



Thesis submitted in partial fulfillment of the requirements for the
degree of Master of Science in Management

AllSimEx: Simulating Dynamic Aspects of Organ Allocation

Ali Parsai

0601164

Academic year 2024-2025

Promotor: Sam VERBOVEN

Social Sciences & Solvay Business School

DECLARATION OF ORIGINALITY

The signed declaration of originality and publication is an integral component of the written work (e.g. bachelor or master thesis) submitted by the student.

With my signature I confirm that:

- I am the sole author of the written work here enclosed¹;
- I have compiled the written work in my own words;
- I have not committed plagiarism as defined in article 118 of the Teaching and examination regulations of the VUB; of which the most common forms of plagiarism are (non-exhaustive list):
 - form 1: to copy paragraphs/sentences of other authors with a reference, but without the use of appropriate quotation marks when needed;
 - form 2: to copy paragraphs/sentences of other authors, whether literally or not, without any sources;
 - form 3: to refer to primary sources where both text and references, whether literally or not, are copied from a non-referenced secondary source;
 - form 4: to copy paragraphs/sentences of other authors, whether with a reference or not, with only minimal or misleading adaptations.
 - form 5: to copy own previous work without proper citation (self-plagiarism).
 - form 6: to generate or substantially modify text fragments with software tools, without explicit mention of this
- I have referenced fully in the text and in the reference list any internet sources, published or unpublished works from which I have quoted or drawn reference;
- I have indicated clearly all text from which I have quoted literally;
- I have documented all methods, data and processes truthfully;
- I have not manipulated any data;
- I have mentioned all persons and organisations who were significant facilitators of the work, thus all the work submitted for assessment is my own unaided work except where I have explicitly stated otherwise;
- this work nor part of it has been submitted for a degree to any other programme, university or institution;
- I am aware that the work will be screened for plagiarism and the misuse of generative AI;
- I am aware that I need to disclose the use of generative AI in accordance with the relevant guidelines;
- I will submit all original research material immediately at the Dean's office when it is requested;
- I am aware that it is my responsibility to check when I am called for a hearing and to be available during the hearing period;
- I am aware of the disciplinary sanctions in case of irregularities as described in article 118 of the Teaching and examination regulations of the VUB;
- (Only for master theses) I agree to the publication of my master thesis. If I do not agree, I will attach the 'consent form for publicising master thesis' form together with my master thesis. (to be found on the VUB website).

Student's surname, first name: Parsai, Ali

Date: 04/08/2025

Signature: *Ali Parsai*

¹ For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.

Disclosure on the Use of AI Tools

During the preparation of this work, the author have used Google's Gemini and OpenAI's ChatGPT as tools in order to help find related work to the topic of the thesis, fix defects in the code of AllSimEx, find information about the topic of the thesis and related concepts, help with $\text{\LaTeX}$ code needed to compose the thesis document, charts, and formatting. In addition, AI was used to proof-read parts of the thesis. There are no verbatim AI text used in this thesis, and the thesis remains author's original work.

Acknowledgements

I would like to thank Dr. Sam Verboven and Alessandro Marchese for their invaluable insights, patience, and timely feedback for this thesis. The data used in the process of research for this thesis have been supplied by the United Network for Organ Sharing as the contractor for the Organ Procurement and Transplantation Network. The interpretation and reporting of these data are the responsibility of the author and in no way should be seen as an official policy of or interpretation by the OPTN or the U.S. Government. The training data for the models investigated in this thesis are based on OPTN data as of May 15th, 2025.

Ali Parsai

Antwerp, August 4, 2025

Management of organ transplantation is an important topic due to the scarcity of donated organs and the long wait lists across the world. Therefore, research into organ allocation policies is required in order to achieve an equitable outcome and increase the general life expectancy of the patients waiting for an organ. One of the problems facing this research is the lack of realistic simulators capable of assessing policies where degradation of the quality of the donated organ due to factors such as decision time are taken into account. In this study we present AllSimEx, an extended version of AllSim, a generic simulation platform for organ allocation. We implement DynamITE, a dynamic organ allocation policy, and compare the results of our experiment with those of DynamITE's original study to demonstrate the feasibility and reliability of creating a simulation platform capable of assessing such policies. Our results show that the experimental outcomes of AllSimEx are similar to those of DynamITE's original study.

Keywords: Organ Transplantation, Organ Allocation, Simulator

Het beheer van orgaantransplantaties is een belangrijk onderwerp vanwege de schaarste aan gedoneerde organen en de lange wachtlijsten wereldwijd. Daarom is onderzoek naar orgaantoewijzingsbeleid noodzakelijk om een rechtvaardige uitkomst te bereiken en de algemene levensverwachting van patiënten die op een orgaan wachten, te verhogen. Een van de problemen waarmee dit onderzoek kampt, is het gebrek aan realistische simulatoren die beleid kunnen beoordelen waarbij rekening wordt gehouden met de verslechtering van de kwaliteit van het gedoneerde orgaan door factoren zoals beslissingstijd. In deze studie presenteren we AllSimEx, een uitgebreide versie van AllSim, een generiek simulatieplatform voor orgaantoewijzing. We implementeren DynamITE, een dynamisch orgaantoewijzingsbeleid, en vergelijken de resultaten van ons experiment met die van de oorspronkelijke studie van DynamITE om de haalbaarheid en betrouwbaarheid aan te tonen van het creëren van een simulatieplatform dat dergelijk beleid kan beoordelen. Onze resultaten tonen aan dat de experimentele resultaten van AllSimEx vergelijkbaar zijn met die van de oorspronkelijke studie van DynamITE.

Trefwoorden: Orgaantransplantatie, Orgaantoewijzing, Simulator

La gestion de la transplantation d'organes est un sujet important en raison de la rareté des dons d'organes et des longues listes d'attente à travers le monde. Par conséquent, des recherches sur les politiques d'attribution d'organes sont nécessaires afin d'obtenir des résultats équitables et d'augmenter l'espérance de vie générale des patients en attente d'un organe. L'un des problèmes rencontrés par cette recherche est le manque de simulateurs réalistes capables d'évaluer les politiques prenant en compte la dégradation de la qualité des organes donnés due à des facteurs tels que le délai de décision. Dans cette étude, nous présentons AllSimEx, une version étendue d'AllSim, une plateforme de simulation générique pour l'attribution d'organes. Nous mettons en œuvre DynamITE, une politique d'attribution dynamique d'organes, et comparons les résultats de notre expérience à ceux de l'étude originale de DynamITE afin de démontrer la faisabilité et la fiabilité de la création d'une plateforme de simulation capable d'évaluer de telles politiques. Nos résultats montrent que les résultats expérimentaux d'AllSimEx sont similaires à ceux de l'étude originale de DynamITE.

