:** Kidneys are transported from the donor's local center to the recipient's center [12].

8. **Kidney Exchange Program South Transplant Alliance (Transnational)**

   - **Program:** *Kidney Exchange Program South Transplant Alliance* (KEPSAT) was launched in 2017 and involves Portugal, Spain, and Italy [22].
   - **Matching:** Unlike the fully merged registry of STEP, KEPSAT operates as a sequential pool; it combines the unmatched patient-donor pairs from their respective national KEPs after their respective national match runs are completed. The matching algorithm's general principal is to prioritize for hardest-to-match patients [22].
   - **Logistics:** Kidneys are transported from the donor's local center to the recipient's center [22].

2.1.3 Key Success Metrics

Evaluating KEP performance requires a thoughtful analysis of which results are relevant and to whom. The literature has proposed various metrics, based on different stakeholder priorities:

- **Total Number of Transplants:** The baseline metric to evaluate the success of a KEP is the total number of transplants facilitated. This metric is consistent with utilitarian goals and is widely reported. However, it does not differentiate between standard matches and those involving hard-to-match patients.

- **Transplant Rate:** Normalizing the total number of transplants based on the size of the pool (the proportion of matched pairs) allows for comparison between programs of different sizes [12].

- **Number of Cycles and Chains:** Evaluating and comparing the structural breakdown of matches, such as the ratio of two-way cycles, three-way cycles, and the number and size of transplants facilitated by chains [23].

- **Transplants of Hard-to-match Patients:** Growing recognition of fairness in healthcare allocation led to the development of specific metrics for equity. Number of transplants based on patient characteristics (blood type O, high sensitization level, compatibility status) reveal whether certain groups are systematically disadvantaged [4–6].

- **Wait-list Attrition:** The number of pairs removed from the waiting list due to death or becoming too sick for surgery while waiting in the KEP pool [24].

- **Time-to-Transplant:** Average time from pool entry to transplantation measures program responsiveness. Shorter waiting times generally improve patient health outcomes. However, waiting time averages may hide substantial inequality: patients who are difficult to match may wait much longer than average [24].

- **Graft Survival:** Long-term transplant success (kidneys still functioning at 1-, 5-, and 10-year intervals) provides the measure of clinical benefit. However, this outcome is rarely used for real-time program evaluation due to long follow-up periods required [12, 25].

2.2. Compatibility in Kidney Transplantation

The medical feasibility of a kidney transplant depends on strict compatibility requirements between the donor and the patient [4]. In the mathematical modeling of KEPs, a potential transplant is represented as a directed arc (i, j) in a compatibility graph, originating from the donor of pair i and terminating at the patient of pair j. However, this arc can only exist if two primary prerequisites are satisfied: ABO compatibility and HLA compatibility, typically assessed through crossmatch testing.

2.2.1 Blood Type Compatibility

ABO compatibility is the fundamental immunological constraint determining whether a donor kidney can be successfully transplanted into a recipient. The ABO classifies individuals based on the antigens expressed on their red blood cells: type A (A antigen), type B (B antigen), type AB (both antigens), and type O (neither antigen). Because the immune system naturally produces antibodies against absent antigens, strict ABO compatibility rules for transplantation emerge, following the same logic as blood transfusions:

- **Type A:** Can donate to Type A or AB patients.
- **Type B:** Can donate to Type B or AB patients.
- **Type AB:** Can receive from any blood type (Universal Recipient), but can donate only to Type AB patients.
- **Type O:** Can donate to any blood type (Universal Donor), but can receive only from Type O donors because their immune system produces both anti-A and anti-B antibodies.

These biological asymmetries create natural structural patterns in KEP. Blood type O patients who need blood type O donors face the most severe matching disadvantage: they are strictly limited to receiving kidneys from blood type O donors; however, these same donors are universally compatible and frequently utilized by patients of all other blood types. On the other hand, blood type AB donors can only donate to the rarest blood type in the population (only blood type AB patients), which creates an excess of this type of donors. Roth et al. [4] demonstrated that blood type distribution is the primary determinant of feasible exchange rate, with blood type O patients continuously underserved in comparison to the rest of the population. This observation motivates the design of targeted prioritization policies that specifically address the inherent disadvantages faced by certain blood types.

2.2.2 HLA Compatibility and Crossmatching

Even if a donor and patient are ABO-Compatible (ABOc), a transplant may fail due to tissue incompatibility. The immune system identifies foreign tissues through HLA proteins, which can trigger immune responses when identified as non-self. Unlike ABO antibodies, normal individuals do not usually have HLA antibodies; they are developed if a patient

has been previously exposed to foreign HLA—such as through pregnancy, blood transfusions, or a previous organ transplant—, a process that can lead to sensitization and the development of pre-existing anti-HLA antibodies, which significantly reduce the likelihood of a compatible transplant. In KEP optimization models, together with ABO compatibility, HLA compatibility and crossmatch tests determine edge existence in the compatibility graph [26–28].

- **HLA Matching:** This involves identifying the specific HLA present in both donor and patient. This identification process is called " HLA typing". In kidney transplantation, the HLA antigens of greatest clinical importance fall into two categories: class I antigens— HLA-A, B, and C—and class II antigens— HLA-DRB1, DRB3/4/5, DQA1, DQB1, DPA1, and DPB1. Class I antigens are found on nearly all somatic cells, including endothelial cells and both B and T lymphocytes. Class II antigens, by contrast, have a more restricted distribution, being expressed only on B lymphocytes, activated T lymphocytes, and antigen-presenting cells (such as monocytes, macrophages, and dendritic cells) [28]. KEPs primarily focus on identifying unacceptable antigens—donor HLA molecules against which the patient has pre-existing antibodies, as these can trigger immune rejection.

- **Panel Reactive Antibody (PRA):** A metric used to quantify a patient's level of immunological sensitization. In modern clinical practice, this is often expressed as calculated PRA (cPRA). This approach consists of gathering a sample of donor HLA phenotypes and directly measuring, for each patient, the proportion of potential donors against whom they have pre-existing anti- HLA antibodies. Patients with PRA $\geq$ 80% are typically classified as highly sensitized (high level) and are incompatible with the majority of potential donors; the other PRA classifications are medium and low level [5, 27, 28].

- **Virtual Crossmatch Test (VXM):** A computational assessment of immunological compatibility that does not require mixing any physical blood samples or cells. It relies on analyzing two laboratory tests performed independently: the recipient's specific profile of unacceptable Human Leukocyte Antigen (HLA) and the potential donor's HLA molecular typing [29, 30]. Therefore, this is, in general, the definitive test considered by optimization models to draw the structural compatibility graph. Relying on virtual crossmatching to determine the existence of directed edges is

mathematically and practically necessary, as physical crossmatching every potential pair in a national program is labor-intensive, and logistically and financially challenging [28, 29]. A negative virtual crossmatch computationally predicts a high probability of a negative physical crossmatch, allowing algorithms to efficiently generate feasible cycles and chains [30].

- **Physical Crossmatch Test:** A laboratory test performed prior to transplantation to confirm immunological compatibility [29]. Because it requires to physically mix the recipient's serum with donor lymphocytes to detect Donor-Specific Antibodies (DSA), it is labor-intensive and logistically demanding to perform [29]. A positive physical crossmatch generally indicates the presence of DSA by showing that the recipient's antibodies react against donor cells (a cytotoxic reaction), contraindicating direct transplantation, whereas a negative physical crossmatch indicates no significant immunological reaction, confirming that the algorithmically determined transplant can safely proceed [27–29].

2.2.3 KEP Participant Compatibility Levels

KEP algorithms can evaluate pairs based on their clinical urgency and graph-theoretic utility. Therefore, operations research models typically classify the KEP pool into distinct compatibility levels based on the biological factors described in Section 2.2.1 and Section 2.2.2:

- **Incompatible Pairs:** These patient-donor pairs cannot undergo a direct transplant due to immunological barriers that cannot be overcome, such as ABO incompatibility or a positive crossmatch. Lacking any outside medical options, these pairs depend entirely on the exchange graph to find a viable organ [2]. Therefore, they represent the foundational demographic of any KEP.

- **Half-Compatible Pairs:** These pairs face an immunological barrier that can be overcome, such as low-level Donor-Specific Antibodies (DSA) or ABO incompatibility [7]. While a direct transplant is medically possible through intensive immunosuppressive desensitization therapy, they often enter the KEP to avoid the clinical risks and financial costs associated with such treatment [31]. Crucially for dynamic modeling, these pairs are not captive; if a KEP match is not found within a reasonable amount of time, they may voluntarily choose to leave the pool in order to pursue

a direct desensitized transplant. This behavior is critical to model in dynamic KEP simulations.

- **Fully Compatible Pairs:** Traditionally, medically compatible pairs undergo direct transplant and bypass the KEP entirely. However, modern programs increasingly encourage these pairs to participate in the exchange pool to secure a mathematically or clinically "better" match, such as a younger donor or a better HLA-antigen match to extend graft survival [10]. The inclusion of compatible pairs is vital for KEPs, as they add highly demanded donors to the pool without adding highly restricted patients, thereby increasing the overall density and efficiency of the compatibility graph [3].

- **Non-Directed Donors (NDD):** Individuals who volunteer to donate a kidney without a designated patient. In the compatibility graph, these donors have vertices with only outgoing arcs. They are mathematically crucial for initiating Non-simultaneous Extended Altruistic Donor (NEAD) chains [32].

By formally classifying participants into these classes, advanced algorithms can apply strict lexicographic priority rules—for example, ensuring that incompatible pairs (who have no other options outside KEP) receive allocation priority over half-compatible or fully compatible pairs.

2.3. Mathematical Optimization Approaches in KEP

2.3.1 Maximum Matching Formulations

At its core, KEP can be formulated as a maximum weighted matching problem on a directed compatibility graph G, $G = (V, A)$, where vertices V represent patient-donor pairs and directed arcs A represent compatibility relationships. The goal is to select a set of cycles and chains in this graph that maximizes some objective function, typically the total number of transplants, subject to the constraint that each pair participates in at most one selected cycle or chain.

This problem falls within the field of combinatorial optimization. In operational KEPs, all surgeries in a cycle must be performed simultaneously to prevent donors from backing out after their patient has received a kidney. With this additional structure, cycle lengths are strictly bounded to a small constant k (typically $k \leq 3$ or 4) to remain logistically feasible. For cycles of length at most $k = 2$, the maximum matching problem can be solved in

polynomial time using the "blossom algorithm" [33]. However, for cycles of length $k \geq 3$ and bounded, Abraham et al. [34] proved that finding the optimal maximum cardinality matching is computationally NP-hard. Despite this computational hardness, Roth et al. [4] proposed the near-universal adoption of $k = 3$ in modern practice. Their work established that restricting exchanges strictly to two-way cycles leaves a significant number of potential transplants unrealized, that allowing $k = 3$ yields considerable efficiency gains, and that allowing larger values of k yields only marginal additional benefits while the operational complexity increases substantially.

2.3.2 Integer Linear Programming (ILP) Formulations

To solve this NP-hard optimization problem, the Operations Research community relies heavily on Integer Linear Programming (ILP), offering the ability to incorporate complex constraints and objective functions while guaranteeing optimal solutions for realistic problem sizes. Some of the formulations presented in the literature are:

- **Cycle Formulation:** The standard cycle formulation [4, 34] introduces a binary variable $x_c \in \{0, 1\}$ for each feasible cycle c, where $x_c = 1$ indicates that the cycle is selected. The objective maximizes total transplants, and constraints ensure each pair appears in at most one selected cycle. This formulation has $O(|C|)$ variables and $O(|V|)$ constraints, where $|C|$ represents the number of feasible cycles.

- **Edge Formulation:** An alternative formulation uses binary variables for each directed edge rather than each cycle [4], introducing additional constraints to ensure selected edges form valid cycles. While the standard cycle formulation requires a quickly growing number of variables, this formulation bounds the number of variables to a polynomial size relative to the pool but requires an exponentially growing number of constraints to prevent paths longer than the maximum cycle length k. To resolve the exponential constraint issue, Constantino et al. [35] proposed new formulations, such as the edge-assignment and extended edge formulations, to strictly bound both the number of variables and constraints of the model to a polynomial size relative to the pool. Furthermore, it is formally known that the cycle formulation strictly dominates the edge formulation in terms of the linear relaxation bound [34].

- **Position-Indexed Formulation:** For KEPs that rely heavily on NEAD chains, Dickerson et al. [36] introduced the Position-Indexed Chain-Edge Formulation (PICEF),

which tracks the position of an edge within a chain, drastically improving computational scalability and reducing computation time for larger instances.

Recent algorithmic advances have further improved the scalability of KEPs under complex priority rules. Delorme et al. [16] developed a set of specialized algorithms and variable-fixing techniques (such as cycle and chain deactivation) designed to target hierarchical KEP optimization more efficiently. By systematically exploiting relaxations of these models, their algorithms achieve exact solutions for instances substantially larger than what prior methods could handle. More recently, Barkel et al. [37] provided a comprehensive state-of-the-art survey of these formulations, structurally classifying the various cycle and chain models and demonstrating empirically how modern ILP matching models navigate the immense computational scale of real-world KEP datasets.

2.3.3 Cycle and Chain Identification Algorithms

Before the ILP can be solved, the feasible exchanges within the compatibility graph must be identified. This identification process is a graph-theoretic enumeration problem that serves as a prerequisite for cycle and chain formulations.

- **Cycle Enumeration:** The standard approach for identifying feasible cycles is to use a Depth-First Search (DFS) starting from each vertex in the pool. By tracking paths and identifying cycles when a path returns to its starting vertex, the algorithm enumerates all cycles of length $2, 3, ..., k$. For $k \leq 3$, which is the most common practical constraint, this enumeration is efficient even for pools of several hundred pairs [38].

- **Chain Enumeration:** Chains initiated by NDD are structurally distinct from cycles as they are direct paths that do not return to the starting vertex. While some non-simultaneous systems allow for chains of arbitrary length, identifying these long chains requires highly complex algorithms, such as adaptations of the Traveling Salesperson Problem [39] and Branch-and-Price [40]. However, because extended chains carry a higher logistical risk of donor withdrawal, many national programs bound chain lengths. Consequently, in this work, chain lengths are bounded to a maximum length of three, which avoids the computational costs of long chains. Therefore, all feasible chains can be efficiently enumerated using a DFS algorithm originating from each NDD donor vertex.

Real KEPs impose additional constraints beyond cycle and chain lengths. To completely prevent donor withdrawal, the surgeries within a cycle must be simultaneous, which strictly bounds feasible cycle lengths due to logistical difficulties and medical center capacity limitations regarding how many exchanges a given hospital can support simultaneously [16, 38].

2.3.4 Static versus Dynamic Simulation Frameworks

While static ILP formulations successfully clear a single snapshot of the compatibility graph, real-world KEPs are inherently dynamic, characterized by the continuous stochastic arrival and departure of patient-donor pairs. This introduces a fundamental temporal tradeoff between immediate (greedy) matching to minimize waiting times versus delaying match runs to "thicken" the pool for more complex cycles — a dilemma theoretically formalized by Ünver [41] and later quantified regarding match frequency by Anderson et al. [42]. Recent simulation studies reinforce that expanding structural flexibility — such as incorporating three-way exchanges and non-directed donors — is critical not only for increasing the total number of transplants but for incentivizing full institutional participation and preventing hospital fragmentation in dynamic pools [14]. To evaluate allocation policies in this stochastic environment, researchers increasingly rely on discrete-event simulations, such as the 5-year evaluations of hierarchical rules by Druzsin et al. [13] and the 6- and 12-year dynamic simulations by Monteiro et al. [8], or, at the transnational level, the work of Klimentova et al. [43], dividing the timeline into discrete match runs, or "periods" (e.g. daily or monthly optimization runs). Building on this foundation, the present work adapts the dynamic simulation methodology of Monteiro et al. [8] to test new hierarchical policies over extended time horizons.

2.4. Fairness and Equity Considerations

The main objective of early KEPs was to maximize the total number of transplants. However, by strictly maximizing the total number of transplants, optimization models naturally favor patient-donor pairs with common blood types and low tissue sensitization. Because these are easily matched pairs, they present the highest number of feasible incoming and outgoing arcs so systems with purely cardinality-based objectives naturally build cycles and chains around them.

Consequently, equity is frequently directly sacrificed in order to maximize efficiency. Managing this inherent tension — between maximizing transplant efficiency and enforcing equity — remains a central challenge in KEP market design, requiring policies that explicitly address this efficiency-equity tradeoff [44]. This is formally defined as the "price of fairness", which is the cost, measured in total transplants lost, required to implement equity policies that prioritize matches for vulnerable groups of patients [5]. In order to manage this trade-off, modern KEPs must first identify the disadvantaged pairs, establish clinical policies to help them, and transform those policies into computational constraints.

2.4.1 Identifying Vulnerable Groups

A persistent concern in KEP design is that purely efficiency-oriented policies systematically disadvantage structurally hard-to-match groups, particularly blood type O patients [4] and highly sensitized patients [5, 6].

- **Blood Type O Patients:** Due to ABO compatibility rules (Section 2.2.1), blood type O patients can only receive kidneys from blood type O donors. Because blood type O patients are the hardest type to match, while blood type O donors are the easiest, a significant directional asymmetry is created in the compatibility graph. In order to complete longer cycles, policies that purely maximize the number of transplants often match blood type O donors with non-O patients, consequently retaining blood type O patients in the pool.
- **Highly Sensitized Patients:** Patients with a high PRA level possess highly reactive immune systems (Section 2.2.2), meaning that there is a very high probability that their body will reject the vast majority of potential donor kidneys. Therefore, their corresponding vertices in the compatibility graph have, on average, significantly fewer incoming edges when compared to patients with low PRA levels. Without policy measures, they accumulate substantially in KEP pools.

2.4.2 Waiting Time Prioritization

Since KEPs operate dynamically, meaning that patients who are not matched in a given match run continue in the pool to the next, unmatched pairs can wait years for a transplant. A key principle in medical ethics for organ allocation is preventing certain groups of patients from being consistently left without transplantation. A purely maximum-cardinality

algorithm ignores this temporal aspect completely, treating a patient who recently entered the pool identically to one who has been waiting several years. This situation must be avoided because it leaves hard-to-match patients waiting indefinitely while resources are allocated to newer, easy-to-match pairs [45].

Therefore, clinical policy requires that accumulated waiting time is taken into account in the kidney allocation process. National programs have formally established policies that mandate the prioritization of pairs based on their time in the system, utilizing point-based weights or tie-breakers to ensure that pairs who have been waiting longer are not consistently overlooked in favor of newly arrived pairs [3, 38]. Furthermore, Monteiro et al. [8] showed that waiting time in dynamic KEP simulations represents a critical fairness dimension, demonstrating that sequential priority algorithms can drastically reduce average and maximum waiting times for these disadvantaged pairs.

2.4.3 Compatibility-Based Prioritization

Beyond temporal equity, KEP policies must also address the underlying biological disparities present in the pool and discussed in Section 2.4.1. Even if a highly sensitized or blood type O patient has not been waiting long, their mathematical probability of finding a match in the near future remains statistically very low without targeted intervention. As a result, prioritizing waiting time alone is not enough for pairs whose vertices lack incoming arcs.

To address this, clinical governing bodies make a deliberate, ethical policy choice to absorb the “price of fairness” [5]. Modern KEP guidelines increasingly establish formal policies that guarantee that these hard-to-match pairs receive equitable consideration relative to easy-to-match pairs [46]. This policy consensus shifts the objective of KEP from simply maximizing the total number of transplants to ensuring a structural protection mechanism for the hardest-to-match pairs in the pool.

2.4.4 Algorithmic Mechanisms for Fairness

To operationalize the temporal and biological policies described, modern KEP matching models explicitly incorporate equity into their mathematical formulations. Instead of relying on manual triage, fairness is enforced through two primary mechanisms, Weighted Fairness and Lexicographic Fairness [5]:

- **Weighted Fairness:** Instead of treating every potential transplant as mathematically equal, matching models apply a re-weighting function to the edges in the compatibility graph. A multiplicative or additive re-weighting function dynamically scales the weight of an edge ending in a highly sensitized vertex, forcing the ILP formulation to value matches involving marginalized patients highly [46]. This approach intentionally trades off the total number of transplants in favor of maximizing a weighted objective function, accepting a reduction in efficiency to promote equitable allocation [44].

- **Lexicographic Fairness:** Advanced national programs move beyond basic weighting schemes, implementing multi-stage optimization that enforces lexicographic priorities across different classes of vertices. In this approach, the policy implemented sets a strict requirement that a certain group of pairs be included in the matching. If and only if this primary constraint is fulfilled, the policy considers a secondary objective (e.g. maximizing total number of transplants). This approach ensures that clinical equity is treated as the algorithm's principal focus.

Methodological Application in This Thesis:

To address these fairness and equity considerations, the present work implements two variations of a compatibility-based prioritization hierarchy: one relying strictly on compatibility levels, and a second that incorporates accumulated waiting time as a tiebreaker. Methodologically, both models adapt the lexicographic fairness structure proposed by Monteiro et al. [8] and Klimentova et al. [43]. By always maximizing transplants for the incompatible pairs with no options outside KEPs, and comparing the outcomes with and without the waiting-time tiebreaker, we can isolate how different allocation rules enforce patient-level equity. Crucially, this hierarchical approach reflects part of the operational reality of major European KEPs [3] and is supported by strong empirical evidence: Klaassen et al. [15] demonstrated that strict priority allocation algorithms for incompatible pairs increased their median transplant rates from 44% to 53% in the Dutch KEP, proving that targeted prioritization yields substantial gains for vulnerable groups without requiring pool expansion.

2.5. Immunosuppressive Therapy in KEP

2.5.1 Medical Protocols and Clinical Outcomes

The development of efficient immunosuppressive protocols has changed the world of organ transplantation. Clinically, this treatment, also known as “desensitization”, is divided into two distinct phases: pre-transplant desensitization procedures to overcome immunological barriers, and lifelong post-transplant medication to prevent organ rejection.

While roughly all transplant recipients need lifelong maintenance medication, the clinical pathway for blood/tissue-type incompatible transplants is considerably more burdensome. These patients must undergo intense desensitization, starting almost a month before the surgery can even take place — combining therapies like plasmapheresis to remove circulating antibodies and block the production of antibodies using agents like rituximab (anti-CD20) and intravenous immunoglobulin (IVIG) [21]. By adopting this technique, patients can receive a kidney from a donor with whom they were previously incompatible. Furthermore, following a desensitized transplant, patients usually need a more stringent and intensified lifelong maintenance regimen when compared to those receiving compatible kidneys. This heavier lifelong immunosuppression carries distinct clinical disadvantages, exposing the patient to higher long-term risks of serious infections, malignancies, and side effects from the medications.

Despite these heavy regimens, the clinical outcomes are highly encouraging. The long-term survival rates for ABO-Compatible (ABOc) and ABO-Incompatible (ABOi) transplants, in South Korea, are approximately similar, with the five-year survival rate of ABOc transplants being 96.3%, while that of ABOi transplants being 95.3% [21]. Furthermore, novel research makes use of the combinations of plasmapheresis/immunoadsorption, high-/low-dose of IVIG, and other medications that suppress the immune system, such as bortezomib and imlifidase (IdeS), to further improve the success rates of these types of transplants [21]. These outcomes confirm that incompatible transplants serve as a reliable, medically viable alternative when compatible exchange matching fails.

2.5.2 Formal Models for Kidney Exchange with Immunosuppressants

The formal integration of immunosuppressants into KEP optimization has generated a distinct and growing body of literature. Rather than treating desensitization purely as an independent clinical option, a more recent strand of literature introduces immunosuppressive therapy as an active and integrated component of the KEP system.

A core divergence in the modeling literature is whether immunosuppressants are treated as universal compatibility enablers or as compatibility enablers only specifically to certain patient-donor pairs. Heo et al. [47] proposed a matching mechanism using the Pairwise Cycles and Chains (PCC) model that achieves Pareto efficiency (ensuring no patient matching can be made better off without making another worse off) and monotonicity (guaranteeing that the availability of immunosuppressants does not disadvantage any patient who would have been successfully matched in a standard exchange), assuming that immunosuppressants can essentially make any donor compatible with any recipient. By using this wide perspective, they demonstrated that integrating immunosuppression treatment with exchange matching strictly dominates their independent use, saying that a group of patients requiring three suppressants for standalone desensitized transplants may achieve the same outcomes with only one or two when exchange is incorporated. Heo et al. [21] extend the previous work, documenting a concerning clinical trend in South Korea (addressed in Section 2.1.2) where there was a rapid increase of ABOi living-donor transplants (from 4.7% in 2009 to 26.3% in 2018), arguing that the observed elimination of kidney exchange cycles and chains by direct desensitized transplantation is suboptimal and that a hybrid policy combining exchange with selective suppressant use would achieve a number of transplants equivalent to relying strictly on ABOi procedures, while minimizing the medical and financial burden of the suppressants.

In contrast, Andersson and Kratz [7] proposed a more restricted and clinically realistic model: the introduction of *half-compatible* pairs. In this framework, immunosuppression is a pair-specific option available only to certain patient-donor pairs, the half-compatible pairs (e.g. ABOi but HLA-compatible with low antibody titers). For these pairs, direct transplantation is clinically feasible if the patient accepts an immunosuppressive protocol. Andersson and Kratz demonstrate, through analytical results and a static simulation, that integrating these half-compatible pairs into the pool—allowing them to participate in exchanges with the fallback option of desensitization treatment with their own donor—substantially increases the total number of transplants and particularly benefits blood-type O patients.

On the algorithmic side, incorporating these clinical options into ILP frameworks introduces significant computational complexity. Aziz et al. [48] presented a broad computational treatment for this problem by introducing the Kidney Exchange with Immunosuppressants (KEI) model. Their ILP formulation extends the PICEF model [36] by incorporating thick (fully compatible matches) and thin (half-compatible matches that require

immunosuppression) arcs, subject to a strict budget constraint on the number of allowable desensitization treatments. They demonstrated that significant efficiency gains in the number of matches are achievable even when a small hospital's budget for suppressants is available. Recently, Delorme et al. [23] advanced the algorithmic frontier for this problem by developing exact algorithms for two variants: KEP with Reserve Arcs (KEP-RA) and KEP with Half-Compatible Arcs (KEP-HCA). Both variants extend the classical formulation by allowing the selection of a limited number of medically incompatible transplants that can be made viable through specialized immunosuppressive protocols. The fundamental difference lies in their operational structure: while KEP-RA assumes that any incompatible arc can potentially be converted into a successful transplant, KEP-HCA restricts the algorithm's choice to a subset of arcs that have been deemed clinically viable [23]. By proving the structural property that any optimal KEP-RA solution contains at most one reserve arc per cycle—meaning that within any cycle, a maximum of one transplant will require medical intervention to overcome an initial incompatibility—, they designed efficient variable-reduction techniques and branch-and-price procedures capable of scaling to static instances of up to one thousand pairs, proving that the formulations are computationally tractable for realistic pool sizes.

More recently, Yang et al. [49] investigated the system-level impact of organ blood type conversion on KEP performance. By developing an ILP model and conducting a Monte Carlo dynamic simulation using Canadian Transplant Registry data, they showed that converting just 15% of non-O donors to the universal blood type O resulted in a 26% increase in transplant rates and a 2.3-month reduction in average waiting times. Critically, their model does not select donors for blood type conversion based on individual medical feasibility; instead, the ILP framework determines which non-O donors to convert, based solely on maximizing the overall objective of the optimization system, which is unrealistic because the study does not account for medical feasibility (such as the organ's biological response to the enzymatic treatment), which dictates whether the conversion is actually viable.

2.5.3 Patient Acceptance and Behavioral Modeling

Even though desensitized transplantation is clinically feasible and has shown survival benefits, patient acceptance of this procedure is not guaranteed. From both an economic and a behavioral perspective, the willingness of half-compatible pairs to accept this arrangement is determined by several interconnected factors:

- **Risk Perception and Quality of Life:** Patients must weigh the benefits of a quicker transplant against the significant lifetime costs associated with the procedure. Because desensitized transplants necessitate a more aggressive lifelong medication regimen than standard compatible transplants, patients have higher long-term risks of infections and severe drug-related side effects [50]. This physical and psychological burden, combined with the fact that patients often overestimate the risks of desensitized transplants relative to the mortality risks of continued dialysis, causes significant hesitancy when opting for the immunosuppression treatment [50].
- **Economic Barriers:** Desensitization protocols involve resource-intensive medical procedures that require high costs or specific insurance coverage, which can create a significant financial barrier to the individual patient [50].
- **Clinical Profile:** Individual characteristics, including the PRA levels [28] and the nature of the relationship between the recipient and the donor [51], influence the degree of acceptance rate variability.

To capture this complex behavioral environment, the simulation developed in this thesis models patient willingness dynamically, exploring acceptance rates from 0% to 100%. This comprehensive coverage allows for a methodical analysis of this specific withdrawal behavior, wherein half-compatible pairs willingly leave the exchange pool via immunosuppressive transplants when finding a transplant through KEP exchanges proves difficult for them.

2.5.4 Economic Considerations

While Section 2.5.3 highlights the economic hurdles facing individual patients, the economic argument from a healthcare system perspective is strongly in favor of expanding transplantation through desensitization.

Dialysis is one of the most common and resource-intensive treatments, costing approximately $80,000 to $90,000 per patient per year in the United States [52]. While desensitized transplants carry high initial costs—requiring additional pre-transplant protocols estimated to add approximately $28,000 to standard transplant costs—the procedure generates net savings of about 7.9% within three to four years by eliminating the need for dialysis [52]. From a system-level perspective, expanding KEP participation through the integration of immunosuppressive therapy reduces overall dialysis costs by eliminating

the need for this long-term treatment for patients that might wait years for a compatible transplant through the KEP matching system, improves patient quality of life, and frees critical clinical dialysis resources for patients who are completely ineligible for desensitized transplantation [52].

2.6. Research Gaps

The preceding review establishes a clear structural tension: while desensitized transplants yield excellent clinical outcomes and high economic benefits from a healthcare system perspective, actual patient acceptance of undergoing immunosuppressive treatment remains highly variable due to perceived risks, lifelong medication burdens, and economic barriers. Since this acceptance varies across individuals and over time, it is practically and mathematically necessary to model immunosuppression integration as a dynamic, probabilistic choice. However, the analysis of the existing literature reveals some important gaps in tackling this problem, which this thesis directly addresses.

The most significant gap in the literature is the limited body of research integrating the immunosuppressive therapy decision within a long-horizon dynamic KEP simulation that simultaneously accounts for patients' receptivity to immunosuppressive therapy and compares multiple prioritization strategies. Desensitization is an established clinical practice and its potential application in KEP has been theoretically acknowledged and established by the formal optimization models proposed to date [7, 21, 23, 47, 48], which operate primarily in static or single-period contexts. Consequently, due to the lack of a time dimension, these models cannot capture or quantify how the pool-level immunosuppressive treatment integration—whereby half-compatible pairs, after experiencing prolonged waiting time, leave the exchange pool and opt to undergo direct desensitized transplantation—affects the composition and efficiency of the exchange pool over time. Furthermore, previous studies cannot quantify how this patient attrition interacts with different matching and prioritization policies, and how they, specifically, impact the incompatible pairs: as half-compatible pairs exit the system with their donors, they strip the pool of critical nodes, which can drastically reduce the matching chances for the most vulnerable patients who have no alternative treatment options outside the KEP.

To address these gaps, this thesis develops and implements a comprehensive simulation framework that models dynamic KEP operation over time horizons of 6, 12, and 24 years, utilizing immunosuppressive treatment as an integrated decision variable for half-compatible pairs. The simulation framework tests behavioral uncertainty by varying the

immunosuppressive acceptance rates of half-compatible pairs after each unmatched period (ranging from 0% to 100%), extending the framework established by Andersson and Kratz [7], while implementing and comparing four distinct allocation policies ranging from pure maximum cardinality efficiency to compatibility-and-waiting-time-based equity models that adapt the lexicographic fairness structure proposed by Monteiro et al. [8].

3. Methodology

This chapter explains the formal framework, the design of the simulation, and the computational implementation underlying the empirical analysis of this thesis. The methodology is anchored by the mathematical formulation of the KEP as a directed compatibility graph, establishing the notation, and base formulations and constraints used throughout the work. Building upon this foundation, we introduce four distinct matching policies, positioning them on an efficiency–equity spectrum based on their prioritization rules. We then describe the dynamic (optimization at each matching period) and static (a single match run execution) simulation frameworks, which include the discrete-event mechanism that controls pair arrivals and departures, match run execution, and the immunosuppression decisions faced by half-compatible pairs. To populate these simulations, we detail a synthetic data generation process, including the distributional assumptions and immunological criteria used to construct realistic patient-donor pairs pools. Finally, the chapter defines the performance metrics used to evaluate outcomes, concluding with the three-dimensional factorial experimental design that tests how matching policies, immunosuppression acceptance rates, and time horizons affect each other.

3.1. Problem Definition and Mathematical Framework

Before analyzing the long-term dynamics of the simulated KEP, it is necessary to establish the structural mathematical framework that governs the matching process. The operational core of a KEP is an optimization problem, traditionally modeled as a matching problem on a directed compatibility graph. The objective of the matching engine is to identify a set of cycles and chains that maximizes an objective. The model we present next was adapted from Monteiro et al. [8].

3.1.1 Compatibility Graph Model Formulation

We model the KEP as a directed compatibility graph $G = (V, A)$, where $V = P \cup N$ is the set of vertices representing patient-donor pairs, P, and Non-Directed Donor (NDD), N, and A represents the compatibility arcs between vertices: $(i, j) \in A$ if the donor in vertex $i \in V$ and the patient in vertex $j \in V$ are compatible.

The length of any cycle or chain is limited by a maximum value k, defined as the maximum number of vertices involved.

Let $\mathcal{C}$ be the set of all feasible cycles and chains with at most k vertices in G, and let $V(c)$ denote the set of vertices that belong to cycle or chain c. Each valid cycles or chains $c \in \mathcal{C}$ is associated to a variable x_c such that:

$$x_c = \begin{cases} 1 & \text{if cycle or chain } c \text{ is selected for the exchange,} \\ 0 & \text{otherwise.} \end{cases}$$

3.1.2 Base Priority Classes General Formulation

Let V be the set of vertices in graph G. We partition V into disjoint priority sets, $\mathcal{P}_i$, based on the objective of the model developed, and order them in a hierarchical sequence $\mathcal{H} = (\mathcal{P}_M, \mathcal{P}_{M-1}, \ldots, \mathcal{P}_1)$, with M being the total number of active priority levels and $\mathcal{P}_M$ having the highest priority.

Given $v \in V$, we sequentially solve the following sequence of problems for $z = M, M-1, \ldots, 1$:

$$v_z^* = \max \sum_{c \in \mathcal{C}(k)} w_c^z x_c \tag{3.1a}$$

$$\text{Subject to:} \quad \sum_{c: i \in V(c)} x_c \leq 1, \quad \forall i \in V \tag{3.1b}$$

$$\sum_{c \in \mathcal{C}} w_c^j x_c \geq v_j^*, \quad \forall j \in \{M, M-1, \ldots, z+1\}, z < M \tag{3.1c}$$

$$x_c \in \{0,1\}, \quad \forall c \in \mathcal{C}(k) \tag{3.1d}$$

The equations of the model can be described as follows:

Objective Function (3.1a): The objective is to maximize the number of transplants specifically for the pairs belonging to the current priority class z. The parameter v_z^* denotes the maximum number of transplants for this specific group. The coefficient w_c^z counts how many pairs in the cycle/chain c belong to priority class z; if the cycle/chain contains no pairs from this class, the weight is 0.

Capacity Constraints (3.1b): These constraints ensure that a pair cannot be selected more than once. For each pair or NDD i, the sum of all selected cycles/chains

containing i must be less than or equal to 1. In other words, vertex i participates in a single cycle/chain, or does not participate at all.

Hierarchical Constraints (3.1c): These constraints guarantee that the search for an optimal solution for the current priority level z does not compromise the results obtained in previous iterations (levels M to $z+1$). This ensures that the maximization of objectives in iteration z is achieved while strictly preserving the optimal objective values previously obtained in higher priority iterations.

Variable Domain (3.1d): This binary constraint forces the model to either accept a cycle/chain in its entirety ($x_c = 1$) or discard it completely ($x_c = 0$).

3.2. Matching Policies

We implemented four distinct matching policies that share the formulation described in Section 3.1.2, and differ solely in their objective functions and prioritization strategies.

3.2.1 Policy 1: Hierarchical with Waiting Time

This policy implements a lexicographic prioritization scheme that first prioritizes pairs based on their compatibility level (Incompatible, Half-Compatible, Compatible, NDD), and then by waiting time within each level. This reflects an equity-focused policy that prioritizes medical necessity (compatibility level of the pairs) while promoting fairness through waiting time consideration. As pairs remain in the pool longer, their priority gradually increases, with the aim of preventing indefinite waiting for hard-to-match pairs.

The prioritization criteria for this policy are defined lexicographically as follows:

- Let $u, v \in V$ be two vertices representing pairs. Let $L(v)$ denote the compatibility class of v, where:

$$L(\text{Incompatible}) > L(\text{Half-Compatible}) > L(\text{Compatible \& Non-Direct Donor})$$

- Let $W(v)$ denote the waiting time of v in periods of time. We say that pair u has higher priority than pair v ($u \succ v$) if and only if:

$$L(u) > L(v) \text{ or } (L(u) = L(v) \text{ and } W(u) > W(v))$$

Example of Policy 1: Consider a pool with four incompatible pairs. Pair 2 has a waiting time of 3 periods, while the remaining pairs have a waiting time of 0. The compatibility between pairs is illustrated in the directed compatibility graph shown in Figure 3.1.

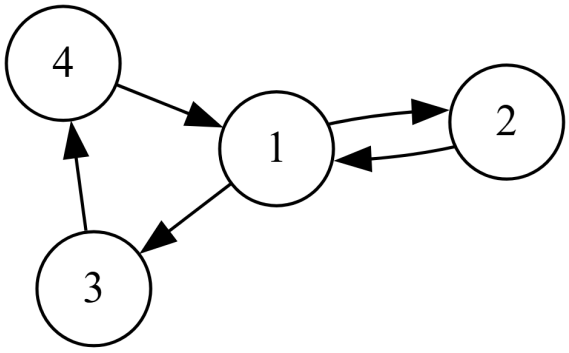

Figure 3.1: Directed compatibility graph with possible cycles for Example of Policy 1.

Since all pairs share the same compatibility class, their priority level is the same and the tie-breaking rule is based on waiting time. In this case, Pair 2 has priority, and the cycle involving this pair (Cycle 1 $\rightarrow$ 2 $\rightarrow$ 1) is selected. However, if Pair 2 had the same waiting time as the other pairs, they would all share the same priority; but, because its cycle is shorter than the alternative, the longer cycle would be preferred (Cycle 1 $\rightarrow$ 3 $\rightarrow$ 4 $\rightarrow$ 1). Finally, if Pair 2 was half-compatible, it would have lower priority than the incompatible pairs, resulting in the selection of the cycle with the incompatible pairs (Cycle 1 $\rightarrow$ 3 $\rightarrow$ 4 $\rightarrow$ 1).

3.2.2 Policy 2: Hierarchical without Waiting Time

This policy implements a lexicographic prioritization scheme that only prioritizes pairs based on their level of compatibility without waiting time consideration. This reflects a pure medical-necessity approach where all pairs within the same compatibility class are treated equally regardless of their time in the pool.

The prioritization criteria for this policy are defined lexicographically as follows:

- Let $u, v \in V$ be two vertices representing pairs. Let $L(v)$ denote the compatibility class of v, where:

$$L(\text{Incompatible}) > L(\text{Half-Compatible}) > L(\text{Compatible \& Non-Direct Donor})$$

We say that pair u has higher priority than pair v ($u \succ v$) if and only if:

$$L(u) > L(v)$$

By comparing Policy 2 with Policy 1, we isolate the impact of waiting time prioritization on outcomes, enabling assessment of whether temporal fairness considerations meaningfully affect transplant allocation.

Example of Policy 2: Consider the scenario described in the example of Policy 1.

Since all pairs share the same compatibility class and the tie-breaking rule based on waiting time is not considered in this policy, Pair 2 has no priority over the remaining pairs, and the cycle involving this pair (Cycle $1 \rightarrow 2 \rightarrow 1$) is not selected, because its cycle is shorter than the alternative, so the longer cycle would be preferred (Cycle $1 \rightarrow 3 \rightarrow 4 \rightarrow 1$). Finally, if Pair 3 and 4 were half-compatible, they would have lower priority than the incompatible pairs, resulting in the selection of the cycle with the incompatible pairs (Cycle $1 \rightarrow 2 \rightarrow 1$).

3.2.3 Policy 3: Maximum Cardinality

This policy implements an objective optimization approach that maximizes transplant volume without prioritization of level of compatibility or waiting time, serving as an efficiency benchmark approach.

This formulation treats all pairs equally and identifies the maximum cardinality matching. Consequently, no prioritization criteria are applied in this policy.

Example of Policy 3: Consider the scenario described in the example of Policy 1, with the only difference being that the Pair 3 and 4 are half-compatible.

If we were considering Policy 2 or Policy 1, Pair 3 and Pair 4 would have less priority when compared to the remaining incompatible pairs, so the cycle selected in both cases would be Cycle $1 \rightarrow 2 \rightarrow 1$, regardless of the waiting time of the pairs. However, for Policy 3 there are no prioritization constrains applied, so the policy selects the cycle that yields the higher number of transplanted pairs regardless of their compatibility level, thereby choosing the Cycle $1 \rightarrow 3 \rightarrow 4 \rightarrow 1$.

3.2.4 Policy 4: Static

This policy applies Policy 3: Maximum Cardinality's formulation to a static scenario (single optimization) considering all pairs present simultaneously, and the simulation consisting of a single matching. It acts as an upper-bound benchmark for exchange transplant

numbers, representing the theoretical maximum that could be attained if all pairs were simultaneously available for matching, without the limitations imposed by dynamic arrivals and departures. For this policy, the scenario and the resulting selection of cycles are equivalent to those detailed in the example of Policy 3.

3.3. Simulation Framework

We implement two simulation frameworks: dynamic and static. The dynamic simulator, used by our dynamic policies (Policy 1: Hierarchical with Waiting Time, Policy 2: Hierarchical without Waiting Time, and Policy 3: Maximum Cardinality), models KEP operation as a discrete-event process over time horizons of 6, 12, and 24 years, capturing the continuous arrival and departure of pairs. The static simulator, used by our static policy (Policy 4: Static), models the KEP as a single optimization problem in which all pairs are placed simultaneously in the pool.

3.3.1 Dynamic Simulator

We implement a discrete-event simulator of KEP operation over time horizons of 6, 12 and 24 years, divided into periods of 90 days (three months), with each period representing a match run, yielding 24, 48 and 96 periods, respectively. The simulation proceeds as follows:

Initialization ($t = 0$): The pool starts empty. A complete dataset is loaded from the pre-generated pool file, containing all patient-donor pairs and NDD that will arrive throughout the simulation horizon. The corresponding pre-computed compatibility arc file encoding all the compatibility relationships between pairs is also loaded.

For each period ($t = 1, 2, ..., T$):

- **Pair Arrival and Departure:** All pairs with arrival time less than or equal to t and greater than $t-1$ join the active pool $P(t)$, and all pairs $i \in P(t-1)$ whose departure time is less than t leave the active pool without being transplanted.
- **Match Run Execution:** The selected matching policy is executed on the current pool composition. The compatibility graph $G(t) = (V(t), A(t))$ is constructed, feasible cycles and chains enumerated, and the ILP solved using the Gurobi Optimizer [53].

- **Transplant Realization:** All pairs in the cycles and chains selected in the optimal solution of the previous step are removed from the active pool and recorded as successfully transplanted.

- **Immunosuppression Decision:** For each half-compatible pair $i \in P(t)$ that remains unmatched in period t, if the pair opts to undergo immunosuppression, it undergoes transplantation with its original donor and exits the pool. Otherwise, it is retained in the pool for the subsequent match run.

3.3.2 Static Simulator

We implement a simulation of KEP where the scenarios are simplified:

- All pairs are placed in the pool at $t = 0$ without acknowledging their arrival or departure times.

- A single match run is executed on the complete pool, utilizing the same procedures for compatibility graph construction, transplant realization, and immunosuppression decisions defined in the dynamic simulator framework (Section 3.3.1).

3.3.3 Cycle & Chain Enumeration

At each period t of the simulation, before the solver is executed, a directed compatibility graph $G(t)$ is constructed from the active pool $P(t)$ by reading the pre-computed arc file. Once $G(t)$ is available, all feasible cycles and chains of length at most k are enumerated prior to solving the ILP, producing the feasible set $\mathcal{C}$ described in Section 3.1. Feasible cycles and chains are identified using a Depth-First Search (DFS) approach. To ensure each cycle is discovered exactly once and to optimize search efficiency, the search is rooted at each non- NDD node v and explores only neighbors with an identifier greater than the root node ($u > v$), recording a cycle whenever the search returns to the root node. For chains, the search is initiated exclusively from NDD nodes, with every partial path of length up to $k-1$ edges being recorded as a feasible chain. NDD nodes participate exclusively as chain initiators and serve no other role in the matching graph.

Algorithms 1 and 2 formalize the cycle and chain enumeration procedures, respectively, with both procedures operating directly on the directed compatibility graph using a set of parameters.

Algorithm 1 receives as input: *CycleNodes*, the set of eligible patient-donor pairs; *NDDset*, the set of NDD nodes identifiers used to prevent NDD from appearing in cycle nodes; *adjList*, the adjacency list representing the directed edges of $G(t)$; *nv*, the total number of valid nodes in the pool; and *k*, the maximum allowable cycle length. It produces as output the set *Cycles* containing all unique feasible cycles of length at most k present in $G(t)$. The auxiliary function Cycle_Step extends a partial path recursively, ensuring that each node participates in at most one path per search root and backtracks immediately upon reaching the depth threshold $k-1$.