Mots Clés: Transplantation d'organes, Attribution d'organes, Simulateur

	Contents
Disclosure on the Use of AI Tools	i
Acknowledgements	ii
Abstract	iii
Samenvatting	iv
Resumé	v
Contents	vi
List of Figures	viii
List of Tables	ix
1 Introduction	1
1.1 Context	1
1.2 Objectives	2
1.3 Research Questions	2
1.4 Thesis Structure	3
2 Background	4
2.1 Organ Transplantation	4
2.1.1 Static Aspects	4
2.1.2 Dynamic Aspects	4
2.1.2.1 Ischemic Time	4
2.1.2.2 Machine Perfusion Parameters	5
2.2 Organ Allocation	5
2.2.1 Treatment Effect	5
2.2.2 Individual Treatment Effect Estimation	6
2.2.3 Allocation Policies	7
2.2.4 Assessment Metrics	8
2.3 Machine Learning and Causal Inference	9
2.3.1 Counterfactual Regression	9
2.3.2 Generative Adversarial Networks	9
3 Related Work	11
4 Experimental Setup	13
4.1 Simulation Environment	13
4.2 Data	14
4.3 Allocation Policies	14
4.4 Simulation Parameters	16

<i>Contents</i>	vii
5 Results and Discussion	17
5.1 RQ1: Is it feasible to extend AllSim by incorporating dynamic organ allocation parameters?	17
5.2 RQ2: Can AllSimEx reliably simulate DynamITE?	17
5.3 Discussion	19
6 Threats to Validity	20
7 Conclusions and Future Work	21
Bibliography	22

List of Figures

5.1 AllSimEx results comparing the average ITE gained with a rolling window of 300 days 18

5.2 Average CIT and ITE gained per transplant for each policy 19

	List of Tables
4.1	Arrival Rates by Age Group 14
5.1	Experiment Overall Results 18

1.1 CONTEXT

Organs available for transplantation represent an scarce medical resource with European statistics revealing a persistent gap between supply and demand; 13,498 patients on the waiting list and 6,815 transplants reported in 2023 (*Eurotransplant Annual Report 2023*). A recent Dutch study analyzing 2,235 potential organ donors between 2018 and 2023 found that 617 donors (27.6%) were not utilized, meaning no organs were transplanted despite the donation process being initiated. The study documented a 51% increase in non-utilized donors from 83 cases in 2018 to 125 cases in 2023, while the percentage of utilized donors decreased from 77% to 70% during this period (Chotkan et al., 2025). This showcases the need for management of organ transplantation processes, including better research into the effectiveness of organ allocation policies.

Studying organ allocation is vital for effectively distributing organs given their scarcity. Individual treatment effect (ITE) is the difference in outcomes between receiving and not receiving a treatment for an individual (Imai et al., 2008). It is essential for researchers to be able to estimate ITE for evaluating organ allocation policies (Shalit et al., 2017). In the context of organ transplantation, ITE is usually measured by the increase in life-expectancy of the patient depending on the allocated organ compared to a non-allocation (Yoon et al., 2018). However, estimating ITE requires knowledge of counterfactual outcomes, namely what would have happened to a patient under an alternative scenario, which are inherently unobservable in real-world clinical data. To address this challenge, researchers train models where these counterfactuals are considered, and thus give accurate ITE estimates for a potential organ-patient match (Berrevoets, Jordon, et al., 2020). Using criteria such as MELD (Freeman, 2005) can determine who receives organs and when, directly affecting survival rates; e.g. prioritizing the sickest young patients with biliary atresia has reduced wait-list mortality in this group, showing how policy can affect outcome (Esmati et al., 2023). It is very important for researchers to be able to simulate the outcome of their policies in a reliable and replicable manner, so that different policies can effectively be compared to each other. To facilitate this, researchers use simulators that take advantage of the aforementioned models to replicate real-life scenarios, e.g. AllSim (Berrevoets, Jarrett, et al., 2023), simKAP (Y. Zhang et al., 2023), or COMET (Rose et al., 2024).

A critical limitation of many of the current approaches is their failure to incorporate the dynamic and time-sensitive aspects of organ viability—particularly the progressive degradation of organs due to cold ischemic time (CIT). The cold ischemic time (CIT) is “a period of time between harvesting an organ for transplant and its reperfusion just after implantation.” (Nečki et al., 2020) Prolonged CIT is associated with poorer graft function, increased rates of organ discard, and reduced patient survival, especially in liver and kidney transplants (Feng et al., 2006; Lozanovski et al., 2020). Effective preservation methods are still highly ineffective against CIT longer than 10 hours (O’Neill et al., 2020). Incorporating these dynamic aspects into organ allocation policies have been studied recently, showing improvements in total ITE in comparison with previous policies, e.g. DynamITE updates the state of an organ after each organ refusal

and improves both ITE and early acceptance rates of organs (Marchese et al., 2025).

1.2 OBJECTIVES

The main objective of this thesis is to introduce the ability to simulate dynamic aspects of organ transplantation, especially organ degradation, in a reliable and reproducible way. Current well-known simulators such as AllSim do not support this. They treat organ viability as static throughout the allocation process, failing to account for CIT. To achieve this objective we (i) extend a simulation framework by incorporating dynamic organ allocation factors, (ii) use extended simulation framework to replicate a policy with dynamic time-sensitive parameters, and (iii) demonstrate correctness of the experiment by comparing results to those of the original authors of the policy. By achieving this objective, we enable researchers to evaluate organ allocation policies with dynamic time-sensitive parameters using a well-known simulation tool, thus making implementation of innovative policies and comparative studies easier to achieve.

To achieve this objective in the scope of this thesis, we select AllSim as the simulator, and DynamITE as the policy to be evaluated. We build AllSimEx, an extended version of AllSim capable of dynamically updating the state of organs, so that it can incorporate dynamic allocation policies such as DynamITE, and check whether the results of our experiment agree with those of Marchese et al., 2025.

1.3 RESEARCH QUESTIONS

In this thesis, we use Goal-Question-Metric approach to formulate and analyze our research questions. This approach allows for a systematic link between high-level research goals to specific, measurable questions, and quantifiable metrics (Basili et al., 1994). For each research question, we provide the intended goal, the question to be answered, and the quantifiable metric by which the question can be answered.

RQ1

- **Goal.** Extending AllSim by incorporating dynamic organ allocation parameters demonstrates the feasibility of enhancing a generic simulation framework, originally designed to work with static parameters, to model dynamic parameters such as CIT.
- **Question.** Is it feasible to extend AllSim by incorporating dynamic organ allocation parameters?
- **Metric.** By implementing this capability in AllSimEx, we demonstrate feasibility through proof-by-construction, meaning we provide a working extension of the simulator that incorporates dynamic organ allocation parameters.

RQ2

- **Goal.** Use extended AllSim to replicate DynamITE as an example of an organ allocation policy with dynamic parameters, thus showing the applicability of the extended simulator to a state-of-the-art policy.
- **Question.** Can AllSimEx reliably simulate DynamITE?
- **Metric.** We use average ITE gained, average CIT, and organ non-use percentage to compare DynamITE with other policies.

1.4 THESIS STRUCTURE

The remainder of this thesis is organized as follows.

- Chapter 2 briefly explains the concepts needed to understand the rest of the thesis.
- Chapter 3 reviews the relevant literature and outlines the current state-of-the-art in areas directly related to the research questions addressed in this work.
- Chapter 4 describes the experimental design and methodology employed to investigate each of the research questions.
- Chapter 5 presents the outcomes of the experiment and answers to our research questions.
- Chapter 6 discusses threats to the validity of the study. It identifies possible biases, methodological constraints, and external factors that may affect the generalizability or reliability of the results. Each threat is considered in the context of how it was mitigated or remains a limitation of the work.
- Chapter 7 summarizes the main contributions of the thesis. It synthesizes the findings, revisits the thesis objectives, and reflects on the overall significance of the work within its broader context. It also explores possible extensions of this research. It outlines open questions, proposes directions for future investigation, and considers how the presented work could be built upon in subsequent studies.

2.1 ORGAN TRANSPLANTATION

Organ transplantation represents the definitive therapeutic intervention for patients with end-stage organ failure with current evidence pointing to significant improvements in both survival and quality-of-life outcomes (Jing et al., 2018; O’Neill et al., 2020). In this section, we explain the organ transplantation concepts related to this thesis, and those that are deliberately left out.

2.1.1 STATIC ASPECTS

Certain static characteristics of organ donors play a significant role in determining the quality and long-term function of transplanted organs. These attributes, which remain unchanged after organ procurement, include donor age, pre-existing comorbidities, and the cause of death. Donor age has been consistently associated with transplantation outcomes; for instance, older donor age correlates with increased risk of graft failure and reduced patient survival across various organ types, although the critical age threshold varies depending on the specific organ being transplanted (Chotkan et al., 2025; Dayoub et al., 2018). Donor comorbidities, such as chronic diseases or acute conditions at the time of death, can also negatively impact graft survival and recipient outcomes, as organs from donors with significant health issues may have diminished functional reserve (Schaapherder et al., 2022). Furthermore, the cause of donor death can influence organ quality, with certain causes, such as prolonged hypoxia or trauma, potentially leading to suboptimal graft function (Hessheimer et al., 2020). While these static donor factors are important, their impact may be less pronounced once an organ is deemed suitable for transplantation and accepted, as dynamic variables also contribute significantly to post-transplant outcomes. Nonetheless, careful assessment of these donor characteristics remains essential in optimizing transplant success and predicting long-term graft performance (Feng et al., 2006).

2.1.2 DYNAMIC ASPECTS

In this subsection, we examine two critical dynamic variables in organ transplantation: ischemic time and machine perfusion parameters. These factors significantly influence graft viability, organ quality, and recipient outcomes, including graft survival and patient recovery (Kalisvaart et al., 2021).

2.1.2.1 Ischemic Time

Ischemic time constitutes a critical determinant of transplant outcomes, comprising both warm and cold ischemic times. Cold preservation functions by slowing the metabolic activity of procured organs and marks the onset of cold ischemic time (CIT), with established thresholds varying significantly across organ types based on their metabolic demands and ischemia tolerance. Acceptable CIT limits typically range from 4 to 6 hours for the heart, less than 12 hours for the liver and pancreas, and up to 24 hours for the kidney, reflecting the differential susceptibility of various organs to ischemia-reperfusion injury (Jia et al., 2020; Moers et al.,

2009). Preservation time with static cold storage (SCS) is inherently limited, as prolonged cold storage increases the risk of early graft dysfunction and chronic complications, necessitating innovations in preservation technology and logistics (Jing et al., 2018). We use CIT as a dynamic measure of degradation of the organ after each rejected offering. While there are many components to prediction of graft failure, CIT has the advantage of being widely tracked and readily available in generic datasets and a good predictor of graft failure (Cesaretti et al., 2024).

2.1.2.2 Machine Perfusion Parameters

Machine perfusion has emerged as a pivotal advancement in organ preservation, offering dynamic control over the preservation environment and enabling real-time assessment of organ viability. Key parameters monitored during machine perfusion include perfusion pressure, flow rate, temperature, oxygenation, and biochemical markers such as lactate and pH, which collectively inform the metabolic state and functional integrity of the organ (Weeder et al., 2015). The optimization of these parameters is organ-specific and aims to minimize ischemia-reperfusion injury by maintaining adequate oxygen delivery and removing metabolic waste products. For example, hypothermic machine perfusion has demonstrated superiority over static cold storage in reducing delayed graft function in kidney transplantation, likely due to improved microcirculation and reduced cellular injury (Moers et al., 2009). Normothermic perfusion, which maintains organs at physiological temperatures, further allows for functional assessment and potential therapeutic interventions prior to transplantation (Nasralla et al., 2018). The integration of machine perfusion into clinical practice continues to evolve, with ongoing research focused on refining perfusion protocols and identifying reliable biomarkers to predict post-transplant outcomes.

These parameters are not yet incorporated into the extended AllSim simulator, despite the technical feasibility of doing so, in order to maintain a focused scope for this thesis. The primary limitation in including these additional parameters is the absence of sufficiently large and comprehensive real-world datasets required to train and validate the underlying models. Simulation platforms in organ transplantation often restrict the number of variables modeled to ensure credible interpretation and to avoid compounding uncertainties that arise from limited or incomplete data sources (Wood et al., 2024).

2.2 ORGAN ALLOCATION

This section provides background information on key organ allocation concepts. We cover treatment effect, which quantifies the benefit of transplantation. We also examine various allocation policies used in this thesis and the assessment metrics used to evaluate their effectiveness.

2.2.1 TREATMENT EFFECT

A treatment effect refers to the impact that a specific intervention or treatment has on an outcome of interest, often quantified as the difference in outcomes between individuals who receive the treatment and those who do not. In literature, this effect is typically interpreted as the causal effect of a treatment by comparing the observed outcome with the counterfactual outcome that would have occurred without the intervention (“Treatment Effects” 2005).

The average treatment effect (ATE) in organ transplantation quantifies the difference in expected outcomes—often measured in survival or quality-adjusted life years—between receiving a transplant and not receiving one, averaged across a relevant population. Accurately estimating ATE is crucial in transplantation research due to challenges such as non-random selection for transplantation, time-dependent treatment assignment, and patient heterogeneity. For example, a Swedish study utilizing inverse-probability-weighted regression adjustment found that

kidney transplantation, compared to dialysis, yielded an average survival benefit of 13.8 years, with the effect persisting across sexes and maintained even after accounting for bias from waiting times (S. Zhang et al., 2025). Similarly, methods like landmark analysis and partly conditional regression have been applied to liver transplantation to estimate ATE from observational data, offering insight into both the treated and control populations when treatment assignment is not randomized (Gong and Schaubel, 2017).

While ATE provides a valuable summary of the overall effectiveness of transplantation, it does not capture the variation in treatment effects across different patient profiles. This limitation motivates the use of the conditional average treatment effect (CATE), which examines how the benefit of transplantation differs across subgroups defined by characteristics such as age, comorbidities, or blood type. Understanding CATE allows researchers and clinicians to tailor treatment decisions according to patient-specific factors, identifying which subpopulations are more or less likely to benefit from a transplant. For instance, in kidney transplantation, stratified analyses have revealed meaningful heterogeneity in survival benefits based on clinical and demographic features (Kastoryano, 2022). Advanced statistical techniques, including risk modeling and landmark analysis, further support the estimation of CATE by incorporating rich covariate data collected before transplantation (H. J. Kim et al., 2019; Gong and Schaubel, 2017). These developments are crucial for aligning treatment strategies with individual patient needs and improving equity and efficiency in organ allocation.

Building on the idea of subgroup-specific effects, the concept of the individual treatment effect (ITE) pushes the analysis further by estimating the unique benefit or risk of transplantation for a single patient. ITE is particularly relevant in the era of precision medicine, where individualized predictions play a central role in clinical decision-making. In organ transplantation, this means estimating the likely outcome of accepting or declining a specific donor organ based on a recipient’s personal characteristics, including genetic compatibility, immunological profile, and comorbid conditions. Recent advances in machine learning and causal inference have enabled the development of decision-support tools capable of modeling these complex interactions. Notably, the OrganITE platform (Berrevoets, Jordon, et al., 2020) employs large-scale transplant registry data and advanced algorithms to estimate ITE at the level of individual candidates, supporting more nuanced and data-driven allocation decisions (W. Zhang et al., 2023). We use the inference model of OrganITE in our experiment to assess the effectiveness of a policy over the course of a simulation.

2.2.2 INDIVIDUAL TREATMENT EFFECT ESTIMATION

Estimating individual treatment effects (ITE) from observational data remains a core challenge in causal inference due to missing counterfactuals. Yoon et al., 2018 introduce GANITE, which addresses this by using generative adversarial networks to model unobserved outcomes. The framework consists of two components: a counterfactual generator and an ITE estimator, trained sequentially. This design enables the model to reduce selection bias and extend to settings with multiple treatments. Empirical evaluations demonstrate superior performance over existing ITE methods. While GANITE effectively reconstructs counterfactual outcomes, it operates on static input features and does not account for temporal dynamics that influence treatment effects. Building on the concept of ITE for organ allocation, Berrevoets, Jordon, et al., 2020 propose OrganITE, a method that matches patients to donor organs by jointly modeling compatibility and organ-type scarcity. The approach estimates ITE in a latent space learned through domain-adversarial training and prioritizes allocation strategies that maximize total life-years across the population. Their results show that accounting for organ rarity improves allocation outcomes compared to static policies. However, OrganITE assumes a static distribution of organ availability and does not address time-dependent degradation. Our work with DynamITE extends this framework by introducing CIT as a dynamic constraint. Unlike

OrganITE, which models long-term scarcity, DynamITE enables policy evaluation under conditions where organ viability degrades over time. This extension is critical for ensuring that ITE-based decisions remain valid under realistic clinical constraints.