Algorithm 1: **DFS-Based Cycle Enumeration**

```
1:  function FIND_CYCLES(CycleNodes, NDDSet, adjList, nv, k)
2:      visited[v] ← False for all v ∈ {0, ..., nv − 1}
3:      for each v ∈ CycleNodes do
4:          if v ≥ nv then
5:              continue
6:          end if
7:          stack ← [v]; visited[v] ← True; depth ← 0
8:          for each u ∈ adjList[v] do
9:              if u < v then
10:                 continue                          ▷ Avoid duplicate cycles
11:             else if u = v then
12:                 Cycles.add(copy(stack))           ▷ Self-loop: 1-cycle
13:             else if ¬visited[u] and u ∉ NDDSet then
14:                 visited[u] ← True                 ▷ Mark node as part of current path
15:                 CYCLE_STEP(u, v, stack, 0, k)
16:                 visited[u] ← False                ▷ Unmark node for backtracking
17:             end if
18:         end for
19:         visited[v] ← False                        ▷ Release root node for other searches
20:     end for
21: end function

22: function CYCLE_STEP(u, root, stack, depth, k)
23:     if depth ≥ k − 1 then
24:         return                                    ▷ Depth limit reached: cannot extend cycle further
25:     end if
```

```
26:     stack.append(u); depth ← depth + 1
27:     for each i ∈ adjList[u] do
28:         if i = root then
29:             Cycles.add(copy(stack))                    ▷ Root edge closes the cycle
30:         else if i > root and ¬visited[i] and i ∉ NDDSet then
31:             visited[i] ← True
32:             CYCLE_STEP(i, root, stack, depth, k)
33:             visited[i] ← False
34:         end if
35:     end for
36:     stack.pop(); depth ← depth − 1                     ▷ Backtrack: remove node from current path
37: end function
```

Algorithm 2 maps the feasible chains originating from NDD. It receives as input: *NDDNodes*, a subset containing exclusively NDD; alongside the matching *NDDSet*, *adjList*, *nv*, and the maximum chain length *k*. It produces as output the set *Chains* containing all unique feasible chains of length between 1 and $k-1$ patient-donor pairs initiated by an NDD node. Unlike the cycle procedure, every partial extension of a chain is immediately recorded as a valid chain, since any path of length 1 to $k-1$ beyond the NDD node constitutes a feasible exchange sequence. The auxiliary function Chain_Step extends the current chain recursively, applying the same backtracking mechanism as Cycle_Step to ensure correctness across all search paths.

Algorithm 2: **DFS-Based Chain Enumeration**

```
1: function FIND_CHAINS(NDDNodes, NDDSet, adjList, nv, k)
2:     visited[v] ← False for all v ∈ {0, ..., nv − 1}
3:     for each v ∈ NDDNodes do
4:         stack ← [v]; visited[v] ← True; depth ← 0
5:         for each u ∈ adjList[v] do
6:             if ¬visited[u] and u ∉ NDDSet then
7:                 visited[u] ← True                      ▷ Mark node as part of current path
8:                 CHAIN_STEP(u, stack, 0, k)
9:                 visited[u] ← False                     ▷ Unmark node for backtracking
10:            end if
11:        end for
12:        visited[v] ← False                             ▷ Release NDD for other searches
13:    end for
```

14: **end function**

15: **function** CHAIN_STEP(u, *stack*, *depth*, k)
16: *stack*.append(u); *depth* $\leftarrow$ *depth* $+ 1$
17: *Chains*.add(copy(*stack*)) ▷ Every prefix is a valid chain
18: **if** *depth* $< k - 1$ **then**
19: **for** each $i \in$ *adjList*$[u]$ **do**
20: **if** $\neg$*visited*$[i]$ **and** $i \notin$ *NDDSet* **then**
21: *visited*$[i] \leftarrow$ **True** ▷ Mark node as part of current path
22: CHAIN_STEP(i, *stack*, *depth*, k)
23: *visited*$[i] \leftarrow$ **False** ▷ Unmark node for backtracking
24: **end if**
25: **end for**
26: **end if**
27: *stack*.pop(); *depth* $\leftarrow$ *depth* $- 1$ ▷ Backtrack: remove node from current path
28: **end function**

The two properties of both enumerations that are worth highlighting are: For cycles, restricting exploration only to neighbors with an higher index ensures each cycles is discovered exactly once, as every cycle has a unique minimum-index node serving as its root. For chains, every partial path extending from an NDD node is immediately recorded as a valid chain.

The complete sets of cycles and chains produced by Algorithms 1 and 2 constitute the feasible set $\mathcal{C}$ used in the integer programs of Section 3.1.2. In all experiments presented in this work, the maximum length is set to $k = 3$, meaning that cycles contain at most three patient-donor pairs and chains contain at most one NDD followed by two patient-donor pairs. This choice is consistent with the findings discussed in the Chapter 2, where the literature establishes $k = 3$ as near-universally adopted in operational KEPs, as the gains from allowing longer exchanges are marginal while the operational complexity increases substantially [4].

3.3.4 Immunosuppression Treatment Mechanism

Half-compatible pairs that remain unmatched after a match run face a decision: proceed with immunosuppressive transplantation or remain in the pool for future match runs. In our model, after each match run, each half-compatible pair opts for immunosuppression

with probability $p \in \{0\%, 25\%, 50\%, 75\%, 100\%\}$. We perform independent simulations for each of these five scenarios across all matching policies.

3.4. Data Generation

Our implementation incorporates a comprehensive set of KEP participants: incompatible, half-compatible, and compatible pairs, alongside Non-Directed Donor (NDD). Each instance is generated with different seeds, yielding different pool sizes and temporal dynamics. To make our instances more realistic, we use a Poisson Process to model participant arrivals using the following rates adapted from Monteiro et al. [8]:

- *Pairs Arrival Rate* = $1/3.6$
- *Non-Directed Donors Arrival Rate* = $1/75$

Although the pair arrival rate is applied uniformly across every type of pair (incompatible, half-compatible, compatible), it is unrealistic to assume that every compatible pair that arrives will participate in the KEP, given that they can undergo direct transplantation with their own donor. To account for this, our model assumes that only 10% of the generated compatible pairs opt to enter the KEP in order to find a better matching or to help others find a compatible solution.

3.4.1 Pairs Characteristics

To generate realistic KEP instances, each pair is assigned a set of clinical and logistical attributes based on reported distributions from the literature. These characteristics are defined as follows:

- **Blood Type:** Patient and donor blood types are independently sampled according to the distribution reported in Roth et al. [4]: $P(O)$ = 0.48, $P(A)$ = 0.34, $P(B)$ = 0.14, and $P(AB)$ = 0.04.
- **Panel Reactive Antibody (PRA) Level:** Patient PRA levels are sampled according to the distribution reported in Roth et al. [4]: $P(low)$ = 0.70, $P(medium)$ = 0.20, and $P(high)$ = 0.10.
- **Spouse Donor:** The spouse donor probability is sampled according to the distribution used in Andersson and Kratz [7]: $P(non-spouse)$ = 0.90 and $P(spouse)$ = 0.10.

- **Patient Titers:** The probability of the patient having low or high antibodies concentration is sampled according to the distribution reported in Andersson and Kratz [7]: $P(low)$ = 0.75 and $P(high)$ = 0.25.
- **Compatibility Type:**
  - *Fully Compatible:* ABO-compatible and HLA-compatible;
  - *Half-compatible:* ABO-incompatible but HLA-compatible, with low titers level;
  - *Incompatible:* ABO-incompatible or ABO-compatible, and HLA-incompatible.
- **Departure Time:** Pairs departure times are generated using a geometric distribution, with a type-specific daily departure probability p based on probabilities used in Monteiro et al. [8]: $p = 0.1$ for Compatible pairs, $p = 0.01$ for NDD, and $p = 0.001$ for Incompatible and Half-Compatible pairs.

3.4.2 Compatibility between Pairs

For all pairs in a generated instance, pairwise compatibility (ABO and HLA) is verified. If the donor of one pair is compatible with the patient of another, this directed relationship is stored. During the simulation, these connections are then used to construct the compatibility graph and find feasible cycles and chains.

3.4.3 Instances Creation

For each time horizon (6, 12, and 24 years), we generated 100 independent instances. Each generated instance is then used as the input for all five immunosuppression acceptance rate scenarios, meaning that each scenario within a given time horizon operates on the identical set of generated pairs — the only difference across scenarios being the per-period probability with which half-compatible pairs accept to undergo immunosuppressive transplantation rather than remain in the pool, which is drawn according to the scenario-specific probability p described in Section 3.3.4.

3.5. Performance Metrics

We evaluate the performance of our policies using metrics that capture the effectiveness, equity, and behavioral outcomes of each policy:

Efficiency Metrics:

- **Exchange Transplants via KEP:** Average number of pairs matched through cycles or chains.
- **Total Transplants:** Average number of pairs that received a kidney (via exchange in KEP or immunosuppression treatment).
- **Timeouts:** Average number of pairs that depart the active pool without transplantation due to death or becoming too sick for surgery while waiting.

For clarity and conciseness throughout the results, we adopt the following notation to refer to specific transplantation outcomes. At the pair-compatibility level: **Incompatible pairs Matched in the KEP (IMK)**, **Half-compatible pairs Matched in the KEP (HMK)**, and **Half-compatible pairs Directly Transplanted with Immunosuppression (HDTI)**. At the system level: **All pairs Matched in the KEP (AMK)** and **All pairs Transplanted Overall (ATO)**, where $\text{ATO} = \text{AMK} + \text{HDTI}$. These acronyms are used consistently from this point onward, particularly in figures and tables throughout Chapter 4.

Equity Metrics:

- **Transplant by Pairs Characteristics:** Average number of pairs matched by specific subgroups (e.g., compatibility type, blood type, PRA level).
- **Average Waiting Time:** Mean time from pool arrival to transplantation for matched pairs.
- **Waiting Time for Worst Cases:** Average waiting time for the 10% of pairs experiencing the longest times in the pool until they get matched, capturing worst-case equity.

System Dynamics Metrics:

- **Pool Evolution:** Tracking active pool sizes over time, alongside phase-specific evaluations comparison.
- **Time-to-Event Distributions:** The measurement of how quickly pairs are processed (matched, timeout, or left unmatched at the simulation's conclusion) relative to their arrival time.

- **Cycles and Chains Utilization:** Percentage of successful matches achieved using cycles and chains.
- **Average Policy Iteration Time:** The mean computational time required by each policy to execute a single match run instance.

These metrics facilitate a comprehensive evaluation of the KEP simulation: efficiency metrics capture overall matching volume, equity metrics evaluate fairness in access to transplantation, and system dynamics metrics characterize the temporal behavior of both the pool and the matching policies.

3.6. Experimental Design and Implementation

3.6.1 Experimental Factors

The framework of experimental design is structured around three primary factors:

- **Matching Policy (4 levels):** Policy 1: Hierarchical with Waiting Time, Policy 2: Hierarchical without Waiting Time, Policy 3: Maximum Cardinality, and Policy 4: Static.
- **Immunosuppression Acceptance Rate (5 levels):** 0%, 25%, 50%, 75%, and 100%.
- **Time Horizon (3 levels):** 6, 12, and 24 years.

A full factorial design across these parameters would yield 60 distinct scenarios. However, Policy 4: Static was excluded from the 24-year simulation due to computational limitations. Because the static model evaluates all pairs simultaneously, the larger compatibility graph generated over a 24-year period prevented the program from reaching optimality within the allocated solver time limit (900 seconds). Consequently, Policy 4: Static was only evaluated for the 6- and 12-year horizons, resulting in 55 unique scenario combinations. With 100 different instances generated per scenario, this yields a total of 5,500 simulation runs.

3.6.2 Computational Implementation

All simulation frameworks and matching policies were implemented in Python 3.11.0.

To ensure computational reliability and verify intended behavior across diverse operational scenarios, the implementation was rigorously validated using the `pytest` framework, achieving 100% statement coverage across all matching policies and simulation dynamics. The underlying integer linear programs were solved using Gurobi Optimizer 12.0.3, an optimization solver chosen for its high performance on ILP problems [53]. We used default Gurobi parameters with a 900-second time limit per match run, meaning that if the solver does not find a proven optimal solution within that time window, it returns the best feasible solution found up to that point.

All simulations were executed on a Windows 11 system with an AMD64 processor ($\sim 3.6\,GHz$) and 16 GB of RAM.

4. Results and Analysis

This chapter presents and analyzes the empirical findings obtained from the 5,500 simulation runs conducted across the full experimental design. Unless stated otherwise, all analyses are conducted on the 24-year simulation horizon, as it provides the most robust characterization of long-term policy behavior. Results are visualized using spaghetti plots where each thin line represents a single instance and the thick black line represents the mean across all 100 instances. The analysis focuses primarily on the performance outcomes for incompatible pairs, as these represent the only pool elements for whom KEP matching constitutes the sole viable path to transplantation. Policy 4: Static is excluded from most analyses given its non-dynamic nature. It is included only as a reference, representing the theoretical best achievable outcome regarding the number of realized transplants if the entire pool were available simultaneously. We begin by characterizing pool composition and dynamics, before examining transplantation outcomes and the efficiency-equity trade-off across the hard-to-match subpopulations. This is complemented by an analysis of timeout behavior, waiting time distributions for matched pairs, and computational performance. The chapter concludes by examining temporal dynamics, contrasting outcomes across simulation phases to isolate how policies behavior evolves as the pool matures.

4.1. Descriptive Statistics of Generated Pools

Table 4.1 summarizes the characteristics of generated pools across the simulated time horizons. On average, the 6-, 12-, and 24-year simulations generated 368, 737, and 1475 pairs per pool, respectively. The pool composition remains consistent across all immunosuppression acceptance rates scenarios, as these rates only affect pair behavior during the simulation rather than influencing the initial data generation process.

Simulation Years	Compatibility	Average of Pairs per Pool	Percentage (%)
6 Years	Incompatible	181	49.19
	Half-compatible	128	34.78
	Compatible	30	8.15
	Non-Direct Donors	29	7.88
12 Years	Incompatible	364	49.40
	Half-compatible	255	34.60
	Compatible	59	8.00
	Non-Direct Donors	59	8.00
24 Years	Incompatible	731	49.55
	Half-compatible	510	34.58
	Compatible	119	8.07
	Non-Direct Donors	115	7.80

Table 4.1: Average pool composition and percentage distribution of pairs by compatibility type across the 6-, 12-, and 24-year simulation horizons.

4.2. Active Pool Size Dynamics

For the 24-year simulation, Figure 4.1 illustrates the evolution of the active pool size over time for each policy under 25% immunosuppression acceptance rate. Solid lines represent the total number of pairs in the pool; dashed lines reflect only incompatible pairs. We focus our analysis only on this specific scenario because the underlying behavioral trends remain highly consistent across all acceptance rates. The only standard observable difference across these scenarios is that higher acceptance rates result in a smaller average pool size.

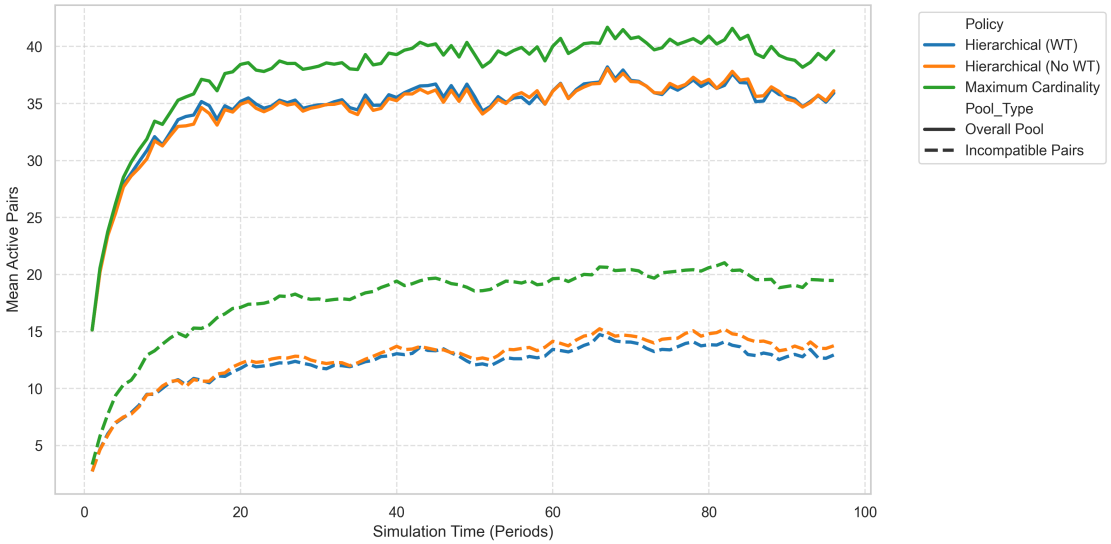

Figure 4.1: Mean active pool size per period for each policy, under 25% immunosuppressive acceptance rate (solid lines represent the total number of pairs in the pool; dashed lines represent incompatible pairs).

All models exhibited a ramp-up (warm-up) phase that lasted around the first 25 periods (approximately 6.25 years). After this initialization period, pool sizes stabilized at steady-state levels that vary according to each policy: Policy 1: Hierarchical with Waiting Time and Policy 2: Hierarchical without Waiting Time maintain the smallest average pool size with numbers very similar to each other (between 34 and 37 pairs, on average). Policy 3: Maximum Cardinality shows the worst results, with an average of between 38 and 41 pairs in the pool; Policy 4: Static was not included in pool dynamics analysis given its single-run nature. Additionally, we can also observe that, even during the warm-up period, Policy 3: Maximum Cardinality presents the highest average number of pairs in the pool.

The evolution of the number of incompatible pairs mirrors the results of the overall pool: a warm-up phase followed by a steady-state, with Policy 1 and 2 consistently maintaining fewer incompatible pairs in the pool over the time horizon when compared to Policy 3. This indicates that prioritization accelerates clearance of this subpopulation in particular. The most notable distinction is that Policy 1 maintains a slightly lower average than Policy 2, suggesting that the additional prioritization constraints further accelerate the clearance of incompatible pairs.

The only scenario that diverges from the others is the one with 0% immunosuppression acceptance rate (Figure 4.2), in which we observe a complete inversion of policy behavior regarding the overall pool evolution compared to the other four acceptance rates. In this

scenario, Policy 1: Hierarchical with Waiting Time consistently maintains the largest overall pool size, while Policy 3: Maximum Cardinality, in the majority of the simulation horizon, yields the lowest pair counts. This behavioral inversion is structurally expected. Without the alternative option of direct immunosuppressive treatment, half-compatible pairs are entirely captive to the KEP. Because the hierarchical policies are explicitly designed to prioritize incompatible pairs first, they unintentionally increase the size of the active pool by making the half-compatible pairs to wait and accumulate. The corresponding pool evolutions for the remaining immunosuppressive acceptance rate scenarios (50%, 75%, and 100%) are provided in **Appendix A.1**.

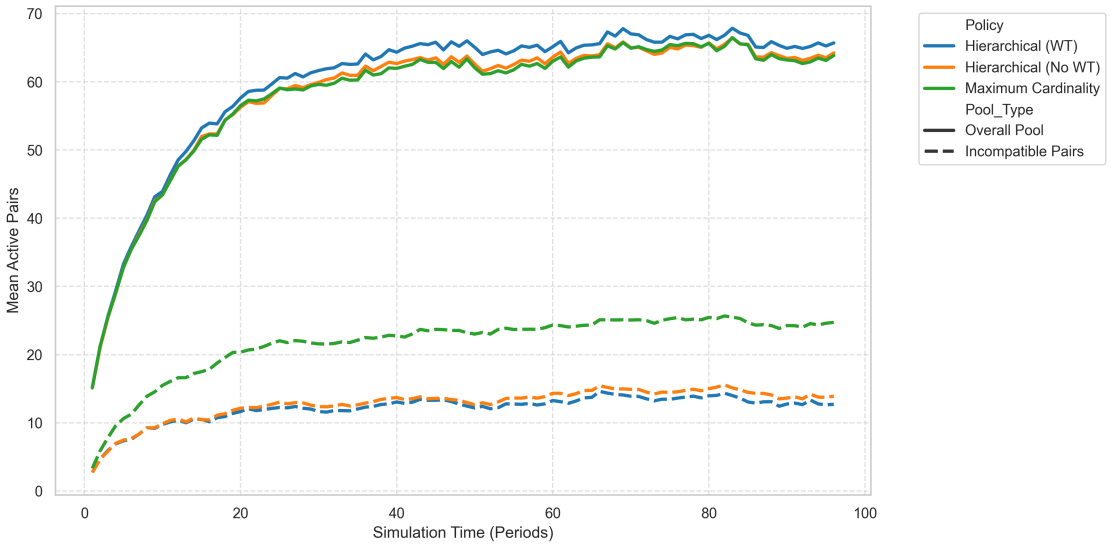

Figure 4.2: Mean active pool size per period for each policy, under 0% immunosuppressive acceptance rate (solid lines represent the total number of pairs in the pool; dashed lines represent incompatible pairs).

To ensure that our results remain consistent under scenarios with higher pool density, we conducted a sensitivity analysis on the 24-year time horizon (Figure 4.3) by doubling the baseline pair arrival rate shown in Figure 4.1.

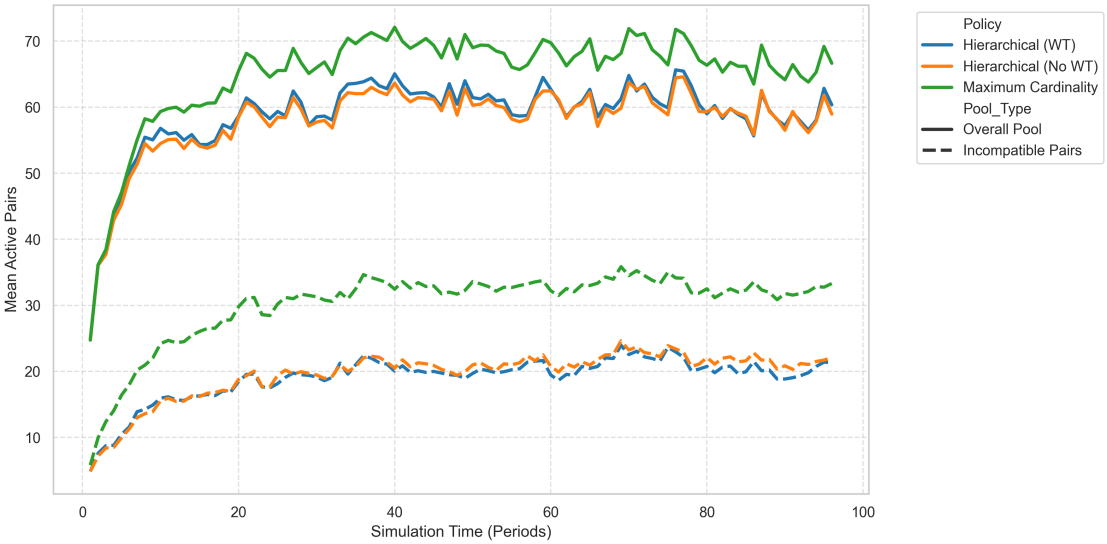

Figure 4.3: Mean active pool size per period for each policy with higher pair arrival rate, under 25% immunosuppressive acceptance rate (solid lines represent the total number of pairs in the pool; dashed lines represent incompatible pairs).

As expected, under steady-state conditions, the scenario with the doubled arrival rate yields pools with, approximately, twice the size of the baseline pool shown in Figure 4.1, affecting both the overall pool and the incompatible equally. Regarding the policies, we can see that they have similar behavior when compared to the previous observations: Policy 3: Maximum Cardinality consistently retains the highest average number of pairs in the pool, while the two hierarchical policies maintain lower numbers, following a similar trend to that shown in Figure 4.1. This confirms that our comparative policy dynamics are robust and scale consistently regardless of the arrival volume of pairs in the pool.

4.3. Efficiency Analysis

This section evaluates the policies based on their ability to maximize overall resource utilization. We analyze the number of transplants facilitated through KEP exchange (AMK), the total transplants realized (ATO), and the frequency with which pairs depart the pool unmatched (timeouts).

4.3.1 Transplants Performed via KEP

We first compare the average number of transplanted pairs strictly achieved through cycle and chain matching (AMK), excluding direct transplants between donors and patients in half-compatible pairs, as a function of the immunosuppression acceptance rate. This

analysis focuses on the 24-year simulation, as it provides the most robust and reliable representation of the pool behavior compared to the 12- and 6-year simulations (Figure 4.4). Unlike the shorter horizons, 24 years gives the simulation enough time to stabilize, generates a much larger dataset, and allows us to accurately observe how pairs behave over time.