Shalit et al., 2017 address assignment bias in ITE estimation by introducing a generalization error bound that depends on both predictive accuracy and the covariate shift between treated and control populations. Their Counterfactual Regression (CFR) framework learns balanced representations by minimizing prediction loss along with an integral probability metric between treatment groups. This approach reduces selection bias without requiring counterfactual generation. Unlike GANITE, which models counterfactuals directly via adversarial training, CFR enforces balance through distributional regularization. OrganITE builds on this line of work by incorporating domain-adversarial training to learn representations under allocation constraints, while also extending to settings where treatment involves matching patients with scarce resources rather than categorical labels. Compared to CFR, OrganITE integrates ITE estimation into a structured allocation problem, reflecting the real-world consequences of non-random treatment assignment.

2.2.3 ALLOCATION POLICIES

Organ allocation policies govern how scarce organs are distributed to patients on transplant waiting lists. Their design significantly impacts patient outcomes and the efficiency of transplant systems. Allocation policies are dynamic and regularly updated in response to medical, ethical, and technological developments. Changes are often informed by clinical evidence, registry data, and observed disparities such as organ discard rates or uneven access to continually improve equity, transparency, and patient survival (Ruf et al., 2022).

MELD and MELD-Na

The Model for End-Stage Liver Disease (MELD) and MELD-Na (MELD incorporating serum sodium) scores are widely utilized for liver allocation (Freeman, 2005). Figure 2.1 shows how we calculate MELD and MELD-Na in this thesis, taken from default values in AllSim (Berrevoets, Jarrett, et al., 2023; Brown et al., 2022)¹. These scores predict short-term mortality risk in patients with liver disease. Higher scores indicate greater medical urgency and prioritize patients for transplantation, thus aiming to reduce mortality among the critically ill. The MELD and MELD-Na systems provide an objective, standardized method for prioritizing liver transplant candidates. Their development addressed limitations of prior allocation methods, leading to a decrease of patient wait-list for the first time in 2002 (Kamath and R. W. Kim, 2007). However, recent studies cast doubt over the effectiveness of these policies, namely, the MELD-Na score, effective for predicting cirrhosis mortality, is less accurate as hepatitis C prevalence declines and nonalcoholic fatty liver disease and alcohol-associated liver disease rise, limiting its discriminative ability (Trivedi, 2022).

MaxAcc

The MaxAcc (Maximizing Acceptance) policy is a research-driven allocation strategy focused on reducing organ discards by statistically predicting which patients are most likely to accept an organ offer. By ranking candidates based on acceptance probability, this approach speeds up placement, shortens cold ischemic time, and decreases unnecessary organ discard. Marchese et al., 2024 suggest that MaxAcc can lower the number of unused organs and improve overall transplant rates without compromising patient outcomes, though it is typically evaluated in pilot settings rather than as a widely implemented national policy.

¹A handy tool to calculate both is available at <https://shorturl.at/RZy8S>

$$\begin{aligned}
MELD &= 3.79 \times \text{Total Bilirubin (mg/dl)} \\
&+ 11.2 \times \text{INR} \\
&+ 9.57 \times \text{Creatinine (mg/dl)} \\
&+ 6.43
\end{aligned} \tag{2.1}$$

$$\begin{aligned}
MELD\text{-}Na &= 1.59 \times (135 - \text{Sodium (mmol/l)}) \\
&+ MELD
\end{aligned} \tag{2.2}$$

2.1: MELD and MELD-Na Score Formulas

DynamITE

DynamITE is a state-of-the-art allocation policy that accounts for the evolving acceptance likelihood of candidates as offers are made. It uses real-time predictive models to dynamically update which patients are most likely to accept an organ, optimizing allocation for both quicker acceptance and maximized patient life expectancy. DynamITE has shown promising results by reducing cold ischemic times and increasing effective matches (Marchese et al., 2025).

2.2.4 ASSESSMENT METRICS

The performance of each policy is evaluated using three metrics. These are (i) Accumulated Individual Treatment Effect, a measure of clinical effectiveness; (ii) the Average Decision Time, a proxy for operational efficiency; (iii) and the Organ Non-use Percentage, which reflects resource utilization.

Accumulated Individual Treatment Effect

Accumulated Individual Treatment Effect quantifies the expected benefit for a transplant candidate of receiving a specific organ, compared to not receiving the transplant at that time. In organ allocation, summing ITEs across all allocations provides a metric for evaluating and comparing allocation strategies based on total population benefit, such as life-years gained. Recent advances use machine learning and causal inference methods to estimate individualized outcomes for transplant candidates, allowing allocation policies to optimize accumulated ITE. Empirical studies indicate these data-driven approaches can improve overall outcomes in organ allocation compared to traditional ranking systems (Papalexopoulos et al., 2024; Berrevoets, Jordon, et al., 2020; Deshpande, 2024).

Average Decision Time

Average Decision Time is an operational metric measuring the typical time elapsed from the moment an organ becomes available to the point a final allocation decision is made (either acceptance or rejection by candidates) (Schold, 2019). This metric captures the efficiency of allocation systems by averaging the time required for each offer cycle to result in a transplant or a discard. Shorter average decision times are linked to reduced cold ischemic time, better organ preservation, lower discard rates, and potentially improved patient outcomes. Prolonged decision times can arise from serial rejections, delays in offer response, or logistical issues, increasing the risk of non-utilization. Studies analyzing waiting times and offer cycles in national cohorts demonstrate that allocation policies designed to reduce average decision time—such as predictive recipient ranking or optimized offer sequencing—can decrease wait-list mortality and improve utilization (Schwab et al., 2024).

Organ Non-use Percentage

Organ non-use percentage measures the proportion of recovered organs that are discarded rather than transplanted, serving as a key indicator of system efficiency and resource utilization. This rate varies by organ type, with kidneys generally exhibiting lower non-use rates than hearts or lungs, but recent data show that in 2022, 19% of all recovered organs in the United States were not transplanted, a figure that has increased annually since 2018 (Israni et al., 2024). We simulate non-use percentage for MaxAcc and DynamITE. By default, MELD and MELD-Na do not measure organ non-use in AllSim.

2.3 MACHINE LEARNING AND CAUSAL INFERENCE

This section discusses the application of machine learning and causal inference methods in organ transplantation research. Specifically, we examine counterfactual regression, which estimates the causal effect of a treatment by modeling potential outcomes for individuals under both treatment and control conditions. We also explore generative adversarial networks (GANs), which create realistic, synthetic patient data to augment limited datasets, enabling more robust model training and evaluation. These approaches facilitate the robust estimation of ITEs and enhance organ allocation through synthetic data generation.