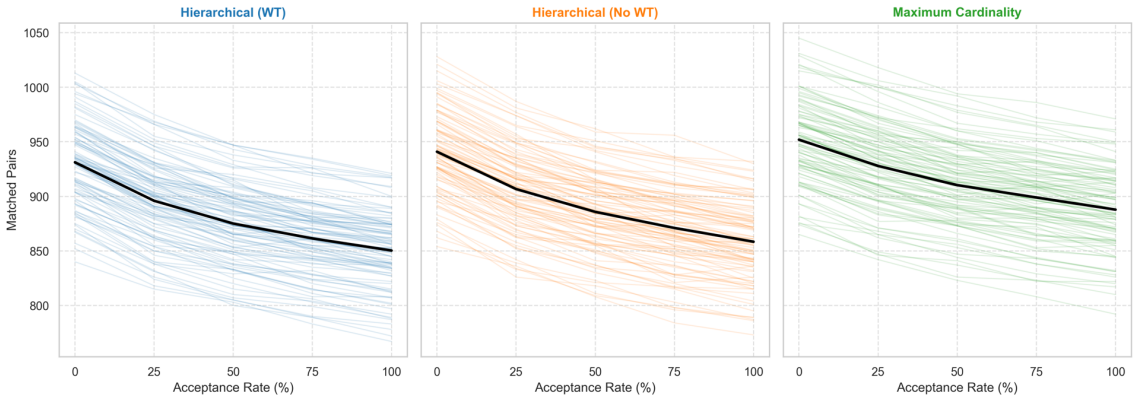

Figure 4.4: Average number of All pairs Matched in the KEP (AMK) (24-year simulation).

As anticipated, Policy 3: Maximum Cardinality consistently outperforms the hierarchical approaches achieving the highest absolute number of matches across all acceptance rate levels. Because it operates without any equity or prioritization constraints, it maintains the largest feasible solution space. The decrease in number of AMK as immunosuppression acceptance rate increases is also to be expected. Higher acceptance probabilities naturally reduce the number of half-compatible pairs remaining in the active pool after each period, subsequently reducing the number of HMK(Figure 4.5).

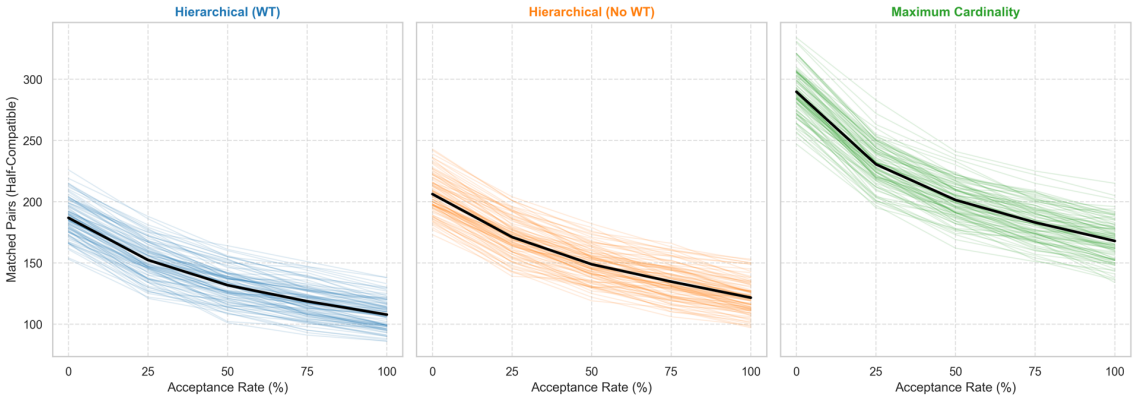

Figure 4.5: Average number of Half-compatible pairs Matched in the KEP (HMK) (24-year simulation).

When evaluating the static upper bound (Figure 4.6), Policy 4: Static achieved a superior

performance ranging from 5.5% to 13%, depending on the immunosuppressive acceptance rate, over Policy 3: Maximum Cardinality in the 6- and 12-year simulations, and outperformed the two hierarchical policies by up to 19%. This gap quantifies the inherent but necessary cost of operating in a dynamic environment, where patients have a limited amount of time that they can wait for a transplant and, therefore, matches have to be performed overtime.

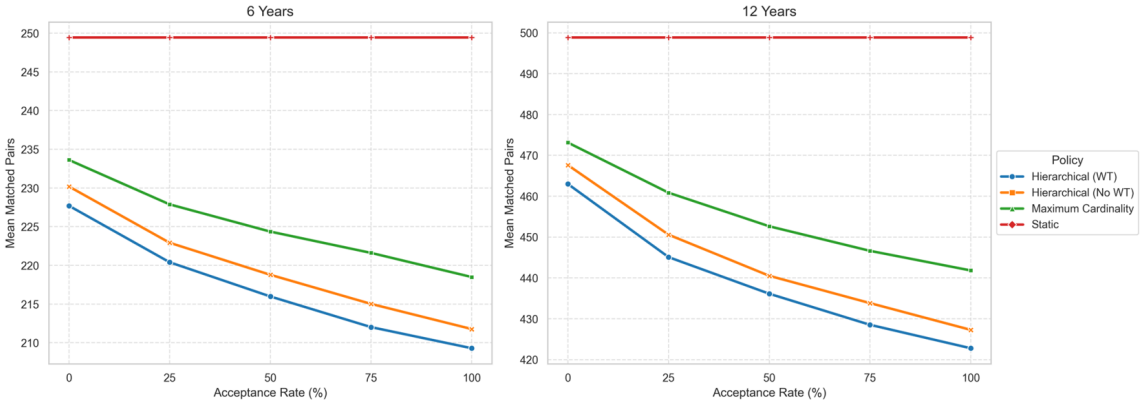

Figure 4.6: Average number of All pairs Matched in the KEP (AMK), for 6 and 12 year simulations.

4.3.2 Total Number of Transplants

The previous section evaluates the number of transplants achieved exclusively within the KEP pool (AMK), disregarding direct transplantation between donor and patient of half-compatible pairs. However, to assess the efficacy of the immunosuppression requires analyzing the total number of transplants of the system (ATO). This metric captures the average absolute number of arriving pairs that successfully receive a kidney throughout the simulation horizon, whether through a KEP match or via direct immunosuppressive treatment (Figure 4.7).

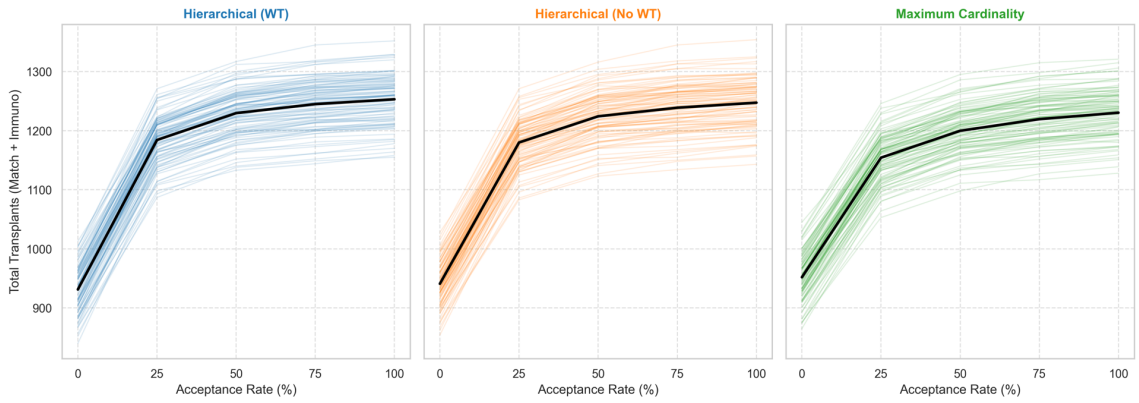

Figure 4.7: Total number of All pairs Transplanted Overall (ATO) (24-year simulation).

As illustrated in Figure 4.7, policies with more constrains benefit more from the immunosuppression alternative, consistently outperforming Policy 3: Maximum Cardinality in the number of ATO even in the lowest acceptance rate. This shows a completely opposite behavior from what was seen in the previous analysis for KEP matches (Figure 4.4) where Policy 3 was the one with the best results (highest number of AMK) across all acceptance rates, and the increase in the immunosuppression acceptance rate resulted in lower number of All pairs Matched in the KEP (AMK) across all policies.

This behavior can be explained as follows. Policy 3 only optimizes the number of transplants, it greedily matches easier-to-match half-compatible pairs inside the KEP, leaving incompatible hard-to-match pairs—who have no alternative treatment options—unmatched. The hierarchical policies, by contrast, actively force the KEP to match those hard-to-match pairs first, which leaves more half-compatible pairs unmatched within the KEP, pushing them toward immunosuppressive transplantation as an alternative pathway. By actively directing the KEP toward the hard-to-match pairs and allowing half-compatible pairs to exit the matching pool via immunosuppression, the equity-constrained policies inevitably achieve the highest efficiency when all transplantation alternatives are considered.

Notably, the total number of ATO scales up as the immunosuppression acceptance rate increases. This is an expected behavior: as the acceptance rate increases, a larger number of half-compatible pairs leaves the pool by opting for a direct desensitized transplantation, receiving a kidney without relying on the KEP system (HDTI). Figure 4.8 presents results for those cases.

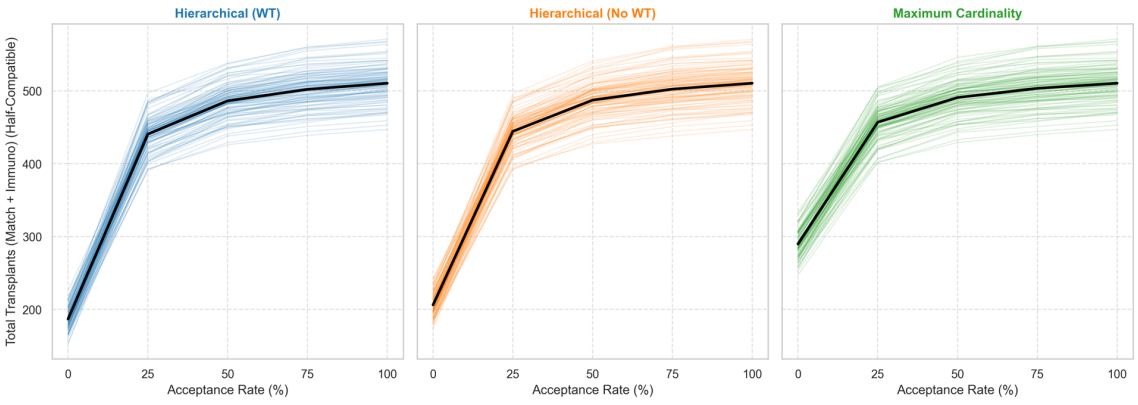

Figure 4.8: Total number of Half-compatible pairs Matched in the KEP (HMK) + Half-compatible pairs Directly Transplanted with Immunosuppression (HDTI) (24-year simulation).

4.3.3 Timeout Pairs

System efficiency is also measured through the number of timeout pairs—pairs that reach their maximum waiting time tolerance and depart the pool without ever being transplanted. We evaluated these average departures for both the overall pool (Figure 4.9a) and incompatible pairs (Figure 4.9b).

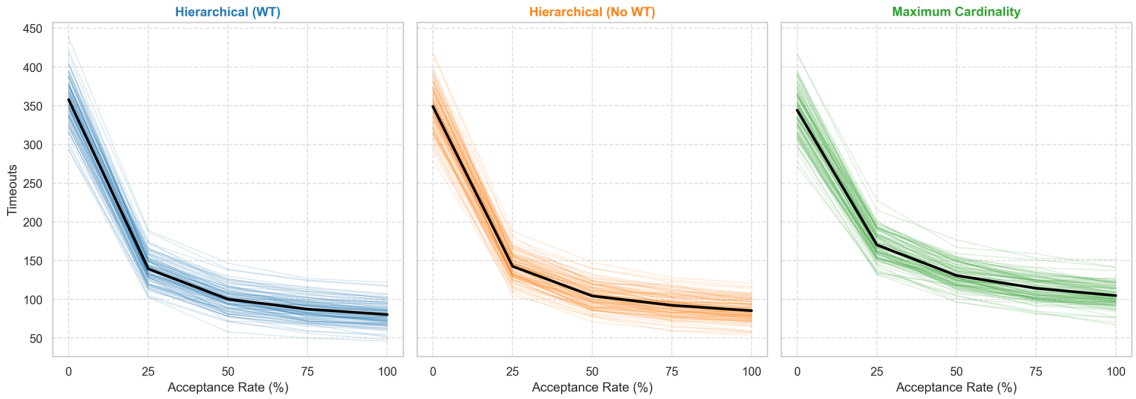

(a) Considering all pairs participating in the pool.

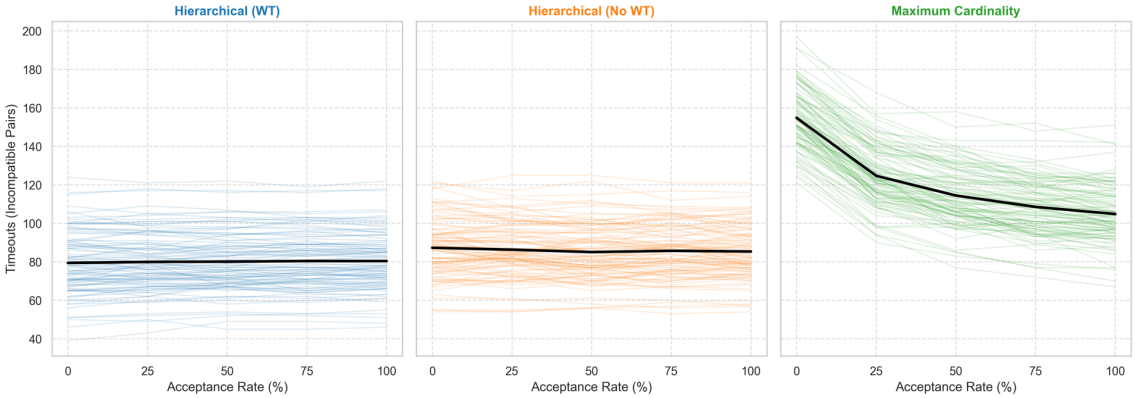

(b) Considering only incompatible pairs in the pool.

Figure 4.9: Average number of pairs leaving the pool before being transplanted.

In terms of minimizing timeouts, the hierarchical models outperformed the Maximum Cardinality policy, both for the overall pool and when looking only at the incompatible pairs. For the overall pool, an increase in the immunosuppression acceptance rate naturally decreases the average number of timeouts, as half-compatible pairs exit the system proactively. However, when isolating incompatible pairs, the average number of timeouts under the hierarchical policies remains stable regardless of the acceptance rate. By contrast, Policy 3: Maximum Cardinality accumulates a growing number of unmatched hard-to-match pairs that eventually depart without a transplant, resulting in a substantially higher

average timeout numbers. This confirms that policies without targeted prioritization strategies systematically sacrifice the most vulnerable patients to achieve marginal gains in overall transplant volume.

4.4. Equity Analysis

This section evaluates how effectively the different policies enforce equity within the matching process. To quantify this, we assess the policies' ability to match hard-to-match demographic groups and their capacity to minimize the waiting time pairs experience before being matched.

4.4.1 Transplants by Pair Characteristics

To assess the efficiency-equity trade-off, we first isolate the matching outcomes for the baseline vulnerable demographic: incompatible pairs (Figure 4.10). We then further analyze this group regarding the most disadvantaged subpopulations: incompatible pairs with Blood Type O patients (Figure 4.12) and incompatible pairs with high PRA levels (Figure 4.13).

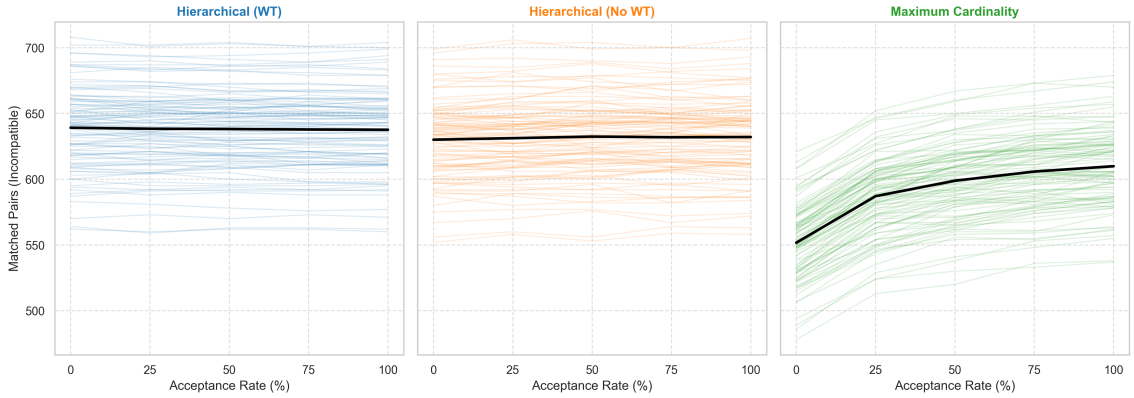

Figure 4.10: Average number of Incompatible pairs Matched in the KEP (IMK) (24-year simulation).

When isolating the incompatible pairs, the conclusions drawn from the efficiency analysis reverse. The hierarchical policies (Policy 1 and Policy 2), which have strict compatibility-based prioritization, consistently outperform the unconstrained Policy 3: Maximum Cardinality. By strictly enforcing lexicographic prioritization, the hierarchical policies force the ILP to select cycles and chains containing incompatible pairs, even if doing so reduces the total number of AMK in a given period. The consistency in the number of IMK across all immunosuppression acceptance rates indicates that the presence of half-compatible pairs in the pool does not negatively impact incompatible pairs. Instead, it benefits the

half-compatible pairs by allowing them to have a chance to find a match without resorting to immunosuppression treatment.

As established in the efficiency analysis, Policy 4: Static achieved the highest overall number of AMK by optimizing across the entire pool simultaneously (see Figure 4.6). However, because Policy 4 is fundamentally a maximum cardinality formulation without equity constraints, the hierarchical policies completely invert the performance hierarchy when isolating incompatible pairs (Figure 4.11).

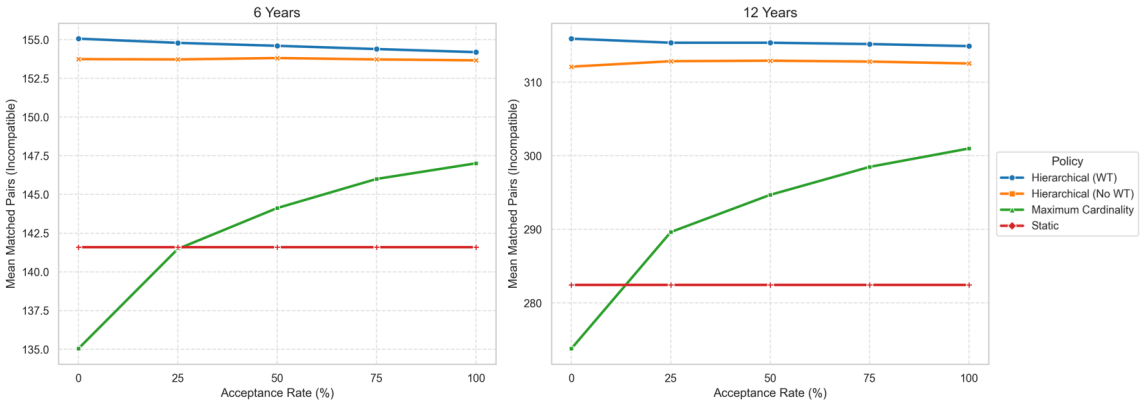

Figure 4.11: Average number of Incompatible pairs Matched in the KEP (IMK) for 6 and 12 year simulations.

As we can observe, in both the 6- and 12-year simulation periods, Policy 1: Hierarchical with Waiting Time and Policy 2: Hierarchical without Waiting Time clearly outperform Policy 4. Policy 1 consistently has the best overall performance surpassing Policy 4 by approximately 9% to 12%, which translates into a dozen additional transplants for incompatible pairs.

Beyond the aggregate incompatible pool, disadvantaged demographics subgroups, such as incompatible pairs with Blood Type O patients (Figure 4.12), also benefit from the prioritization rules. Patients with Blood Type O face disadvantages in matching processes; they can only receive kidneys from donors with Blood Type O, while donors with Blood Type O are universally compatible and frequently used by policies to close easier cycles for patients who are not blood type O.

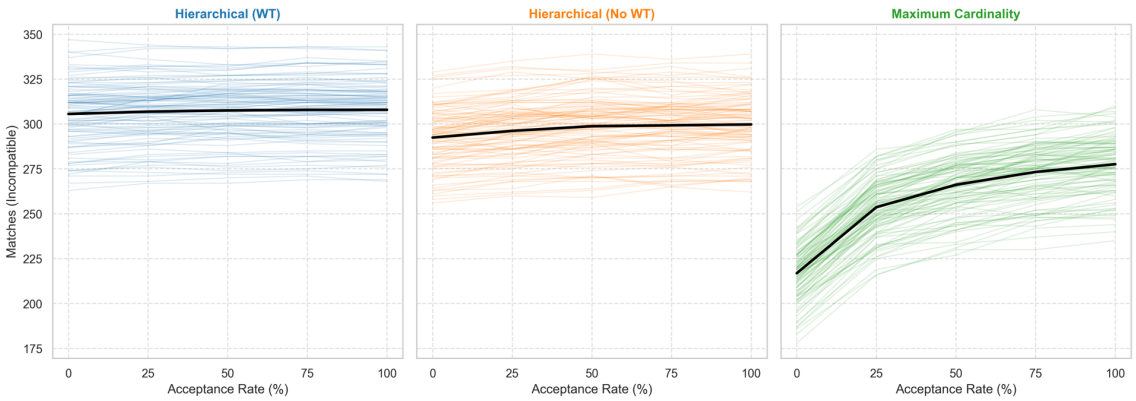

Figure 4.12: Average number of Incompatible pairs Matched in the KEP (IMK) with blood type O patients (24-year simulation).

As illustrated in the Figure 4.12, both hierarchical policies consistently facilitate a higher number of IMK for Blood Type O patients compared to Policy 3. By enforcing prioritization, the hierarchical models implicitly prevent Blood Type O donors from being incorporated into lower-priority matches, maintaining them as an essential resource for Blood Type O patients who would otherwise remain unmatched.

Other disadvantaged demographics subgroups are the incompatible pairs with highly sensitized (High PRA level) patients.

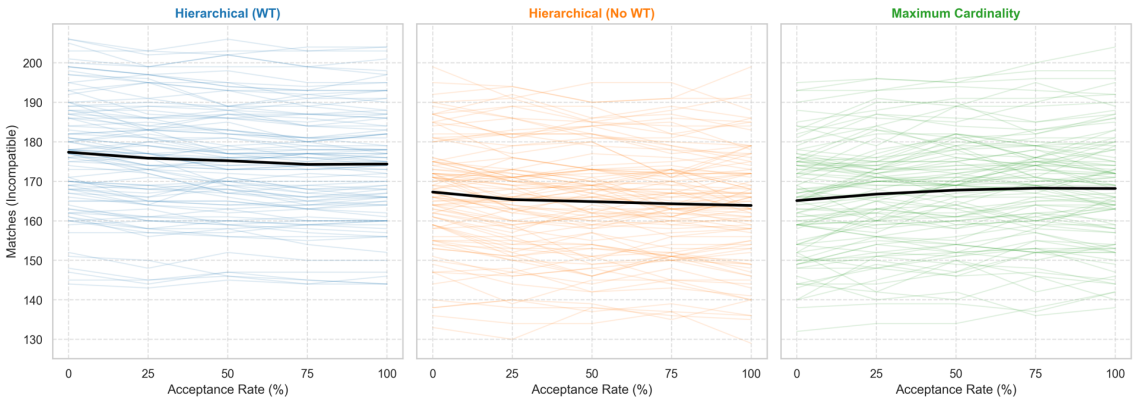

Figure 4.13: Average number of Incompatible pairs Matched in the KEP (IMK) with high PRA level patients (24-year simulation).