2.3.1 COUNTERFACTUAL REGRESSION

Counterfactual regression is a statistical method designed to estimate potential outcomes that remain unobserved in real-world scenarios, particularly in observational studies where treatment assignment is not randomized (Dickerman and Hernán, 2020). This approach addresses selection bias, which occurs when individuals receiving different treatments differ systematically in their characteristics, thereby compromising direct comparisons between treatment groups (Prosperi et al., 2020). The fundamental objective is to predict an individual’s potential outcome under alternative treatment conditions—the counterfactual scenario—which is essential for estimating ITEs that quantify outcome differences for a single individual across different treatments (Liu et al., 2024). The core challenge lies in the fact that only the factual outcome under the received treatment is observable, while all alternative treatment outcomes must be inferred through modeling techniques. In organ transplantation applications, counterfactual regression addresses the complex allocation problem where the “treatment” represents receipt of a specific donor organ and the “outcome” measures patient life expectancy. The OrganITE framework exemplifies this application by estimating patient-specific survival probabilities under different organ allocation scenarios (Berrevoets, Jordon, et al., 2020). Methods such as GANITE frame unobserved potential outcomes for different organs as missing data requiring statistical estimation, enabling optimization beyond simple donor-recipient matching algorithms. By predicting these counterfactual outcomes, transplantation models can optimize allocation decisions to maximize objectives such as accumulated ITE across a population.

2.3.2 GENERATIVE ADVERSARIAL NETWORKS

Generative Adversarial Networks (GANs) constitute a machine learning framework comprising two neural networks: the Generator and the Discriminator; trained simultaneously through competitive adversarial processes (Goodfellow et al., 2014). The generator starting with random input creates synthetic data that mimics the real data, while the discriminator works to differentiate between real and generated data. Through backpropagation, the discriminator’s classification provides a signal that the generator uses to update its weights, enabling the Generator to learn underlying data distributions and produce high-quality synthetic samples without explicitly defining probability density functions. This adversarial training process allows GANs to model complex data distributions effectively, particularly advantageous in high-dimensional

spaces where traditional density estimation approaches become computationally expensive. In counterfactual inference applications, GANs address the fundamental challenge of estimating unobserved treatment outcomes through specialized architectures. The GANITE method is motivated by the possibility that uncertainty in counterfactual distributions can be captured by attempting to learn them using a GAN, employing a Counterfactual Generator to produce plausible outcomes for untreated conditions and a Counterfactual Discriminator to identify original factual outcomes within completed datasets (Yoon et al., 2018). The generator observes the factual part of real data and creates the counterfactuals (missing part) conditioned on observed factual data, while the discriminator inputs the real dataset and tries to determine which part was actually observed. Related techniques such as domain adversarial training, utilized in OrganITE models, employ discriminator-like components trained to be unable to predict treatment types from learned representations, thereby removing treatment-assignment bias from model predictions (Berrevoets, Jordon, et al., 2020). Alternative approaches like OrganSync bypass GAN architectures entirely, instead utilizing synthetic control methods for counterfactual estimation in organ allocation scenarios (Berrevoets, Alaa, et al., 2021).

AllSim. AllSim (Berrevoets, Jarrett, et al., 2023) is a benchmarking environment for simulating and evaluating resource allocation policies in multi-user systems where users compete for scarce resources. It tackles a key limitation in policy evaluation: historical data is biased by the allocation policy active during its collection. AllSim addresses this by learning unbiased, modular components from the biased data, including distributions for users and resources, their respective arrival times, and a utility function. The framework is customizable, enabling users to modify system conditions, such as arrival rates or user types, to stress-test policies. To assess policies that differ from historical ones, AllSim uses counterfactual inference methods to predict potential outcomes, thereby correcting for allocation bias.

This approach marks an improvement over existing state-of-the-art evaluation methods, which are often limited to being either fully synthetic, which can inject bias, or unable to deviate from historical data, which precludes the testing of new policies. While other simulators may not support diverse users and resources or are confined to steady-state, single-user models, AllSim is explicitly designed for non-steady-state, and multi-user problems. By extracting realistic components from observational data and then generating new, counterfactual datasets based on a user’s modifications, AllSim allows for a more systematic and methodical pre-deployment evaluation of allocation policies compared to current approaches.

simKAP. The simKAP (Y. Zhang et al., 2023) framework is a simulation model designed to replicate the kidney allocation process with high fidelity by incorporating both dynamic waitlist modeling and a shared decision-making mechanism. It simulates three key phases: candidate list generation using Poisson processes to model arrival and departure, application of allocation rules based on national and state-level criteria, and recipient-donor matching influenced by probabilistic models of clinical decision-making. The shared decision-making component is novel in this context and factors in pediatric status, organ quality matching (via EPTS and KDRI), blood type compatibility, and geographic proximity. simKAP accommodates custom algorithms and allows incorporation of jurisdiction-specific rules, offering flexibility beyond existing models like SAM.

Compared to existing state-of-the-art models, simKAP introduces key enhancements in realism and adaptability. Unlike prior models that often omit acceptance behavior or assume static candidate pools, simKAP integrates recipient behavior and dynamic wait list turnover, resulting in simulated outcomes that more closely align with actual allocation patterns. The evaluation of three simulation configurations showed that only the full model (including dynamic lists and shared decision-making) accurately reproduced real-world metrics such as national-to-state allocation ratios and recipient wait times. Moreover, the framework enables in-silico evaluation of hypothetical risk-based allocation algorithms, which demonstrated improved post-transplant outcomes and equity for high-risk recipients. This positions simKAP as a robust tool for policy evaluation and optimization in organ allocation systems.

OASim. OASim (Tim Weaver, Josh Pyke, 2023) is a robust software framework for simulating the organ allocation system (OAS). It functions primarily as an agent-based simulation, modeling individual candidates and donors as they progress through a sequence of discrete events

within the system. The tool’s focus is on the allocation process that begins once candidates are listed and organs are recovered, with the match run (MR) being the central event it simulates. OASim incorporates Monte Carlo techniques to introduce stochasticity, allowing researchers to randomize elements like the arrival times of donors and candidates to create novel, hypothetical match runs for analysis. The software is designed to answer a wide range of research questions, from creating predictive data for unobserved scenarios to helping researchers understand the complex behaviors of the current system.

This framework builds upon and expands the capabilities of historical simulation tools, such as the simulated allocation models (SAMs). While past studies were often limited to backward-looking, counterfactual analyses of how different allocation rules would have performed in a historical period, OASim’s broader toolset supports more varied inquiries. A significant enhancement is the ability to model potential future data conditions, such as changing trends in organ donation or candidate listing rates, to make predictions about the system’s future state. The system also allows any of its component models, like the organ placement mechanism, to be varied or set to extreme values. This flexibility enables researchers to test a wider range of possibilities and gain more confidence in the validity of the simulation for making real-world inferences.

COMET. COMET (Rose et al., 2024) is a new framework for simulating organ allocation policies using an agent-based model that projects population-level outcomes by simulating the interactions of individually generated synthetic donors and candidates over time. A key feature is its use of data-driven probability models to create these synthetic populations, which allows for the exploration of hypothetical future trends in donor availability or candidate characteristics. The framework is designed to be modular and open-source, with distinct components for processes like donor generation, matching, and risk assessment, which facilitates flexible adaptation and collaborative development.

COMET addresses significant limitations of the current state-of-the-art, the Simulated Allocation Models (SAMs), which rely on direct sampling of historical registry data. This reliance restricts the ability of SAMs to forecast the effects of evolving clinical practices or shifting population demographics, a gap COMET fills by using adaptable synthetic populations. Furthermore, COMET’s open-source nature and independence from sensitive data aim to broaden researcher access. In a validation study, COMET-Lung accurately predicted a significant drop in transplants for type O lung candidates following a major policy change, an outcome that the existing TSAM model failed to anticipate, demonstrating COMET’s potential for more accurate and robust policy evaluation.

Experimental Setup

This chapter provides a comprehensive and step-by-step description of the experimental setup used to compare DynamITE with other policies within AllSimEx.

4.1 SIMULATION ENVIRONMENT

The simulation is implemented using the `AllSimEx` framework, which builds upon AllSim framework (Berrevoets, Jarrett, et al., 2023). It provides a modular structure for modeling complex resource allocation systems.

Overall Simulation Flow: The central `Sim` class orchestrates the simulation over a defined time horizon T . In each time step t , the simulation performs the following sequence:

1. Remove patients whose simulated time-to-live (TTL) has expired.
2. Sample new patients arriving in the wait list.
3. Sample new organs becoming available, and increment the organ counter.
4. The current allocation `policy` is invoked via its `select` method to choose a recipient for the available organ.
5. For each successful transplant, the organ transplant counter is incremented. Patient and organ data, along with relevant outcomes like time-to-live and Individual Treatment Effect gained are logged.
6. All patients currently on the wait list are “aged” by one day, decrementing their TTL.

Patient and Organ Arrival System: Patient arrivals are modeled using a Poisson process (Kingman, 2002). Donors and Patients arrive based on distributions defined for 4 different age groups, similar to AllSim experiment (Berrevoets, Jarrett, et al., 2023). These groups are (i) under 30, (ii) between 30 and 45, (iii) between 45 and 60, (iv) and over 60. The rate $\lambda_X(t)$ at which new patients arrive at time t is explicitly defined per group (Table 4.1).

Upon arrival, the covariates of new patients and new donors are sampled from Kernel Density Estimators (KDE), respectively `KDE_x` and `KDE_o`, which are fitted to a subset of the historical patient dataset. As no conditioning bins are applied during its initialization, these KDEs represents the overall population’s distribution. Organ quality is a dynamic factor. If a patient rejects an organ offer, its Cold Ischemic Time (CIT) increases, reflecting the passage of time and potential degradation. The duration a patient takes to accept or reject an offer (Estimated Decision Time) is modeled using an exponential distribution. The distribution has a scale parameter of 0.5 hours, which skews probabilities towards shorter decision times. The sampled value is truncated to a maximum of 2 hours. This mathematical model for decision time is represented as:

$$\text{Estimated Decision Time} = \min(\text{Exponential Distribution}(\lambda = \frac{1}{\text{Scale}}), 2.0), \text{Scale} = 0.5$$

Table 4.1: Arrival Rates by Age Group

#	Description	Lambda Function
1	Under 30	$\lambda_1(t) = \frac{3}{1 + e^{-\frac{t-450}{150}}}$
2	Between 30 and 45	$\lambda_2(t) = \frac{2}{1 + e^{-\frac{t-350}{150}}}$
3	Between 45 and 60	$\lambda_3(t) = \frac{2}{1 + e^{-\frac{t-150}{150}}} + \frac{2}{1 + e^{-\frac{t-650}{100}}}$
4	Over 60	$\lambda_4(t) = \frac{1}{1 + e^{-\frac{t-150}{150}}} + \frac{1}{1 + e^{-\frac{t-650}{100}}}$

When a rejection occurs, the organ’s CIT covariate (which is tracked in hours within the dataset) is directly incremented by Estimated Decision Time. This accumulation of CIT is crucial as it negatively impacts both subsequent acceptance probabilities and the ITE.

4.2 DATA

The data for the training of the inference model is extracted from UNOS dataset based on OPTN data as of May 15th, 2025.

Core Patient and Organ Data: The main simulation loop uses real patient and organ characteristics from the UNOS dataset, using `UNOSDataModule` supplied with AllSim. This module handles necessary preprocessing steps. This results in distinct patient covariate sets, which include prefixed names like `INIT_BILIRUBIN`, `INIT_AGE` and organ covariate sets, including `AGE_DON`, `COLD_ISCH`, `BMI_DON_CALC`.

Synthetic Outcomes: To allow for controlled experimentation and the calculation of counterfactuals (like ITE), the transplant outcomes are generated synthetically rather than using observed real-world outcomes. This enables the evaluation of policies against a known ground truth. A dedicated synthetic dataset is utilized for training the ITE model. This dataset is independent of the main simulation’s patient/organ sampling. It includes combined patient and organ features and a binary acceptance label ($A \in \{0, 1\}$), indicating whether a patient accepted a given organ.

Two key machine learning models are trained and integrated into the simulation to enable advanced policy evaluation:

OrganITE serves as the causal inference engine, predicting the potential life expectancy of a patient under different scenarios. Specifically, for a patient x and organ o , it predicts $Y^O(x, o)$ (life expectancy if they receive organ r) and $Y^0(x)$ (life expectancy if they do not receive an organ). The network’s input dimension is the combined size of patient and organ covariates. The Individual Treatment Effect (ITE), $\tau(x, o)$, is derived as the difference between these two predicted outcomes:

$$\tau(x, o) = Y^O(x, o) - Y^0(x)$$

4.3 ALLOCATION POLICIES

Five distinct organ allocation policies are simulated and compared to understand their performance across various metrics. All policies are implemented as classes inheriting from the `Policy` abstract base class.

- **DynamITE (Dynamic Individual Treatment Effect):** DynamITE aims to maximize the total ITE gained across all transplants, while explicitly accounting for patient acceptance behavior and organ deterioration (CIT accumulation) during the offer process. For each newly available organ, DynamITE first identifies all patients currently on the wait list. For each patient x on the wait list, DynaScore is computed with respect to the current state of the organ o . This score is a composite metric defined as:

$$\text{DynaScore}(x, o) = \frac{P(\text{Accept}|x, o)^{\alpha_{acc}} \cdot \tau(x, o)^{\alpha_{ite}}}{(T_k(x, o) + \epsilon)^{\alpha_{dec_time}}}$$

- $P(\text{Accept}|x, o)$: The acceptance probability for patient x and organ o .
- $\tau(x, o)$: The Individual Treatment Effect (ITE) gained from patient x receiving organ o , calculated as $Y^O(x, o) - Y^0(x)$.
- $T_k(x, o)$: The estimated decision time (in hours) for patient x and organ o .
- ϵ : A small constant (10^{-6}) added to the denominator of the decision time term to prevent division by zero or numerical instability for very small decision times.
- $\alpha_{acc}, \alpha_{ite}, \alpha_{dec_time}$: Tunable weighting coefficients for acceptance probability, ITE, and decision time, respectively. These are set as $\alpha_{acc} = 0.6$, $\alpha_{ite} = 0.7$, and $\alpha_{dec_time} = 1.0$ in this experiment.

After computing DynaScores for all eligible patients, the patients are sorted in descending order based on these scores. Organs are then offered sequentially to patients following this DynaScore-ranked order.

- If a patient accepts (based on a random draw against $P(\text{Accept}|x, o)$), that patient receives the organ, and the organ is removed from list.
- If a patient rejects, the organ’s CIT is increased by the predicted decision time $T_k(x, o)$, simulating deterioration. The deteriorated organ is then offered to the next patient in the list.
- This process continues until the organ is accepted or all ranked patients have been offered and rejected. If the organ is not accepted, it is marked as non-used.
- **MaxAcc:** MaxAcc sorts the patients based on their acceptance probability, aiming to increase the chance of a transplant for an arriving organ.
- **MELD:** A traditional urgency-based policy that ranks patients solely on their MELD score. This score is calculated from patient covariates using the standard MELD formula (Equation 2.1). Organs are offered to the highest-MELD score patient on the waitlist. If rejected, it proceeds to the next highest MELD score.
- **MELD-Na:** An extension of the MELD policy that enhances patient ranking by incorporating serum sodium levels into the MELD score calculation. The formula is a modification of the standard MELD score (Equation 2.2). Similar to MELD, offers proceed down the MELD-Na-ranked list.
- **FIFO (First-In, First-Out):** A simple baseline policy that offers organs strictly based on the order patients joined the wait list (chronological order of entry). The organ is offered to the patient who has been on the wait list the longest. If rejected, the organ’s CIT increases, and it is offered to the next patient in chronological order of wait list entry. This policy does not perform any complex matching or prioritization based on medical factors or acceptance probabilities.