In Figure 4.13, we can observe that as the immunosuppression acceptance rate increases, half-compatible pairs leave the pool more frequently by opting for immunosuppression treatment, reducing the number of feasible cycles and chains available to Policy 3, forcing it to opportunistically include high- PRA incompatible pairs whenever a rare compatible cycle arises. Policy 2: Hierarchical without Waiting Time, by contrast, applies a rigid class-based prioritization without a waiting-time tiebreaker, which means it has no mechanism

to distinguish between recently-arrived and long-waiting incompatible pairs, in which high-PRA incompatible pairs may be included, within the same priority class. As pool density decreases, Policy 2 may consistently select easier-to-match incompatible pairs over high-PRA ones, leaving the most sensitized patients unmatched. Policy 1: Hierarchical with Waiting Time avoids this structural problem entirely; its temporal tiebreaker naturally elevates long-waiting high- PRA pairs in the priority hierarchy, which explains its consistently superior performance for this subgroup across all acceptance rates.

4.4.2 Waiting Times Analysis

A key aspect of fairness in any dynamic matching system is the time between a pair's arrival and their successful transplantation. Because aggregate average waiting time can obscure the severe delays experienced by hard-to-match pairs, this section evaluates temporal equity from two distinct perspectives. We first analyze the the average waiting time for successfully matched pairs across the 24-year simulation horizon (Figure 4.14 for AMK, and Figure 4.15 only for IMK), and then we isolate the worst-case scenarios by examining the top 10% of the highest waiting times (Figure 4.16 for AMK, and Figure 4.17 only for IMK) to capture the true worst-case equity.

4.4.2.1 Average Waiting Times

For the average waiting time analysis, we evaluate AMK (Figure 4.14) and the IMK (Figure 4.15).

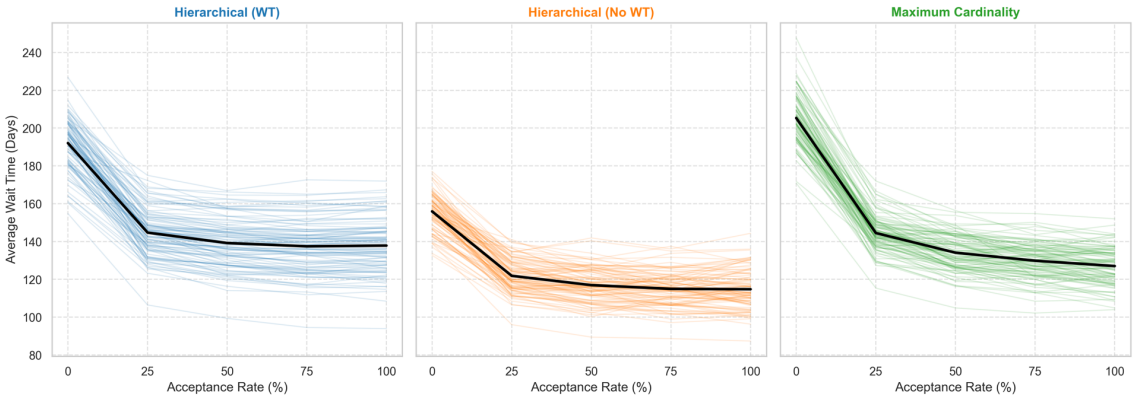

Figure 4.14: Average waiting time distribution for All pairs Matched in the KEP (AMK) (24-year simulation).

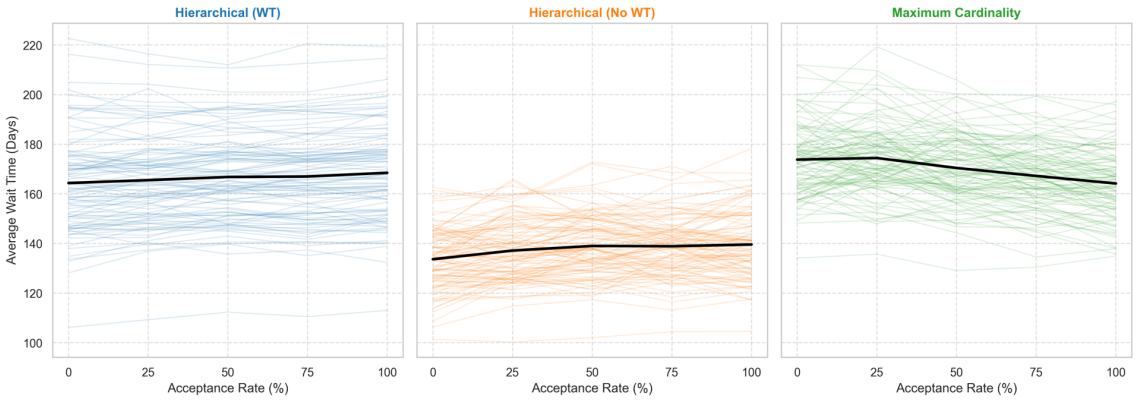

Figure 4.15: Average waiting time distribution for Incompatible pairs Matched in the KEP (IMK) (24-year simulation).

The average waiting time for AMK decrease for all policies as the immunosuppression acceptance rate increases. This is an expected compositional effect: as half-compatible pairs increasingly opt for direct desensitization, they depart the pool rapidly, pulling the overall aggregate waiting time downward.

However, when analyzing the policies performance—particularly within the incompatible pairs—the results reveal an unexpected pattern. At first glance, Policy 2: Hierarchical without Waiting Time appears to yield the most favorable temporal outcomes, maintaining the lowest average wait time (approximately between 135 and 139 days for the IMK). Policy 1: Hierarchical with Waiting Time records a higher average waiting time (approximately between 165 and 169 days), and Policy 3: Maximum Cardinality ranges between 165 and 175 days.

The lower average observed in Policy 2 is not an indicator of superior equity, but rather a statistical artifact of the prioritization rules. Because Policy 2 lacks a temporal tiebreaker, it matches the “easiest” incompatible pairs—typically those who have lower waiting times and possess diverse compatibility options—while leaving the most difficult pairs to match stranded in the pool indefinitely that will eventually leave the pool without a transplant, never being taken into account to this average of matched pairs.

Conversely, Policy 1 explicitly forces the optimization model to rescue these long-waiting, hard-to-match pairs. When Policy 1 successfully transplants a pair that has been waiting in the pool for several years, that massive waiting time is factored into the metric, pulling the policy’s average waiting time upward. Therefore, a higher mean waiting time for matched pairs in an equity-constrained policy is not necessarily a sign of poor performance, but rather indicates a successful clearance of long-waiting pairs. This phenomenon proves

the limitation of average waiting time alone as a fairness metric, and points out the necessity to examine the tail of the distribution—the pairs that wait the longest—to obtain a true picture of worst-case equity.

4.4.2.2 Waiting Times for Worst Cases (Top 10%)

While average waiting times provide a general temporal baseline, true equity is defined by how a system can handle the worst-case scenarios. We evaluate the 10% of pairs that experienced the longest waiting times before being matched, analyzing both AMK (Figure 4.16) and IMK (Figure 4.17).

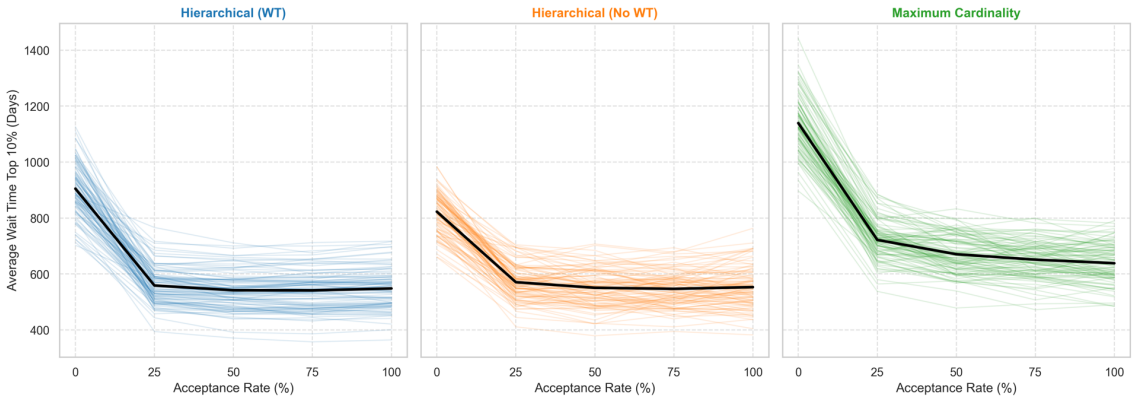

Figure 4.16: Waiting time distribution for the 10% of All pairs Matched in the KEP (AMK) with the longest waiting times (24-year simulation).

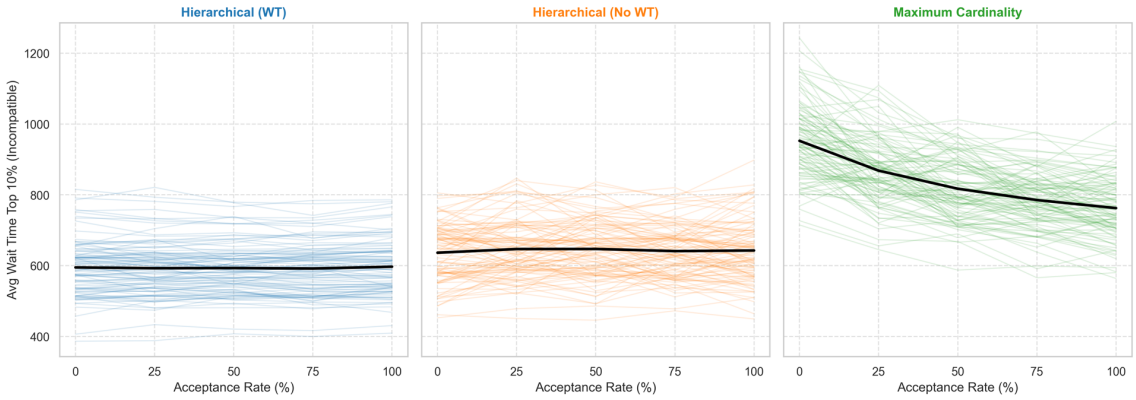

Figure 4.17: Waiting time distribution for the 10% of Incompatible pairs Matched in the KEP (IMK) with the longest waiting times (24-year simulation).

The models with lexicographic restrictions outperform the Policy 3: Maximum Cardinality, both for AMK and IMK by more than 100 days, with Policy 1 achieving slightly better results than Policy 2. The main reason for this is that the incompatible pairs, which are the most difficult to match, will normally experience the longest waiting times. However, the hierarchical models, by actively prioritizing these pairs, help them find matches as quickly

as possible, significantly compressing the waiting time tail compared to the Maximum Cardinality policy.

Additionally, for the hardest-to-match incompatible pairs (Figure 4.17), the waiting time under the hierarchical policies remain constant regardless of the immunosuppression acceptance rate. This temporal stability mirrors our earlier findings regarding the number of transplants. This proves that lexicographic prioritization of incompatible pairs shields them from the structural shifts in the pool caused by half-compatible pair departures. In contrast, the longest waiting times for the overall pool naturally decrease as the immunosuppression acceptance rate rises, simply because the half-compatible pairs, which are included in the overall metric, depart the pool earlier via direct immunosuppressive treatment.

It is worth noting that the waiting time for the 10% worst cases reported in this analysis are substantially higher than those observed in the Average Waiting Times analysis. This divergence is expected: the average metric aggregates AMK, the majority of whom are matched quickly within their first periods, naturally suppressing the mean. The top 10% analysis, by contrast, isolates precisely those pairs that survived multiple match runs without being matched, exposing the actual upper limit of waiting time exposure that is hidden by the overall average. When combined, the two approaches provide a comprehensive picture of temporal equity: the top 10% of the worst case represents the burden carried by the most vulnerable pairs, while the average represents the typical experience for the pairs in the pool.

4.5. System and Temporal Dynamics

In this section, we examine the underlying structural behaviors that drive the efficiency and equity outcomes established in the previous sections. We analyze the matched pairs across different simulation phases, how cycles and chains contribute for the overall number of transplants, how quickly pairs are processed after entering the pool, and what is the computational cost of each policy.

4.5.1 Phase-Specific Matching Outcomes

To understand how pool maturity and structural congestion impact matching efficacy, we conduct a temporal analyses: comparing the number of AMK and IMK between the initial 20 and final 20 simulation periods.

4.5.1.1 Early vs. Late Pool Dynamics: First 20 vs. Last 20 Match Runs

To examine whether policy performance is stable over the course of the simulation or whether it evolves as the pool matures, we compare outcomes observed during the first 20 match runs (periods 1 to 20) against those observed during the last 20 match runs (periods 77 to 96) of the 24-year simulation. The first window captures behavior during the warm-up phase, when the pool is still sparse and the compatibility graph has limited structure. The last window captures behavior at the end of the simulation horizon, when the pool has been operating in steady state for an extended period and the cumulative effects of each model's matching policy are now fully manifest. We analyze the mean matched pairs for both periods, Figure 4.18 for AMK and Figure 4.19 for the IMK.

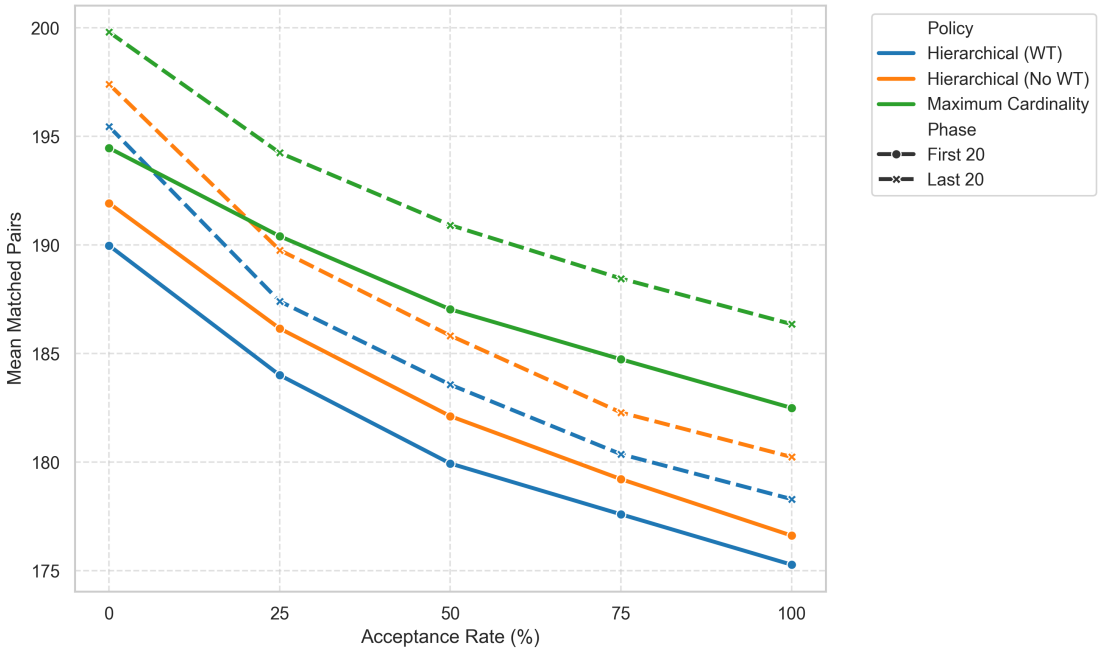

Figure 4.18: Average number of All pairs Matched in the KEP (AMK) (First 20, in solid lines, vs. Last 20 Match Runs, in dashed lines).

Analyzing AMK (Figure 4.18), we can observe that all policies have a higher absolute number of matches during the last 20 match runs compared to the first 20. This increase is a direct mathematical consequence of pool density: the number of pairs in the pool in the steady-state period is substantially higher than the initial warm-up pool, offering all policies a richer graph with more feasible cycles and chains to use. As expected, overall match volumes decrease across both temporal windows as the immunosuppression acceptance rate increases, since the proactive departure of half-compatible pairs leaves fewer total participants available for exchange.

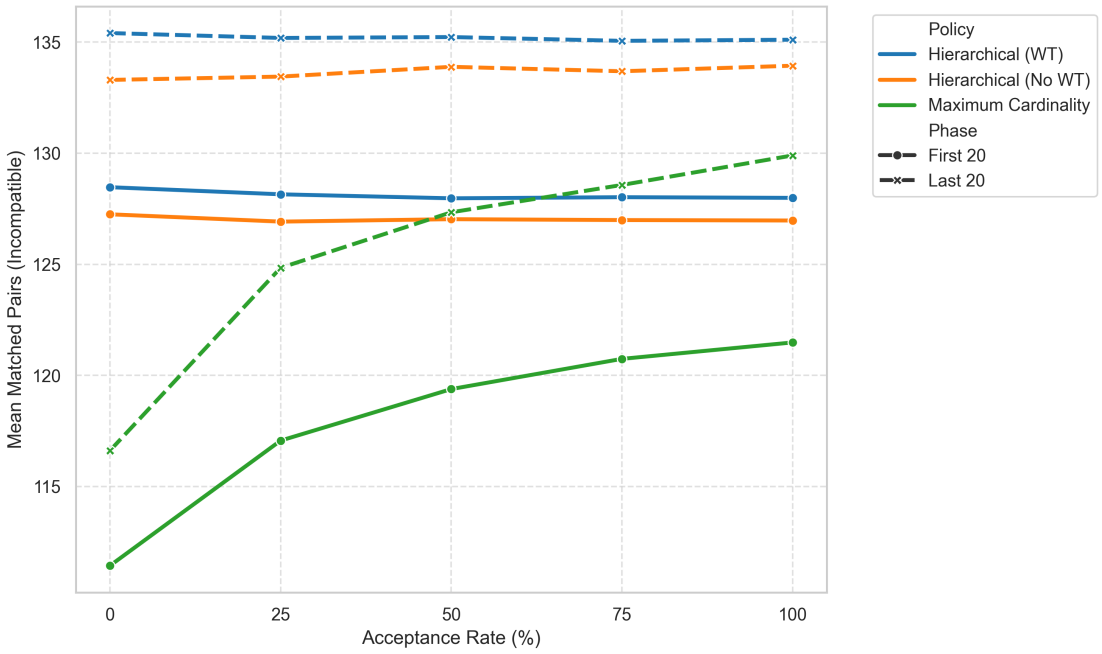

Figure 4.19: Average number of Incompatible pairs Matched in the KEP (IMK) (First 20, in solid lines, vs. Last 20 Match Runs, in dashed lines).

For the IMK (Figure 4.19), we see that the underlying behavioral trends remain consistent with previous analysis of IMK. During both periods, the hierarchical models (Policy 1 and Policy 2) outperform the Policy 3: Maximum Cardinality.

Crucially, while the number of transplants increases in the mature pool, the fundamental behavior of each policy across the different immunosuppression acceptance rates remains identical in both time periods. For both the first 20 and last 20 match runs, the results yielded by the hierarchical policies remain relatively stable, consistent with previous analysis. Policy 3 also exhibits the exact same structural behavior in both time periods: its match volume increases as the acceptance rate increases. Whether in the early or late stages of the simulation, Policy 3 only increases its matching of incompatible pairs when higher acceptance rates proactively remove the easier, half-compatible feasible cycles and chains from the pool.

In conclusion, this temporal comparison confirms three critical aspects of the KEP's long-term operational behavior. At the beginning of the program, there is a noticeable transitory warm-up phase, characterized by a sparse compatibility graph and lower number of matched pairs. As pairs arrivals continue, the system successfully stabilizes and tends toward a steady-state environment. Finally, and most importantly, the system remains highly responsive across both phases. Although there is an inevitable accumulation of

patients who are difficult to pair in the steady state, the dynamic system does not become structurally blocked. Crucially, the results suggest that it is not significantly more difficult for a pair to be matched in the mature steady-state phase than it was during the initial transitional phase.

To explicitly highlight this sustained responsiveness, we compare the time-to-match distributions for incompatible pairs entering the pool at different stages of the simulation. Figure 4.20 presents the matched occurrences as a function of time since pool entry, comparing pairs who enter during the initial transitional phase (Year 1) with those entering during the established steady-state phase (Year 12).

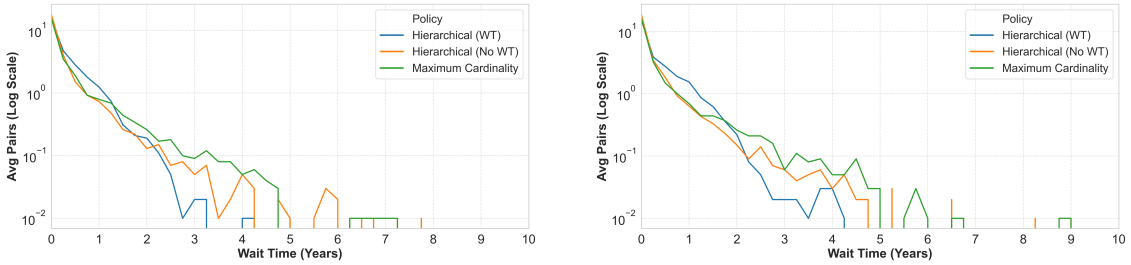

(a) Entry Period: Year 1 (Transitional)

(b) Entry Period: Year 12 (Steady State)

Figure 4.20: Average number of Incompatible pairs Matched in the KEP (IMK) by waiting time in the pool, comparing the transitional entry pairs (a) against the steady-state entry pairs (b) for 25% immunosuppression acceptance rate using a log scale.

As illustrated in Figure 4.20, the structural decay of the waiting time distribution remains remarkably consistent across both entry phases. For pairs entering the pool in Year 12, the system exhibits the same initial matching peak observed for pairs entering in Year 1, indicating that newly arriving easier-to-match pairs are processed just as rapidly in the mature pool as they were in the early pool. Furthermore, the policies successfully manage the structural accumulation of hard-to-match pairs without reducing the matching velocity for the subsequent arrivals, maintaining highly consistent behavioral profiles across both entry periods.

4.5.2 Cycles and Chains Utilization

To evaluate the structural composition of match outcomes, we analyzed the proportional reliance on cycles versus NDD-initiated chains (Figure 4.21), as well as the absolute number of AMK processed by specific topological structures, such as 2-way cycles, 3-way cycles, and chains (Figure 4.22).