4.4 SIMULATION PARAMETERS

Each simulation run spans $T = 1021$ days, providing a sufficient period to observe long-term policy performance and system dynamics. Each simulation starts with an initial wait list population of 100 patients to establish an initial state for the system. These patients are assigned negative IDs (e.g., $-1, -2, \dots$) to distinguish them from new arrivals during the simulation, which receive positive IDs based on the time step. Policy performance is comprehensively assessed using the following metrics, which are logged and aggregated during simulation:

- **Total ITE Gained:** This is a primary outcome metric. It represents the cumulative sum of $\tau(x, o)$ for all successfully transplanted patient-organ pairs across the entire simulation duration. It quantifies the total ITE gained due to the allocation policy’s decisions.
- **Average CIT:** This measures the mean CIT (in hours) of organs at the moment they are successfully transplanted. It provides insight into the quality of organs being placed and the efficiency of the offer process, as higher CIT generally correlates with poorer outcomes.
- **Total Non-use Percentage:** This critical metric quantifies the efficiency of organ utilization. It is calculated as the percentage of sampled organs that were ultimately not transplanted (i.e., discarded after being rejected by all patients in wait list).

In this chapter, we present the results of our experiment with regards to our research questions.

5.1 RQ1: IS IT FEASIBLE TO EXTEND ALLSIM BY INCORPORATING DYNAMIC ORGAN ALLOCATION PARAMETERS?

Motivation. We set out to extend AllSim by incorporating dynamic organ allocation parameters, in order to demonstrate the feasibility of enhancing a generic simulation framework originally designed to work with static parameters to model dynamic parameters such as CIT.

Approach. Implementation of AllSimEx demonstrates the feasibility through proof-by-construction. We take the current implementation of AllSim, and extend its simulation engine and policy templates so that it is able to handle dynamic organ allocation policies. In particular, selection of a patient from the wait list given an incoming organ is extended so that it updates the quality of the organ after each decision.

Findings. The fact that AllSimEx was capable of performing the DynamITE experiment, a dynamic organ allocation policy, shows that it is indeed feasible to model dynamic parameters in a generic simulation framework. Implementation of metrics such as ITE gain in the simulation engine also demonstrates the extendability of the framework for simulating other, more complex dynamic organ allocation policies in the future.

5.2 RQ2: CAN ALLSIMEX RELIABLY SIMULATE DYNAMITE?

Motivation. We set out to use AllSimEx to replicate DynamITE as an example of an organ allocation policy with dynamic parameters, thus showing the applicability of the extended simulator to a state-of-the-art policy. We design an experiment with parameters similar to those of Marchese et al., 2025, so that we are able to compare the results. This demonstrates the sanity of the system, showing the reliability of AllSimEx in simulating dynamic organ allocation policies.

Approach. We chose the parameters of the experiment as close as possible to those of Marchese et al., 2025, in order to make sure the results of these experiments remain comparable. Even though the underlying data is newer than Marchese et al., 2025, and the simulation engine and the organ/patient arrival distributions are different, given the fact that DynamITE takes ITE into consideration, we expect the results to be similar to the original study given the metrics we use. We calculate total ITE gained, average CIT, and total non-use percentage for each policy. Then we compare these results to the original study to see if there are any meaningful differences.

Findings. Figure 5.1 shows the average ITE gained for each policy with a rolling window of 300 days. Predictably, DynamITE has the highest gained ITE over the course of the simulation. The only exception is at the end, where FIFO policy performs better. This is mostly due to the fact that the patients generated near the end of the experiment are of the younger age and the donors of the older age due to the particular Poisson process defined for the experiment, where

the FIFO policy accidentally selects good candidates. There is no significant advantage to using MaxAcc over other policies, even though it performs better than both MELD and MELD-Na in the middle of the simulated period. The surprising result is the fact that MELD-Na performed worse than MELD in this simulation, the reasons for which are not within the scope of this study.





Figure 5.1: AllSimEx results comparing the average ITE gained with a rolling window of 300 days

Table 5.1 presents the results of the experiment. DynamITE performs significantly better than other policies in total ITE gained. This is not surprising, given that DynaScore optimizes for estimated ITE gained before offering an organ to a patient in the list. DynamITE and MaxAcc perform better than others in average CIT, although their results are very close to each other. DynamITE performs better on total non-use percentage than MaxAcc, even though MaxAcc attempts to optimize for maximum acceptance. For other policies, the simulation does not degrade the quality of the organ after each offer, leading to non-use outcome becoming extremely unlikely.


Table 5.1: Experiment Overall Results

Metric (Unit)	DynamITE	MaxAcc	MELD	MELD-Na	FIFO
Total ITE Gained (Years)	9983.12	8158.60	7931.08	8105.11	8695.22
Average CIT (Hours)	4.71	4.74	4.96	5.02	4.92
Total Non-use Percentage (%)	2.32%	2.70%	0.0%	0.0%	0.0%

Figure 5.2 shows the result of the simulation per transplant. This clearly shows that per each transplant, DynamITE increases the ITE by the highest amount compared to other policies. The high value of ITE gain per transplant for the FIFO policy is explained by the distribution of arriving organs and patients as explained in the previous section.



(a) Average ITE gained per transplant per policy



(b) Average CIT per transplant per policy

Figure 5.2: Average CIT and ITE gained per transplant for each policy

5.3 DISCUSSION

Out of 5 policies simulated in this experiment, DynamITE performs best according to the measured metrics. When we compare the results to those of Marchese et al., 2025, it is clear that DynamITE is performing as expected, thus demonstrating the reliability of AllSimEx. The observed underperformance of MaxAcc can be attributed to the fact that MaxAcc in AllSimEx does take into account the degradation of the quality of the organ after each offer, meanwhile MaxAcc in Marchese et al., 2025 only uses the initial values of the organ. This explains why MaxAcc performed worse than DynamITE in non-use percentage, given that it does not update its ranking while the simulation degrades the quality of the organ. Another observation is the underperformance of MELD-Na when compared to MELD. In Marchese et al., 2025 these two policies performed within the same general bounds. However, this is not the case in our experiment. Given that the results of Marchese et al., 2025 have been performed on many iterations of KDE trainings while AllSimEx relied on the same KDE for all the experiments, there is a chance that over many iterations these two policies converge in results.

Given the general similarity of the results, the outcome of the experiment points to AllSimEx as a viable simulation framework for further research into organ allocation policies with dynamic parameters.

INTERNAL VALIDITY

This refers to the degree to which the observed outcomes can be attributed to the factors under study, rather than to confounding variables or alternative explanations. In order to assess whether the results of this study compare fairly with those of DynamITE, we chose the parameters of the experiment as close as possible to the original, while the underlying implementation of the simulation framework differs. In addition, the newer dataset used in this work and different process of patient/organ arrival, still results in similarities between the two studies, thus showing the setting of the experiment did not affect its objective outcome.

EXTERNAL VALIDITY

External validity concerns the extent to which our findings can be generalized beyond the specific context of the study. It remains to be seen whether we can implement the same mechanism for updating the state of an organ during the simulation, however the proof of feasibility shows it is theoretically possible to do so. The degree to which these results are generalizable to other simulation frameworks relies significantly on their implementation specifications.