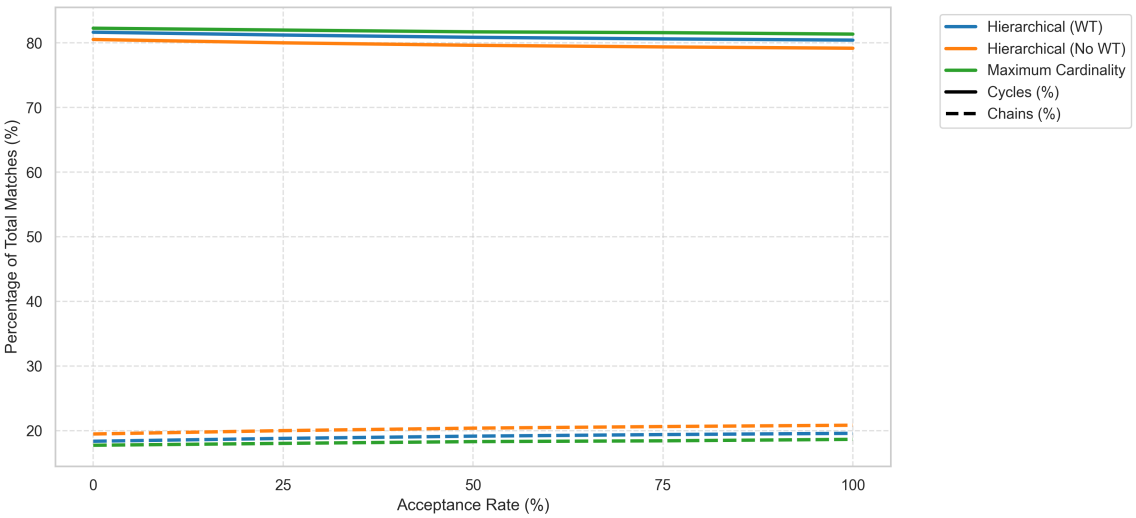

Figure 4.21: Proportional utilization of cycles (solid lines) versus chains (dashed lines) for the 24-year simulation.

Across all dynamic policies, cycles constitute the vast majority of matched structures, accounting for approximately 80% to 82% of all matched pairs. Despite NDD representing only 7.8% of pool arrivals, the chains they initiate account for roughly 18% to 20% of all KEP transplants. This asymmetry highlights the remarkable operational value of NDDs: acting as chain initiators. They possess the capacity to go through thick graph elements and unlock match combinations that otherwise would remain infeasible.

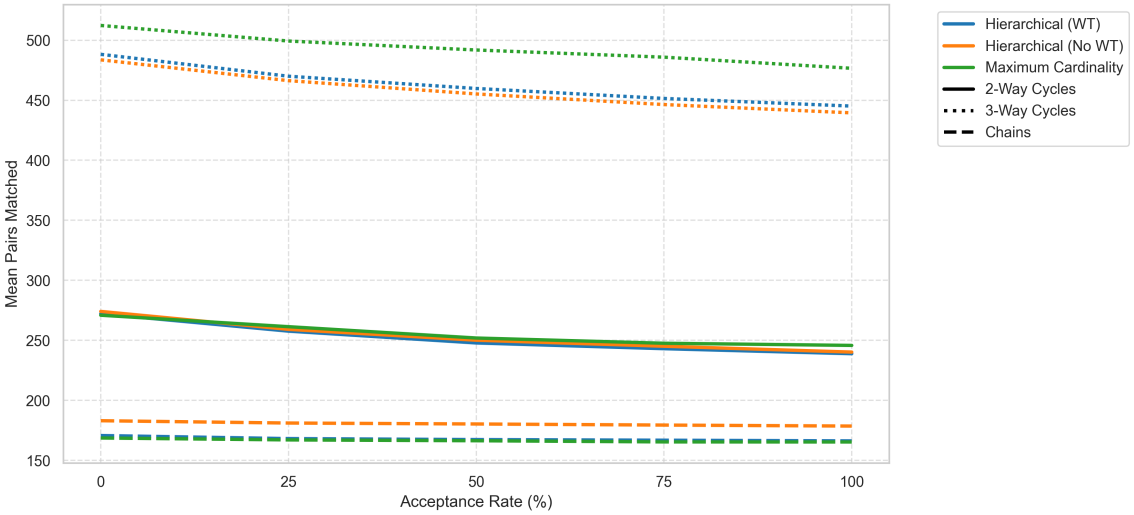

Figure 4.22: Average absolute number of all participating pairs matched by 2-way cycles (solid lines), 3-way cycles (dotted lines), or chains (dashed lines) for the 24-year simulation.

When breaking down the absolute number of matched pairs by specific match sizes (Figure 4.22), 3-way cycles emerge as the primary engine of the KEP, consistently matching nearly double the number of pairs compared to 2-way cycles.

A clear topological difference also emerges between policies regarding chain utilization. Both Policy 1: Hierarchical with Waiting Time and Policy 2: Hierarchical without Waiting Time consistently utilize a higher number of chains compared to Policy 3: Maximum Cardinality. Since Policy 3 optimizes purely for number of transplants, it naturally favors easy-to-form 3-way cycles. The hierarchical policies, by contrast, are forced by their lexicographic constraints to rescue hard-to-match incompatible pairs. Because chains are open-ended and don't require the final donor to close a loop back to the initiating patients, they offer greater structural flexibility. Consequently, the hierarchical policies make greater use of NDDs as chain initiators, focusing on getting them to clear the highest-priority compatibility classes, which explain their greater chain usage.

4.5.3 Time-to-Event Distributions

To further trace the life-cycle of an incompatible pair within the KEP, we analyze the time-to-event distributions for matched, timeout, and end-of-simulation unmatched pairs, examining how quickly incompatible pairs are processed after entering the pool.

Figure 4.23 presents the average number of IMK and of timeout incompatible pairs as a function of time since pool entry, for the 25% immunosuppression acceptance rate. We use a logarithmic scale to accommodate the wide range of values observed: the number of pairs processed is extremely high in the earlier match runs and decays sharply making a linear scale unsuitable for visualizing the full range simultaneously. For matched and timeout pairs, all three policies exhibit a sharp decline in the first match run, reflecting that the majority of incompatible pairs that will ever be matched are matched quickly — most within the first 90 days of entering the pool. After only a couple of years of waiting since pairs entered the pool, the matching rate has dropped to near zero across all policies, indicating that pairs remaining in the pool beyond this window are predominantly hard-to-match. All policies operate on an identical set of pairs, ensuring a fair baseline for comparison. Under these conditions, Policy 1: Hierarchical with Waiting Time stands out as the superior option in this analysis, as it both maximizes initial matches and reduces the number of pairs waiting over four years to nearly zero, while the other policies present higher values, requiring an additional two to three years to reach the same near-zero levels. The corresponding distributions for the remaining immunosuppressive acceptance rate scenarios (0%, 50%, 75%, and 100%) exhibit similar behavior and are provided in Appendix B.1.

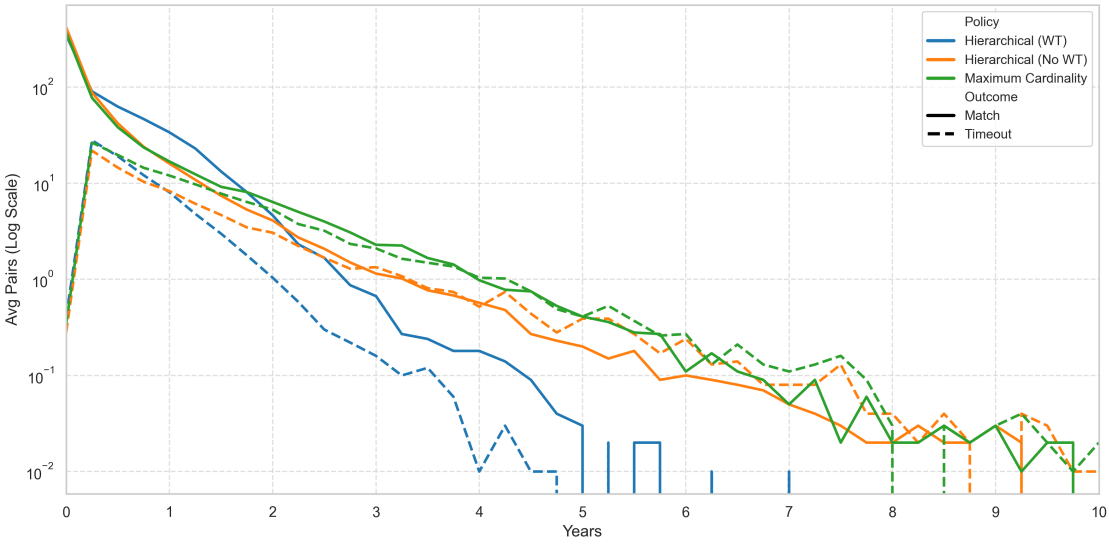

Figure 4.23: Average number of Incompatible pairs Matched in the KEP (IMK) and timeout occurrences, evaluated by time since pool entry, by policy, at a 25% immunosuppression acceptance rate, using a log scale.

Figure 4.24 illustrates the distribution of incompatible pairs that remain unmatched at the end of the simulation. Because the simulation terminates before these pairs can be matched or leave the pool, their waiting times are calculated as the number of days between their entry into the KEP pool and the simulation cutoff. Since the range of these accumulated values is considerably narrower than the matched and timeout distributions, the data is presented on a linear scale rather than requiring a logarithmic transformation for meaningful visualization. As shown, all three policies converge to near-zero unmatched pairs after 3 years since they entered the pool. However, distinct behavioral differences remain with Policy 3: Maximum Cardinality displaying a consistently higher and more persistent tail throughout the observation window, whereas the hierarchical policies—particularly Policy 1: Hierarchical with Waiting Time—achieve near-zero levels significantly earlier. The corresponding distributions for the remaining immunosuppressive acceptance rate scenarios (0%, 50%, 75%, and 100%) exhibit similar behavior and are provided in Appendix B.2.

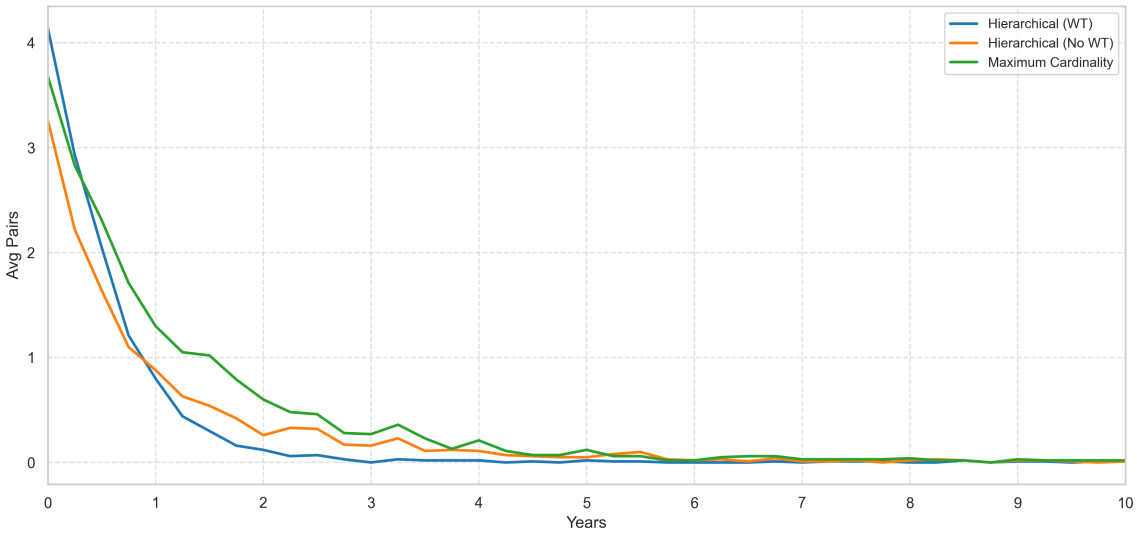

Figure 4.24: Average number incompatible pairs remaining unmatched at end of simulation, evaluated by time since pool entry, by policy, at a 25% immunosuppression acceptance rate.

As observed in both analyses (Figure 4.23 and Figure 4.24), Policy 1: Hierarchical with Waiting Time proved to be the best, reinforcing that compatibility-based prioritization combined with waiting time consideration is the most effective strategy for ensuring that pairs do not remain stranded in the pool indefinitely.

4.5.4 Computational Performance of Each Policy

Beyond clinical outcomes, computational cost is a practically relevant consideration for real-world KEP implementation, where match runs must be executed within operational time constraints. Table 4.2 reports the mean solve time per match run for each policy, broken down by simulation time horizons and immunosuppression acceptance rates.

Policy	Acceptance Rate (%)	6 Years (s)	12 Years (s)	24 Years (s)
Policy 1: Hierarchical with Waiting Time	0	0.0570	0.0770	0.0860
	25	0.0341	0.0373	0.0396
	50	0.0281	0.0290	0.0328
	75	0.0233	0.0257	0.0286
	100	0.0216	0.0242	0.0265
Policy 2: Hierarchical without Waiting Time	0	0.0236	0.0271	0.0290
	25	0.0184	0.0187	0.0196
	50	0.0157	0.0165	0.0179
	75	0.0153	0.0156	0.0169
	100	0.0149	0.0151	0.0164
Policy 3: Maximum Cardinality	0	0.0073	0.0082	0.0086
	25	0.0059	0.0059	0.0063
	50	0.0047	0.0050	0.0054
	75	0.0046	0.0049	0.0052
	100	0.0045	0.0048	0.0050
Policy 4: Static	0	35.5336	474.2291	—
	25	35.5410	474.2294	—
	50	35.5351	474.2304	—
	75	35.5321	474.2291	—
	100	35.5377	474.2327	—

Table 4.2: Mean solve time per match run for each policy, broken down by simulation time horizons and immunosuppression acceptance rates.

The results reveal fundamentally different computational profiles between dynamic and static approaches. All dynamic policies are extremely fast, with mean solve times under 0.09 seconds, confirming that per-period ILP solving imposes negligible computational burden. As expected from their structural complexity (Section 3.2), Policy 3: Maximum Cardinality is the fastest, as it solves a single ILP. Policy 2: Hierarchical without Waiting Time is approximately two to three times slower, and Policy 1: Hierarchical with Waiting Time is the slowest dynamic policy, with times roughly two to three times higher than Policy 2. These differences are consistent with the lexicographic structure described in Section 3.1.2 and Section 3.2: Policy 2 solves M sequential ILP with three priority levels

based on compatibility class, and Policy 1 solves additional levels within each compatibility class due to the waiting time sub-ordering, increasing the number of ILP iterations further.

The most notable pattern for the dynamic policies is the effect of the immunosuppression acceptance rate on solve time. Across all three dynamic policies, mean solve times decrease substantially as the acceptance rate increases, with the largest reductions occurring between 0% and 25%. For Policy 1: Hierarchical with Waiting Time under the 24-year simulation, for instance, the mean solve time drops from 0.0860 seconds at 0% acceptance to 0.0265 seconds at 100% acceptance — a reduction of approximately 69%. This occurs because higher acceptance rates reduce the number of half-compatible pairs remaining in the pool after each match run, producing a smaller and sparser compatibility graph, a reduced feasible set $\mathcal{C}$ of cycles and chains, and consequently a lighter ILP at each period.

Policy 4: Static presents a different computational profile, which results in computational times that are orders of magnitude higher than those of the dynamic policies: approximately 35.5 seconds for the 6-year simulation and approximately 474.2 seconds for the 12-year simulation—a 13-fold increase in solve time between the two horizons, despite only a two-fold increase in pool size — which reflects the exponential scaling typical of ILP problems as the feasible set $\mathcal{C}$ grows: a larger pool generates a disproportionately larger number of feasible cycles and chains, increasing the combinatorial complexity of the problem significantly. Despite these values, the 12-year solve time of approximately 474 seconds, although not exceeding the 900-second Gurobi time limit that was implemented, represents a substantial fraction of it. Therefore, by computing Policy 4 in the 24-year simulation, it exceeded the limit, resulting in a solution that would not reflect the best one possible for the pairs present. Furthermore, the acceptance rate does not affect static solve times, because all pairs enter the pool simultaneously at $t = 0$, executing a single matching run.

Overall, all dynamic policies mean solve times remain well below the 900 second Gurobi time limit, without a single match run reaching the timeout threshold across all of the 5,500 simulation runs. This confirms that the proposed lexicographic prioritization approach is computationally tractable for pool sizes of the scale considered in this study, delivering equity and efficiency gains with minimal computational cost, making it perfectly suited for real-world KEP operations.

5. Conclusion

This thesis's focus was to bridge the critical gap between theoretical mathematical optimization and the dynamic realities of clinical practice in Kidney Exchange Programs (KEPs). While traditional KEP algorithms have achieved remarkable success in maximizing the total number of transplants, this purely utilitarian approach unintentionally marginalizes the most vulnerable participants in the pool. Furthermore, traditional static models fail to capture the dynamic behavior of real-world KEPs, including the explanation for the departure of half-compatible pairs who possess the alternative option of a direct immunosuppressive transplantation. By addressing these challenges of algorithmic equity and dynamic pool behavior, this work provides a framework for designing fairer and more realistic KEPs that incorporate immunosuppressive treatment as part of their systems.

5.1. Summary of Findings and Contributions

Through the development of a comprehensive discrete-event simulation framework, this thesis modeled KEP operations over extended time horizons, explicitly incorporating the clinical decision mechanisms of half-compatible pairs who can opt to undergo direct, immunosuppressive transplantation. Across 5,500 simulation runs spanning four matching policies, five immunosuppressive acceptance rates, and three simulation horizons of 6, 12, and 24 years, we evaluated the performance of hierarchical lexicographic policies (Policy 1: Hierarchical with Waiting Time and Policy 2: Hierarchical without Waiting Time) against standard maximum cardinality matching (Policy 3: Maximum Cardinality and Policy 4: Static).

The results of this thesis support five main conclusions:

- The primary finding is that hierarchical lexicographic rules that prioritize difficult-to-match patients substantially improve outcomes for the most vulnerable pool participants. Policy 1: Hierarchical with Waiting Time and Policy 2: Hierarchical without Waiting Time provide a consistently higher number of Incompatible pairs Matched in the KEP (IMK) compared to Policy 3: Maximum Cardinality with a difference of almost 100 incompatible transplanted pairs between them in some scenarios for the 24-year simulation. More importantly, such results are achieved with only a marginal "price of fairness" in the overall number of All pairs Matched in the KEP

(AMK) with a difference of less than 50 transplanted pairs between them in every scenario for the 24-year simulation. Furthermore, the hierarchical policies even surpass Policy 4: Static for the number of IMK in both the 6- and 12-year simulations. This indicates that a dynamic policy specifically constrained in favor of equity can outperform the theoretical maximum cardinality for the most vulnerable patients—a result that arises because a pure maximum cardinality approach, by optimizing without any priority structure, sacrifices incompatible pairs to form easier and bigger cycles with easier-to-match pairs, whereas the hierarchical policies actively reserve exchange opportunities for incompatible pairs at every match run (see Section 4.4 for detailed discussion). This confirms that the choice of the matching algorithm is a policy decision with profound equity consequences for the most vulnerable pool participants.

- The second finding concerns the role of the waiting time tiebreaker. Comparing Policy 1 and Policy 2 isolates the contribution of the waiting time criterion within the hierarchical framework. While both policies achieve similar overall number of AMK, Policy 1 consistently outperforms Policy 2 for the hardest-to-match pairs, with the advantage growing over time as disparities in waiting time widen. The fact that Policy 2 appears to have lower waiting times is not an anomaly; it is simply the result of selection bias: Policy 2, without the waiting time tiebreaker, matches recently arrived easier-to-match incompatible pairs while leaving the longest-waiting, hard-to-match cases stranded. These pairs eventually depart the pool without contributing to the metric, resulting in Policy 2 having an higher number of incompatible pairs who left the pool without being matched compared to Policy 1. The top 10% worst-case analysis confirms that Policy 1 achieves the lowest extreme waiting times—the only metric that captures the actual impact on the most vulnerable patients.

- The third finding directly addresses one of the main concerns motivating this thesis. The concern was that the withdrawal of half-compatible pairs from the pool to pursue direct immunosuppressive transplantation could strip the pool of critical nodes and harm the incompatible pairs, but this does not hold true in practice. The number of IMK remains stable across all immunosuppressive acceptance rates from 0% to 100%, across both hierarchical policies and all simulation horizons. In fact, the presence of half-compatible pairs in the pool serves primarily to benefit those pairs themselves, granting them a critical opportunity to secure a match without resorting

to immunosuppressive treatment. For Policy 3: Maximum Cardinality, the number of IMK actually increases as the immunosuppressive acceptance rate rises—increasing from an average of 552 incompatible pairs with a 0% acceptance rate to 610 incompatible pairs with a 100% acceptance rate. Because the departure of half-compatible pairs from the pool reduces the available options, the policy is forced to utilize the remaining incompatible pairs to construct feasible cycles and chains. Therefore, instead of considering immunosuppression as a competing alternative that limits the pool, it should be recognized as a complementary pathway that efficiently directs half-compatible pairs to the treatment that best fits them while preserving KEP resources for those with no outside options.

- The fourth finding interprets the efficiency-equity trade-off at the system level. When KEP exchange transplants and immunosuppressive transplants are considered together as a unified system (i.e., all pairs transplanted are considered, ATO), the efficiency costs of the hierarchical prioritization are balanced out by the complementary immunosuppressive pathway. The equity-constrained policies, by directing KEP resources toward incompatible pairs and routing half-compatible pairs toward the immunosuppressive alternative, achieve superior system-level efficiency while simultaneously delivering far better outcomes for the most vulnerable patients.

- The fifth finding indicates that the proposed approaches are computationally viable. All dynamic policies achieve mean solve times well below one second per match run across all experimental scenarios. Therefore, the lexicographic prioritization policies impose no meaningful operational burden, removing one of the potential objections to their adoption in practice.

5.2. Clinical and Policy Implications

The findings of this thesis carry direct implications for the design and operation of real-world Kidney Exchange Programs, addressing both clinical practices and institutional policies.

The most direct clinical implication regards the treatment of half-compatible pairs. The results suggest that the departure of half-compatible pairs from the pool opting for immunosuppressive transplantation does not negatively impact incompatible pairs, and that the presence of half-compatible pairs in the pool serves primarily to benefit those pairs themselves by granting them a critical opportunity to secure a match without resorting to immunosuppressive treatment.

The second implication regards the design of matching algorithms in operational KEPs. The results strongly support the adoption of compatibility-based hierarchical prioritization with a waiting-time tiebreaker as the standard matching policy. The equity gains for incompatible pairs, blood type O patients, and highly sensitized patients are significant, the computational overhead is minimal, and the overall drop in efficiency is small, particularly when the immunosuppressive treatments are counted as part of total system output. These results are consistent with recent evidence from the Dutch KEP [15], where strict priority allocation algorithms for incompatible pairs increased their median transplant rates from 44% to 53%.

Finally, the results support a view of KEP and immunosuppressive transplantation not as competing alternatives, but rather as complementary parts of a complete transplantation system.

5.3. Limitations

It is important to acknowledge and consider a few limitations of this study when interpreting the results.

The most important limitation is our reliance on computer-generated data. All simulation instances are generated with tools used in the KEP literature rather than using real-world patient registries. While the distributional parameters are drawn from established and widely cited sources including Roth et al. [4] and Andersson and Kratz [7], real-world pools may exhibit correlations between patient characteristics, seasonal registration trends, or distinct hospital practices that our synthetic data generator cannot replicate. Therefore, validation against real registry data from existing KEPs would strengthen the real-world applicability of these results.

The second limitation concerns the immunosuppressive decision model. The acceptance decision is modeled as a fixed probability after each unmatched period, without capturing the natural differences between patients, such as their personal preferences, risk perceptions, clinical suitability, financial situation, and how waiting time influences their decision-making. In practice, the probability of accepting immunosuppression almost certainly increases with waiting time — patients who have been waiting for several years without a match are more likely to accept the available alternative than recently registered patients. Currently, the lack of real-world data on this type of decision prevents us from accurately incorporating this real-world behavior into our model.

5.4. Future Work

The findings and limitations of this thesis offer the following roadmap for future research in this field.

- The most direct next step is to incorporate a waiting-time-dependent probability for accepting immunosuppressive treatments. Replacing the fixed acceptance rate with a function that increases with waiting time would introduce a richer and more realistic behavioral model, and would more accurately reflect how the hierarchical prioritization policies interact with the immunosuppressive exit rate.

- Validation against real patient registry data is a second critical direction. Applying the simulation framework to real data from operational KEPs would enable a direct calibration of arrival rates, departure probabilities, and pairs' characteristics (such as blood type or PRA), ensuring the model's results closely simulate real-world clinical environments. This would significantly strengthen the external validity of our findings.

- A third research direction explores the quality of exchange matches. The present simulation optimizes exclusively for the number of transplants, treating all feasible matches as equivalent once compatibility between pairs is confirmed. However, in practice, not all compatible matches are equally beneficial: a younger donor, a lower degree of HLA mismatch, or a higher estimated post-transplant survival score may yield substantially better long-term graft outcomes compared to a standard compatible match. Incorporating a match quality score—derived from factors such as the age of the donor, HLA mismatch load, and Donor-Specific Antibodies (DSA) profiles—into the objective functions would allow the matching models to distinguish between a match that is merely feasible and one that is likely to produce durable long-term outcomes. This new variable could take different forms such as a weighted objective that simultaneously maximizes the number of transplants and aggregate match quality, or a minimum quality threshold below which a potential arc is excluded from the compatibility graph. The incorporation of this additional criterion would be a clinically meaningful extension to improve our simulation framework.

- A fourth possible next step introduces immunosuppressive therapy directly into the exchange structure itself, representing a fundamental shift from the current model

in which immunosuppression is treated exclusively as a unilateral alternative to exchange. In the current framework, a half-compatible pair that accepts immunosuppression exits the pool and undergoes transplantation with their own designated donor. An alternative scenario would allow these pairs to remain within a cycle or chain and have their patient receive a kidney from another pair's donor under an immunosuppressive protocol, rather than from their own, which is illustrated in Figure 5.1, where Pair 3 is the one receiving the kidney through immunosuppressive treatment. This would expand the set of feasible arcs in the compatibility graph by activating thin arcs—directed edges representing transplants that are viable only through clinical desensitization—between pairs who are compatible only if they undergo immunosuppressive therapy, in the style of the Kidney Exchange with Immunosuppressants (KEI) model proposed by Aziz et al. [48], the KEP with Half-Compatible Arcs (KEP-HCA) variant studied by Delorme et al. [23], and the ILP model proposed by Yang et al. [49]. Crucially, such an arc would require the explicit consent of the receiving patient, since they would bear the clinical consequences of the immunosuppressive treatment. However, patient consent would be more likely if the possible match offered meaningful advantages over transplantation with their own designated donor—for example, a younger donor, a better HLA match, or a higher estimated graft survival score. This creates a natural link with the match quality direction described above, in which the decision to accept immunosuppression within an exchange cycle would be governed not only by biological feasibility but by whether the quality of the proposed exchange sufficiently exceeds the quality of the transplant with their own donor to justify the extra clinical burden. Incorporating this joint decision—where patients accept immunosuppression only when it yields an improvement in match quality—within the dynamic simulation framework would represent a significant advance over existing models and would provide the first long-horizon evaluation of how KEP exchange mechanisms with immunosuppression interact with pool dynamics and prioritization policies over time.

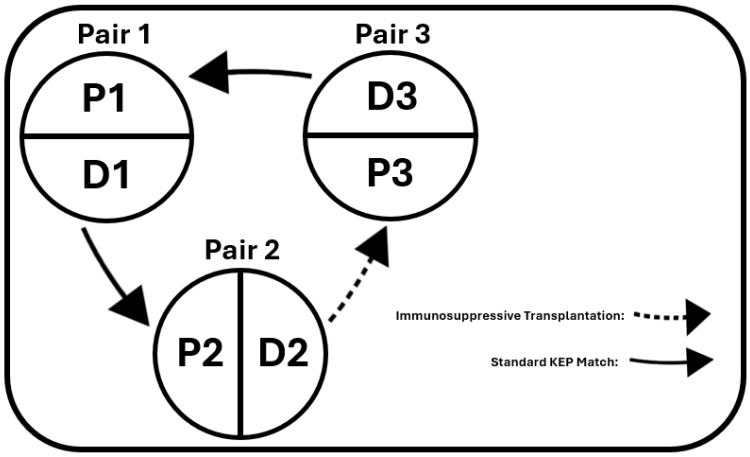

Figure 5.1: Cycle incorporating immunosuppressive transplantation, where Pair 3 receives a kidney via immunosuppressive therapy.

5.5. Final Remarks

Kidney Exchange Programs represent one of the most compelling examples of how rigorous mathematical optimization can be placed directly in service of human life. Every algorithmic decision incorporated into a KEP matching policy—which priority class is treated first, whether waiting time is considered, how immunosuppressive exits are handled—leads to real clinical outcomes for real patients. For an incompatible pair with no alternative option outside the KEP, these specific algorithmic rules, governing how matches are selected, are not merely technical design choices. They determine whether a transplant actually happens.

Historically, KEP algorithms have primarily maximized the total number of transplants. However, as these programs mature and hard-to-match pairs increasingly accumulate on waiting lists, it is no longer sufficient to treat fairness optimization as a secondary objective. As suggested throughout this thesis, the efficiency-equity tension in KEPs, while real, is not as severe as it first appears. When immunosuppressive transplantation is recognized as a complementary pathway for half-compatible pairs rather than a competing alternative, and when the matching algorithms are designed to save KEP resources for pairs who have no outside alternative, the apparent conflict between only maximizing the number of transplants and protecting the most vulnerable patients is substantially mitigated.

Particularly, hierarchical prioritization with waiting time consideration is simultaneously fairer, computationally tractable, and, at the system level (considering KEP matching and immunosuppressive treatment collectively), more efficient than unconstrained maximum cardinality matching. Therefore, the evidence analyzed in this thesis supports a clear and

practical recommendation for the designers of operational KEPs. Balancing efficiency and equity requires a shift in perspective in how we formulate our optimization models: programs must adopt compatibility-based hierarchical prioritization policies with a waiting time tiebreaker, and they must regard immunosuppressive transplantation as an essential component of a functioning system, rather than a threat to the integrity of the pool.

The evidence presented in this thesis indicates that by bridging the gap between Integer Linear Programming (ILP) and the dynamic realities of clinical practice, this work provides a computationally robust and ethical framework. Ultimately, it ensures that every patient who enters a KEP—regardless of blood type, sensitization level, or accumulated waiting time—is provided a fair and equitable chance to receive a life-saving transplant.

References

[1] Robert A. Wolfe, Valarie B. Ashby, Edgar L. Milford, Akinlolu O. Ojo, Robert E. Ettenger, Lawrence Y. C. Agodoa, Philip J. Held, and Friedrich K. Port. Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. *New England Journal of Medicine*, 341(23):1725–1730, 1999. doi:10.1056/NEJM199912023412303. URL https://www.nejm.org/doi/full/10.1056/NEJM199912023412303. [Cited on page 1.]

[2] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119:457–488, 2004. doi:10.1162/0033553041382157. [Cited on pages 2, 6, and 14.]

[3] Péter Biró, Bernadette Haase-Kromwijk, Tommy Andersson, Eyjólfur Ásgeirsson, Tatiana Baltesová, Ioannis Boletis, et al. Building kidney exchange programmes in europe-an overview of exchange practice and activities. *Transplantation*, 103(7): 1514–1522, 2019. doi:10.1097/TP.0000000000002432. [Cited on pages 2, 6, 7, 8, 9, 15, 20, and 21.]

[4] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review*, 97(3):828–851, 2007. doi:10.1257/aer.97.3.828. [Cited on pages 2, 6, 11, 12, 16, 19, 37, 38, and 71.]

[5] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Price of fairness in kidney exchange. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1013–1020, 2014. doi:10.65109/GDRR5031. [Cited on pages 2, 13, 19, and 20.]

[6] Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P. Dickerson, and Vincent Conitzer. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence*, 283:103261, 2020. ISSN 0004-3702. doi:10.1016/j.artint.2020.103261. URL https://www.sciencedirect.com/science/article/pii/S0004370220300229. [Cited on pages 2, 11, and 19.]

[7] Tommy Andersson and Jörgen Kratz. Pairwise kidney exchange over the blood group barrier. *The Review of Economic Studies*, 87(3):1091–1133, 05 2019. ISSN 0034-6527. doi:10.1093/restud/rdz018. URL https://doi.org/10.1093/restud/rdz018. [Cited on pages 3, 4, 14, 23, 26, 27, 38, 39, and 71.]

[8] Tiago Monteiro, Xenia Klimentova, João Pedro Pedroso, and Ana Viana. A comparison of matching algorithms for kidney exchange programs addressing waiting time. *Central European Journal of Operations Research*, 29(2):539–552, 2021. doi:10.1007/s10100-020-00680-y. [Cited on pages 4, 18, 20, 21, 27, 28, 38, and 39.]

[9] F. T. Rapaport. The case for a living emotionally related international kidney donor exchange registry. *Transplantation proceedings*, 18(3 Suppl 2):5–9, jun 1986. URL https://pubmed.ncbi.nlm.nih.gov/11649919/. [Cited on page 6.]

[10] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2):151–188, 2005. ISSN 0022-0531. doi:10.1016/j.jet.2005.04.004. URL https://www.sciencedirect.com/science/article/pii/S0022053105001055. [Cited on pages 6 and 15.]

[11] Paolo Ferrari et al. Kidney paired donation: principles, protocols and programs. *Nephrology, dialysis, transplantation : official publication of the European Dialysis and Transplant Association - European Renal Association*, 30(8):1276–1285, 2015. doi:10.1093/ndt/gfu309. [Cited on pages 6 and 8.]

[12] Ilse Duus Weinreich, Tommy Andersson, Margrét Birna Andrésdóttir, Mats Bengtsson, Alireza Biglarnia, Claus Bistrup, et al. Scandiatransplant exchange program (step): Development and results from an international kidney exchange program. *Transplantation Direct*, 9:1549, 2023. doi:10.1097/TXD.0000000000001549. [Cited on pages 7, 10, and 11.]

[13] Kristóf Druzsin, Péter Biró, Xenia Klimentova, and Rita Fleiner. Performance evaluation of national and international kidney exchange programmes with the ENCKEP simulator. *Central European Journal of Operations Research*, 32:923–943, 2024. doi:10.1007/s10100-024-00914-3. [Cited on pages 7 and 18.]

[14] Itai Ashlagi, Ágnes Cseh, David Manlove, Axel Ockenfels, and William Pettersson. Designing a kidney exchange program in germany: simulations and recommendations. *Central European Journal of Operations Research*, 33(4):1529–1555, 2025. doi:10.1007/s10100-024-00933-0. [Cited on pages 7 and 18.]

[15] Mattheüs F. Klaassen, Marry de Klerk, Marije C. Baas, Hanneke Bouwsma, Bungener Laura F., Maarten H. L. Christiaans, et al. Novel allocation strategies can boost kidney exchange programs: A monte carlo simulation. *Transplant International*, 39: 15423, 2026. doi:10.3389/ti.2026.15423. [Cited on pages 7, 8, 21, and 71.]

[16] Maxence Delorme, Sergio García, Jacek Gondzio, Jörg Kalcsics, David Manlove, and William Pettersson. New algorithms for hierarchical optimization in kidney exchange programs. *Operations Research*, 72(4):1654–1673, 2023. doi:10.1287/opre.2022.2374. [Cited on pages 7, 17, and 18.]

[17] Instituto Português do Sangue e da Transplantação. Circular normativa n.º 001/2022: Programa nacional de doação renal cruzada (pndrc) e programa internacional de doação renal cruzada (pidrc). Technical report, IPST, IP, Lisboa, Portugal, 2022. URL https://ipst.pt. Accessed 10 February 2026. [Cited on page 8.]

[18] Miquel Bofill, Marcos Calderón, Francesc Castro, Esteve Del Acebo, Pablo Delgado, Marc Garcia, Marta García, Marc Roig, María O. Valentín, and Mateu Villaret. The spanish kidney exchange model: Study of computation-based alternatives to the current procedure. In Annette ten Teije, Christian Popow, John H. Holmes, and Lucia Sacchi, editors, *Artificial Intelligence in Medicine*, volume 10259 of *Lecture Notes in Computer Science*, pages 272–277. Springer, 2017. ISBN 978-3-319-59758-4. doi:10.1007/978-3-319-59758-4_31. [Cited on page 9.]

[19] United Network for Organ Sharing. Kidney paired donation (kpd) operational guidelines and policy updates. Technical report, UNOS, Richmond, VA, 2021. URL https://unos.org. Accessed 10 February 2026. [Cited on page 9.]

[20] Mattheüs F. Klaassen, Marry de Klerk, Frank J. M. F. Dor, Sebastiaan Heidt, Stijn C. van de Laar, Robert C. Minnee, et al. Navigating a quandary in kidney exchange programs: A review of donor travel versus organ shipment. *Transplant International*, 38, 2025. ISSN 1432-2277. doi:10.3389/ti.2025.14804. [Cited on page 9.]

[21] Eun Jeong Heo, Sunghoon Hong, and Youngsub Chun. Efficient use of immunosuppressants for kidney transplants. *Journal of Health Economics*, 85:102650, 2022. ISSN 0167-6296. doi:10.1016/j.jhealeco.2022.102650. URL https://www.sciencedirect.com/science/article/pii/S0167629622000698. [Cited on pages 9, 22, 23, and 26.]

[22] JT Francisco, R Carvalho, J Freitas, MT Coimbra, et al. International kidney paired donation - the experience of a single center. *BJT*, 26(1):e3423, 2023. doi:10.53855/bjt.v26i1.531_ENG. [Cited on page 10.]

[23] Maxence Delorme, Wendy Liu, and David Manlove. Mathematical models and exact algorithms for kidney exchange problems with immunosuppressants. *INFORMS Journal on Computing*, 2026. doi:10.1287/ijoc.2024.1071. In press. [Cited on pages 11, 24, 26, and 73.]

[24] Dominique Bertrand, Philippe Gatault, Coralie Poulain, et al. Sensitization and deceased-donor kidney transplant access in france. *Kidney International Reports*, 11(4):106340, 2026. ISSN 2468-0249. doi:10.1016/j.ekir.2026.106340. [Cited on page 11.]

[25] David B. Leeser, Andrew G. Thomas, Alexander A. Shaffer, Jeffrey L. Veale, Alec B. Massie, Michael Cooper, et al. Patient and kidney allograft survival with national kidney paired donation. *Clinical Journal of the American Society of Nephrology*, 15 (2):228–237, 2020. doi:10.2215/CJN.06660619. [Cited on page 11.]

[26] S. U. Akgul, H. S. Ciftci, S. Temurhan, Y. Caliskan, A. Bayraktar, T. Tefik, I. A. Kaya, I. O. Canitez, E. Demir, H. Yazici, H. Bakkaloglu, A. E. Aydin, A. Turkmen, I. Nane, F. Aydin, and F. S. Oguz. Association between HLA antibodies and different sensitization events in renal transplant candidates. *Transplantation Proceedings*, 49(3): 425–429, apr 2017. doi:10.1016/j.transproceed.2017.02.004. [Cited on page 13.]

[27] Nizam Mamode, Oriol Bestard, Frans Claas, Lucrezia Furian, Siân Griffin, Christophe Legendre, Liset Pengel, and Maarten Naesens. European guideline for the management of kidney transplant patients with HLA antibodies: By the european society for organ transplantation working group. *Transplant International*, 35:10511, Aug 2022. doi:10.3389/ti.2022.10511. [Cited on pages 13 and 14.]

[28] Bart Smeulders, William Pettersson, Ana Viana, Tommy Andersson, Catarina Bolotinha, Pavel Chromy, Margherita Gentile, Karine Hadaya, Aline Hemke, Xenia Klimentova, Dirk Kuypers, David Manlove, Matthew Robb, Antonij Slavcev, Paolo Tubertini, María O. Valentin, Joris van de Klundert, and Paolo Ferrari. Data and optimisation requirements for kidney exchange programs. *Health Informatics Journal*, 27 (2), 2021. doi:10.1177/14604582211009918. PMID: 33878984. [Cited on pages 13, 14, and 25.]

[29] M. C. Bhaskaran, S. Heidt, and T. Muthukumar. Principles of virtual crossmatch testing for kidney transplantation. *Kidney International Reports*, 7(6):1179–1188, 2022. doi:10.1016/j.ekir.2022.03.006. [Cited on pages 13 and 14.]

[30] Y. Rocha, A. Jaramillo, J. Neumann, K. Hacke, E. Palou, and J. Torres. Crossmatch assays in transplantation: Physical or virtual?: A review. *Medicine*, 102(50):e36527, 2023. doi:10.1097/MD.0000000000036527. [Cited on pages 13 and 14.]

[31] Robert A. Montgomery, Bonnie E. Lonze, Karen E. King, Edward S. Kraus, Lauren M. Kucirka, Jayme E. Locke, et al. Desensitization in hla-incompatible kidney recipients and survival. *New England Journal of Medicine*, 365(4):318–326, 2011. doi:10.1056/NEJMoa1012376. [Cited on page 14.]

[32] Michael A. Rees, Jonathan E. Kopke, Ronald P. Pelletier, Dorry L. Segev, Matthew E. Rutter, Alfredo J. Fabrega, et al. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine*, 360(11):1096–1101, 2009. doi:10.1056/NEJMoa0803645. [Cited on page 15.]

[33] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. doi:10.4153/CJM-1965-045-4. [Cited on page 16.]

[34] David J. Abraham, Avrim Blum, and Tuomas Sandholm. Clearing algorithms for barter exchange markets: enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, pages 295–304, New York, NY, USA, 2007. Association for Computing Machinery. doi:10.1145/1250910.1250954. [Cited on page 16.]

[35] Miguel Constantino, Xenia Klimentova, Ana Viana, and Abdur Rais. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research*, 231(1):57–68, 2013. ISSN 0377-2217. doi:10.1016/j.ejor.2013.05.025. URL https://www.sciencedirect.com/science/article/pii/S0377221713004244. [Cited on page 16.]

[36] John P Dickerson, David F Manlove, Benjamin Plaut, Tuomas Sandholm, and James Trimble. Position-indexed formulations for kidney exchange. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 25–42, 2016. doi:10.1145/2940716.2940759. [Cited on pages 16 and 23.]

[37] Mathijs Barkel, Rachael Colley, Maxence Delorme, David Manlove, and William Pettersson. Operational research approaches and mathematical models for

kidney exchange: A literature survey and empirical evaluation. *European Journal of Operational Research*, 331(2):325–350, 2026. ISSN 0377-2217. doi:10.1016/j.ejor.2025.08.059. URL https://www.sciencedirect.com/science/article/pii/S0377221725006964. [Cited on page 17.]

[38] David F. Manlove and Gregg O'Malley. Paired and altruistic kidney donation in the uk: Algorithms and experimentation. *ACM J. Exp. Algorithmics*, 19, January 2015. ISSN 1084-6654. doi:10.1145/2670129. [Cited on pages 17, 18, and 20.]

[39] Ross Anderson, Itai Ashlagi, David Gamarnik, and Alvin E Roth. Finding long chains in kidney exchange using the traveling salesman problem. *Proceedings of the National Academy of Sciences*, 112(3):663–668, 2015. [Cited on page 17.]

[40] Lizeth Riascos-Álvarez, Merve Bodur, and Dionne Aleman. A branch-and-price algorithm enhanced by decision diagrams for the kidney exchange problem, 2023. [Cited on page 17.]

[41] M. Utku Ünver. Dynamic kidney exchange. *The Review of Economic Studies*, 77(1): 372–414, 2010. doi:10.1111/j.1467-937X.2009.00575.x. [Cited on page 18.]

[42] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yashodhan Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017. doi:10.1287/opre.2017.1644. [Cited on page 18.]

[43] Xenia Klimentova, Ana Viana, João Pedro Pedroso, and Nicolau Santos. Fairness models for multi-agent kidney exchange programmes. *Omega*, 102:102333, 2021. doi:10.1016/j.omega.2020.102333. [Cited on pages 18 and 21.]

[44] Duncan C. McElfresh and John P. Dickerson. Balancing lexicographic fairness and a utilitarian objective with application to kidney exchange. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, pages 1161–1168, 2018. doi:10.48550/arXiv.1702.08286. [Cited on pages 19 and 21.]

[45] D. Bertsimas, V. F. Farias, and N. Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013. URL http://www.jstor.org/stable/23482075. [Cited on page 20.]

[46] Mingrui Zhang, Xiaowu Dai, and Lexin Li. Fairness-aware kidney exchange and kidney paired donation, 2026. URL https://arxiv.org/abs/2503.06431. [Cited on pages 20 and 21.]

[47] Eun Jeong Heo, Sunghoon Hong, and Youngsub Chun. Kidney exchange with immunosuppressants. *Economic Theory*, 72:1–19, 2021. doi:10.1007/s00199-020-01276-y. [Cited on pages 23 and 26.]

[48] Haris Aziz, Ágnes Cseh, John P. Dickerson, and Duncan C. McElfresh. Optimal kidney exchange with immunosuppressants. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(1):21–29, 2021. doi:10.48550/arXiv.2103.02253. [Cited on pages 23, 26, and 73.]

[49] Cong Yang, Woonghee Tim Huh, Steven Shechter, and James Lan. The impact of organ blood type conversion on kidney paired donation: A simulation and optimization study. *American Journal of Transplantation*, 2026. ISSN 1600-6135. doi:10.1016/j.ajt.2026.04.013. In press. [Cited on pages 24 and 73.]

[50] Jennifer C Gander, Elisa J Gordon, and Rachel E Patzer. Decision aids to increase living donor kidney transplantation. *Current Transplantation Reports*, 4(1): 1–12, 2017. doi:10.1007/s40472-017-0133-1. [Cited on page 25.]

[51] A. D. Waterman, E. A. Schenk, A. C. Barrett, B. M. Waterman, et al. Incompatible kidney donor candidates' willingness to participate in donor-exchange and non-directed donation. *American Journal of Transplantation*, 6(7):1631–1638, July 2006. doi:10.1111/j.1600-6143.2006.01350.x. [Cited on page 25.]

[52] Supreet Sethi, Jua Choi, Mieko Toyoda, Ashley Vo, Alice Peng, and Stanley C. Jordan. Desensitization: Overcoming the immunologic barriers to transplantation. *Journal of Immunology Research*, 2017:6804678, 2017. doi:10.1155/2017/6804678. [Cited on pages 25 and 26.]

[53] Gurobi Optimization LLC. Gurobi optimizer reference manual [software], 2026. URL https://www.gurobi.com. Accessed 10 December 2025. [Cited on pages 33 and 42.]

A. Appendix A

A.1. Supplementary Active Pool Evolutions

This appendix provides the supplementary active pool evolution plots (total pool and incompatible pairs) over the 24-year simulation horizon for the 50%, 75%, and 100% immunosuppressive acceptance rate scenarios. These visualization plots support the discussion presented in Section 4.2, illustrating how the average pool size decreases as the acceptance rates increase.

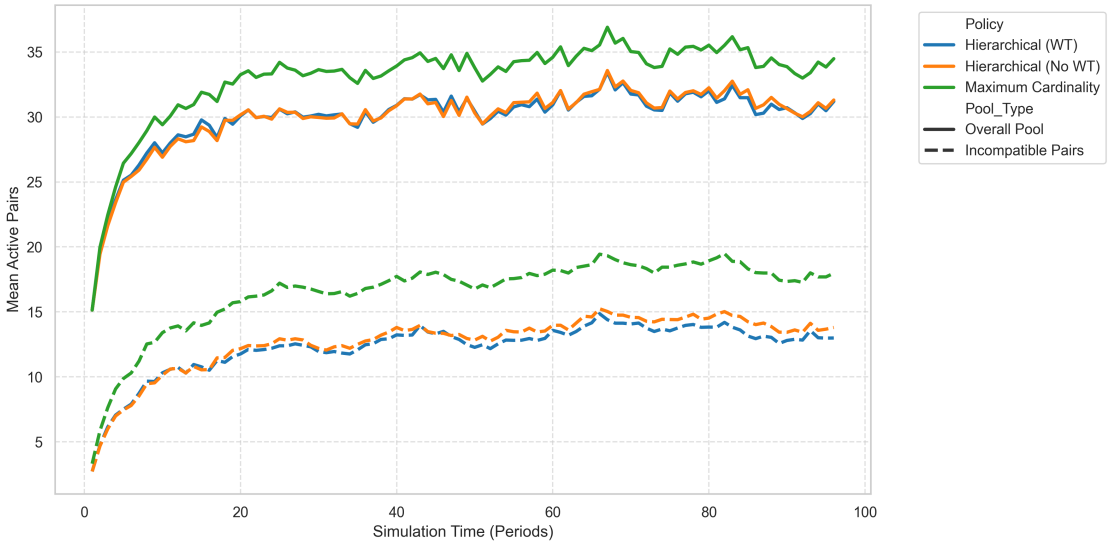

Figure A.1: Mean active pool size per period for each policy, under 50% immunosuppressive acceptance rate (solid lines represent the total number of pairs in the pool; dashed lines represent incompatible pairs).

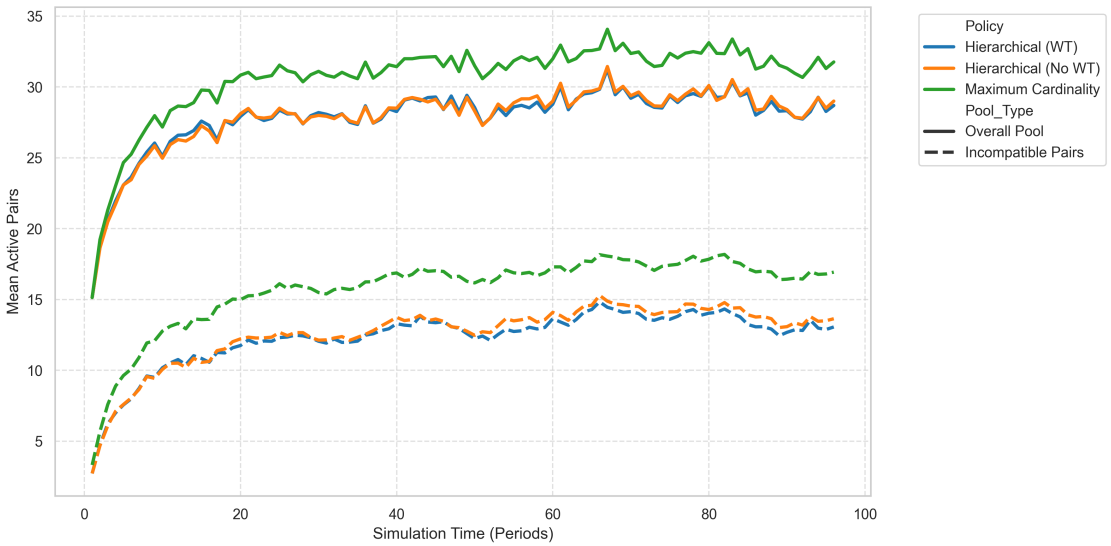

Figure A.2: Mean active pool size per period for each policy, under 75% immunosuppressive acceptance rate (solid lines represent the total number of pairs in the pool; dashed lines represent incompatible pairs).

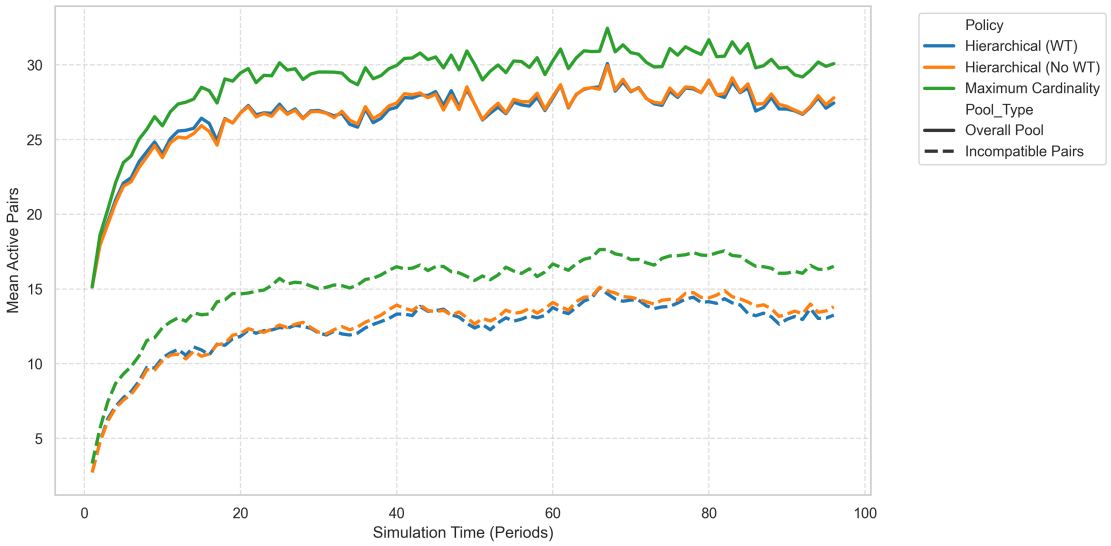

Figure A.3: Mean active pool size per period for each policy, under 100% immunosuppressive acceptance rate (solid lines represent the total number of pairs in the pool; dashed lines represent incompatible pairs).

B. Appendix B

B.1. Supplementary Distributions of Matched and Timeout Pairs

This appendix provides the supplementary distributions of matched and timeout pairs as function of time since pool entry across the remaining immunosuppressive acceptance rate scenarios (0%, 50%, 75%, and 100%). These plots complement the analysis presented in Section 4.5.3, demonstrating that the behavior of the policies remain structurally consistent regardless of the acceptance rate.

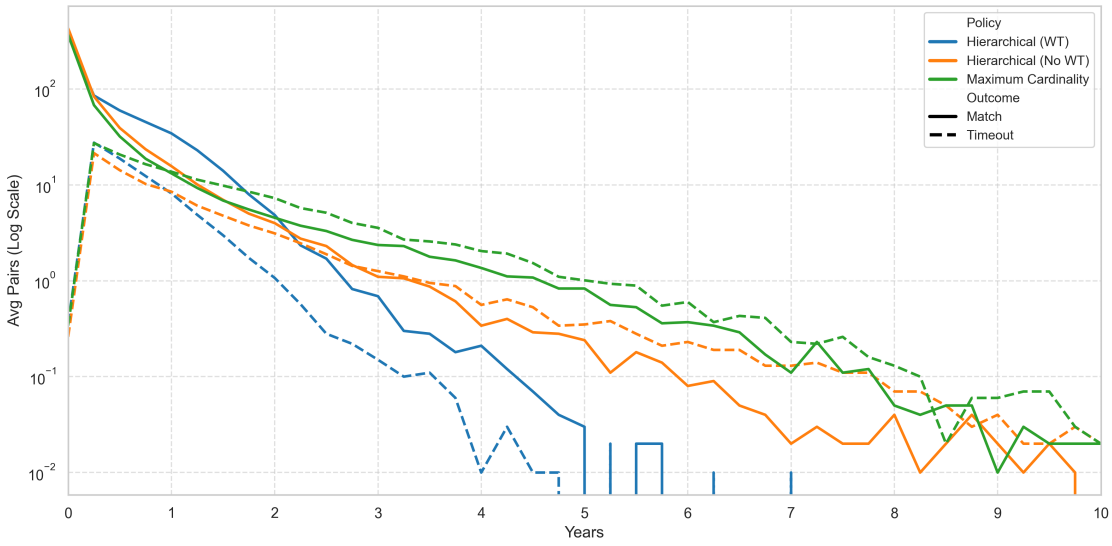

Figure B.1: Average number of Incompatible pairs Matched in the KEP (IMK) and timeout occurrences, evaluated by time since pool entry, by policy, at a 0% immunosuppression acceptance rate, using a log scale.

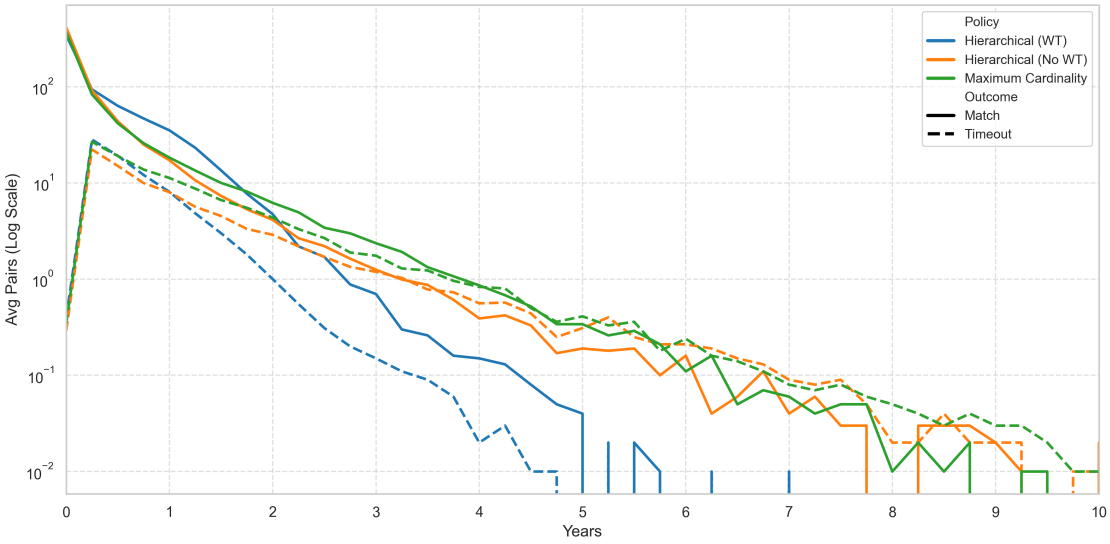

Figure B.2: Average number of Incompatible pairs Matched in the KEP (IMK) and timeout occurrences, evaluated by time since pool entry, by policy, at a 50% immunosuppression acceptance rate, using a log scale.

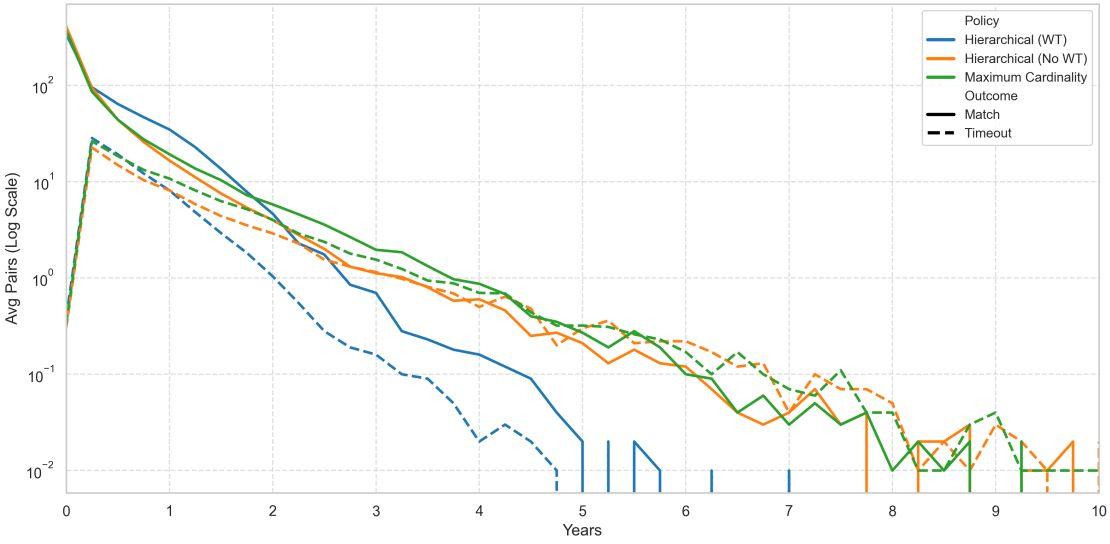

Figure B.3: Average number of Incompatible pairs Matched in the KEP (IMK) and timeout occurrences, evaluated by time since pool entry, by policy, at a 75% immunosuppression acceptance rate, using a log scale.

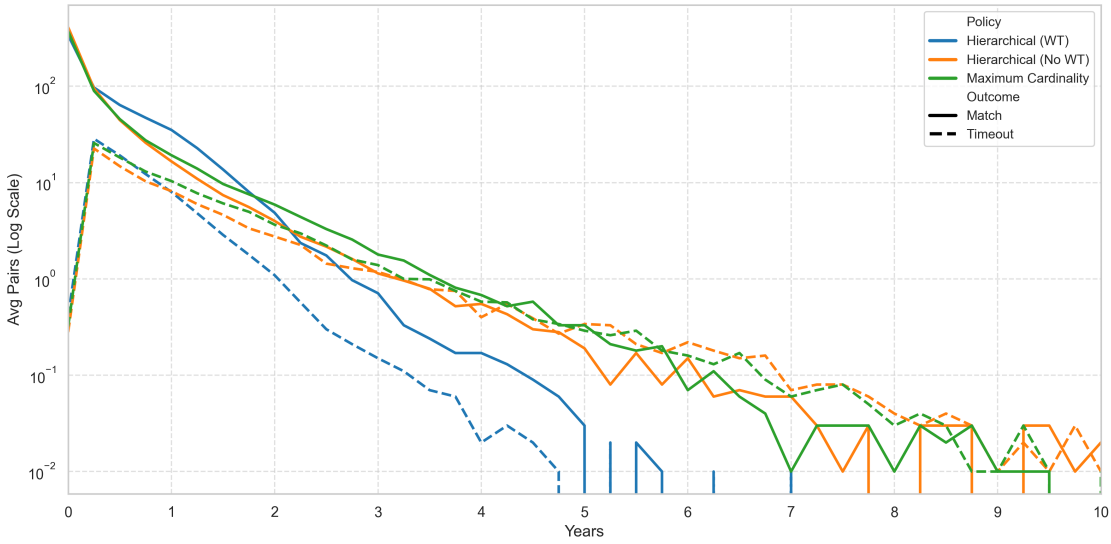

Figure B.4: Average number of Incompatible pairs Matched in the KEP (IMK) and timeout occurrences, evaluated by time since pool entry, by policy, at a 100% immunosuppression acceptance rate, using a log scale.

B.2. Supplementary Distributions of Unmatched Pairs

This appendix provides the supplementary distributions of pairs remaining unmatched at the end of the simulation across the remaining immunosuppressive acceptance rate scenarios (0%, 50%, 75%, and 100%). These plots complement the analysis presented in Section 4.5.3, demonstrating that the behavior of the policies remain structurally consistent regardless of the acceptance rate.

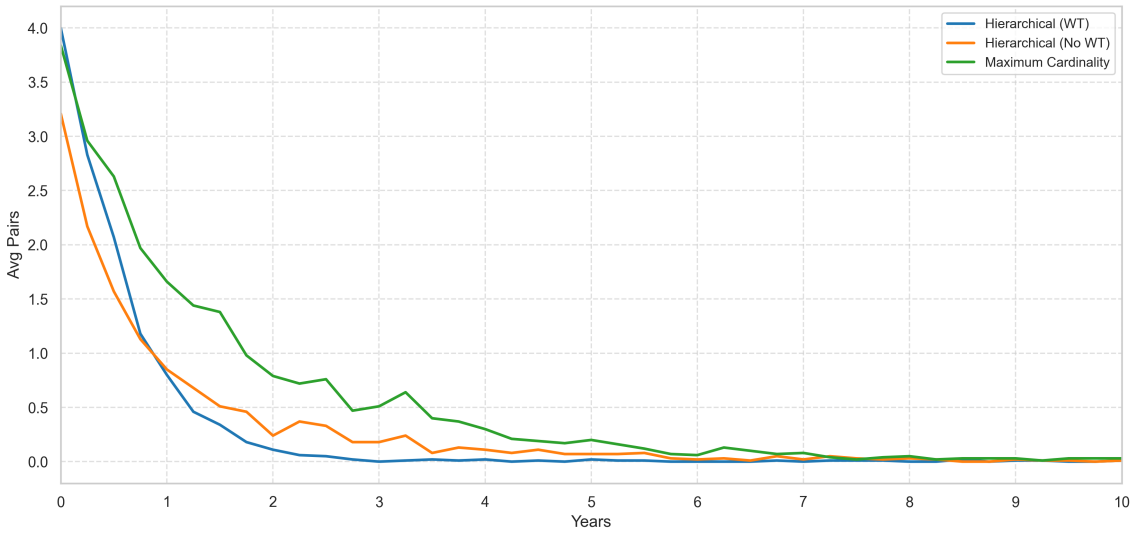

Figure B.5: Average number of incompatible pairs remaining unmatched at end of simulation, evaluated by time since pool entry, by policy, at a 0% immunosuppression acceptance rate.

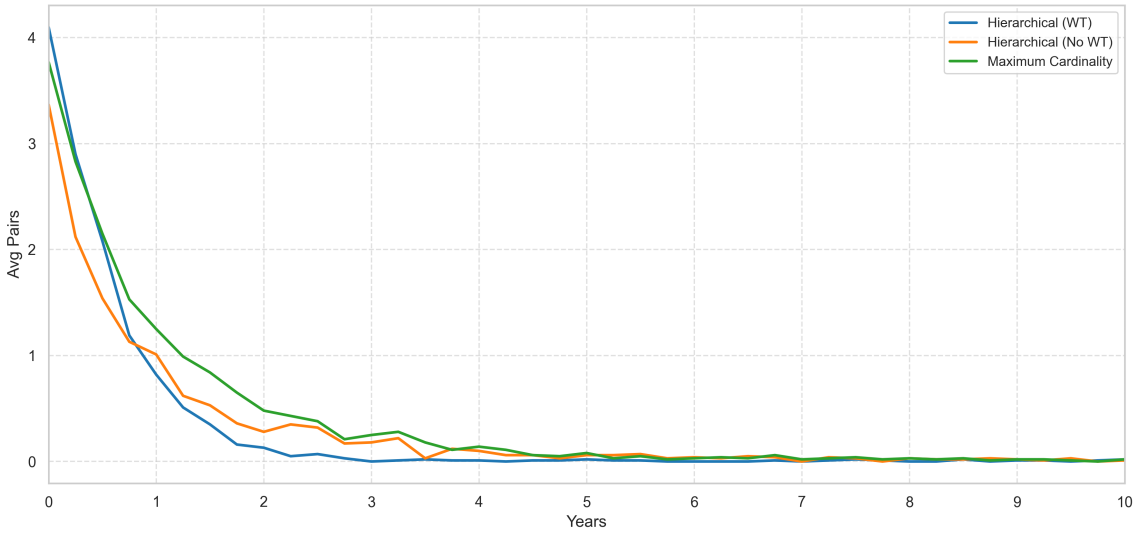

Figure B.6: Average number of incompatible pairs remaining unmatched at end of simulation, evaluated by time since pool entry, by policy, at a 50% immunosuppression acceptance rate.

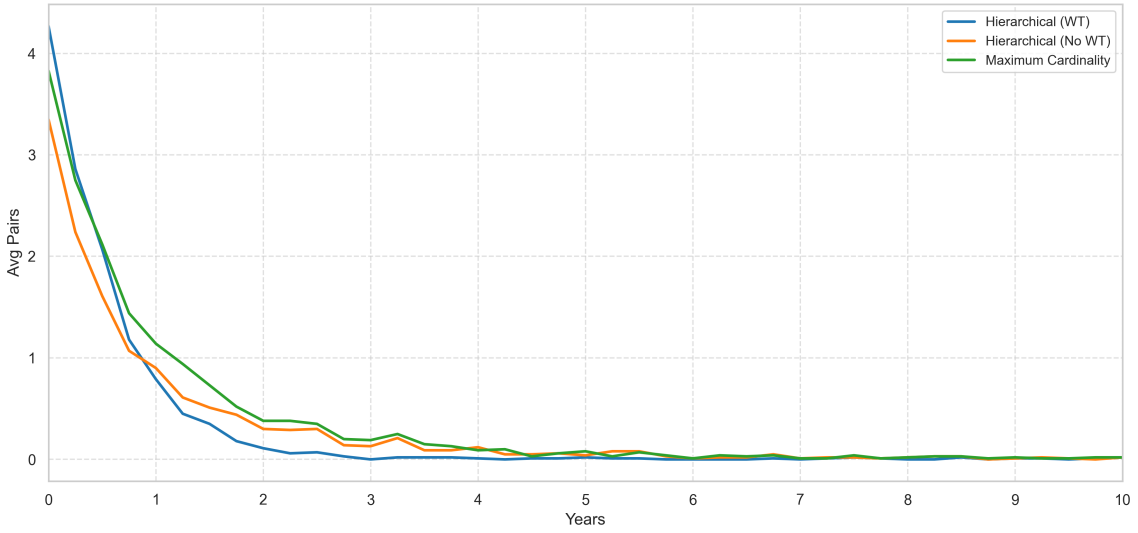

Figure B.7: Average number of incompatible pairs remaining unmatched at end of simulation, evaluated by time since pool entry, by policy, at a 75% immunosuppression acceptance rate.

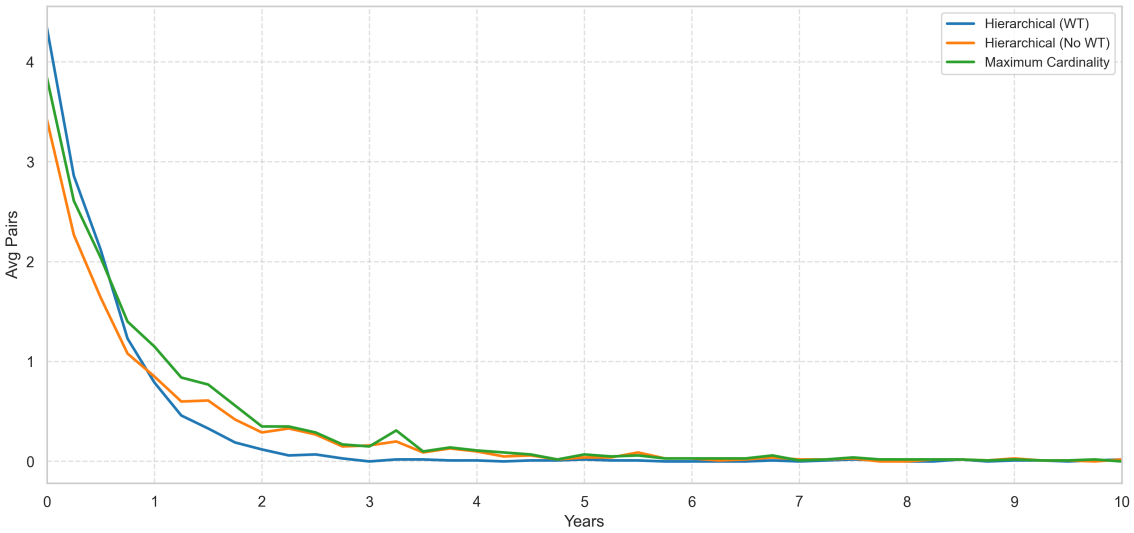

Figure B.8: Average number of incompatible pairs remaining unmatched at end of simulation, evaluated by time since pool entry, by policy, at a 100% immunosuppression acceptance rate.