Conclusions and Future Work

This study describes AllSimEx, an extension of the AllSim simulation framework modified to incorporate dynamic organ allocation policies. The work aimed to demonstrate the feasibility of this extension and validate the new framework’s reliability. Feasibility was established by proof-by-construction. The successful implementation of AllSimEx and its capacity to model time-dependent factors like Cold Ischemic Time (CIT) confirm that a generic simulator can be effectively adapted for dynamic policy evaluation. The framework’s reliability was evaluated by simulating five allocation policies, including the time-sensitive DynamITE policy. As hypothesized, DynamITE yielded the highest total ITE gain (9983.12 years) and the lowest average CIT value (4.71 hours). These outcomes are consistent with DynamITE’s previously known results and demonstrate the correctness of AllSimEx simulator. In summary, AllSimEx is a validated framework that enables the modeling and comparison of organ allocation policies with dynamic parameters. It provides a functional tool for future research into developing more efficient, time-sensitive allocation strategies. There are two immediate directions where the current research can be extended. First, is to attempt implementation of more complex organ allocation policies where more parameters are evaluated dynamically (e.g. machine perfusion parameters and warm ischemic time), and second, more complex policy templates where organ-patient pairing state-space is explored using state-of-the-art methods such as reinforcement learning algorithms.

- Basili, V. R., Caldiera, G., and Rombach, D. H. (1994). “The Goal Question Metric Approach”. In: *Encyclopedia of Software Engineering*. Vol. 1. John Wiley & Sons.
- Berrevoets, J., Alaa, A., Qian, Z., Jordon, J., Gimson, A. E. S., and van der Schaar, M. (July 18, 2021). “Learning Queueing Policies for Organ Transplantation Allocation Using Interpretable Counterfactual Survival Analysis”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by M. Meila and T. Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, pp. 792–802. URL: <https://proceedings.mlr.press/v139/berrevoets21a.html>.
- Berrevoets, J., Jarrett, D., Chan, A., and van der Schaar, M. (2023). “AllSim: Simulating and Benchmarking Resource Allocation Policies in Multi-User Systems”. In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., pp. 851–866. URL: https://proceedings.neurips.cc/paper_files/paper/2023/file/0296e17ec30fc36007edaaa2f96b5f17-Paper-Datasets_and_Benchmarks.pdf.
- Berrevoets, J., Jordon, J., Bica, I., gimson, a., and van der Schaar, M. (2020). “OrganITE: Optimal Transplant Donor Organ Offering Using an Individual Treatment Effect”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., pp. 20037–20050. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/e7c573c14a09b84f6b7782ce3965f335-Paper.pdf.
- Brown, C., Aksan, N., and Muir, A. J. (Nov. 2022). “MELD-Na Accurately Predicts 6-Month Mortality in Patients With Decompensated Cirrhosis: Potential Trigger for Hospice Referral”. In: *Journal of Clinical Gastroenterology* 56.10, pp. 902–907. ISSN: 0192-0790. DOI: 10.1097/mcg.0000000000001642. URL: <https://journals.lww.com/10.1097/MCG.0000000000001642>.
- Cesaretti, M., Izzo, A., Pellegrino, R. A., Galli, A., and Mavrothalassitis, O. (June 27, 2024). “Cold Ischemia Time in Liver Transplantation: An Overview”. In: *World Journal of Hepatology* 16.6, pp. 883–890. ISSN: 1948-5182. DOI: 10.4254/wjh.v16.i6.883. URL: <https://www.wjgnet.com/1948-5182/full/v16/i6/883.htm>.
- Chotkan, K. A., Kuiper, M. A., Alwayn, I. P. J., Heemskerk, M. B. A., Braat, A. E., and Jansen, N. E. (Mar. 11, 2025). “Analysis of Unused Organ Donors in the Netherlands: Older Donor Age Associated With Higher Risk of Non-Utilization”. In: *Transplant International* 38, p. 14157. ISSN: 1432-2277. DOI: 10.3389/ti.2025.14157. URL: <https://www.frontierspartnerships.org/articles/10.3389/ti.2025.14157/full>.
- Dayoub, J. C., Cortese, F., Anžič, A., Grum, T., and De Magalhães, J. P. (Sept. 2018). “The Effects of Donor Age on Organ Transplants: A Review and Implications for Aging Research”. In: *Experimental Gerontology* 110, pp. 230–240. ISSN: 05315565. DOI: 10.1016/j.exger.2018.06.019. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0531556518301864>.
- Deshpande, R. (Feb. 28, 2024). “Smart Match: Revolutionizing Organ Allocation through Artificial Intelligence”. In: *Frontiers in Artificial Intelligence* 7. ISSN: 2624-8212. DOI: 10.3389/

- frai.2024.1364149. URL: <https://www.frontiersin.org/articles/10.3389/frai.2024.1364149/full>.
- Dickerman, B. A. and Hernán, M. A. (July 2020). “Counterfactual Prediction Is Not Only for Causal Inference”. In: *European Journal of Epidemiology* 35.7, pp. 615–617. ISSN: 0393-2990, 1573-7284. DOI: [10.1007/s10654-020-00659-8](https://doi.org/10.1007/s10654-020-00659-8). URL: <http://link.springer.com/10.1007/s10654-020-00659-8>.
- Esmati, H. et al. (Feb. 2023). “Waitlist Mortality of Young Patients with Biliary Atresia: Impact of Allocation Policy and Living Donor Liver Transplantation”. In: *Liver Transplantation* 29.2, pp. 157–163. ISSN: 1527-6465. DOI: [10.1002/lt.26529](https://doi.org/10.1002/lt.26529). URL: <https://journals.lww.com/10.1002/lt.26529>.
- Eurotransplant Annual Report* (June 1, 2023). Eurotransplant. URL: https://www.eurotransplant.org/wp-content/uploads/2024/06/ETP_AR2023_HighRes.pdf.
- Feng, S., Goodrich, N., Bragg-Gresham, J., Dykstra, D., Punch, J., DebRoy, M., Greenstein, S., and Merion, R. (Apr. 2006). “Characteristics Associated with Liver Graft Failure: The Concept of a Donor Risk Index”. In: *American Journal of Transplantation* 6.4, pp. 783–790. ISSN: 16006135. DOI: [10.1111/j.1600-6143.2006.01242.x](https://doi.org/10.1111/j.1600-6143.2006.01242.x). URL: <https://linkinghub.elsevier.com/retrieve/pii/S160061352202651X>.
- Freeman, R. B. (Jan. 2005). “MELD: The Holy Grail of Organ Allocation?” In: *Journal of Hepatology* 42.1, pp. 16–20. ISSN: 01688278. DOI: [10.1016/j.jhep.2004.11.002](https://doi.org/10.1016/j.jhep.2004.11.002). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168827804005161>.
- Gong, Q. and Schaubel, D. E. (Mar. 1, 2017). “Estimating the Average Treatment Effect on Survival Based on Observational Data and Using Partly Conditional Modeling”. In: *Biometrics* 73.1, pp. 134–144. ISSN: 0006-341X, 1541-0420. DOI: [10.1111/biom.12542](https://doi.org/10.1111/biom.12542). URL: <https://academic.oup.com/biometrics/article/73/1/134-144/7537737>.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). “Generative Adversarial Nets”. In: *Advances in Neural Information Processing Systems*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger. Vol. 27. Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2014/file/f033ed80deb0234979a61f95710dbe25-Paper.pdf.
- Hessheimer, A. J. et al. (Aug. 2020). “Donation after Circulatory Death Liver Transplantation: Consensus Statements from the Spanish Liver Transplantation Society”. In: *Transplant International* 33.8, pp. 902–916. ISSN: 0934-0874, 1432-2277. DOI: [10.1111/tri.13619](https://doi.org/10.1111/tri.13619). URL: <https://onlinelibrary.wiley.com/doi/10.1111/tri.13619>.
- Imai, K., King, G., and Stuart, E. A. (Apr. 1, 2008). “Misunderstandings Between Experimentalists and Observationalists about Causal Inference”. In: *Journal of the Royal Statistical Society Series A: Statistics in Society* 171.2, pp. 481–502. ISSN: 0964-1998, 1467-985X. DOI: [10.1111/j.1467-985X.2007.00527.x](https://doi.org/10.1111/j.1467-985X.2007.00527.x). URL: <https://academic.oup.com/jrjsssa/article/171/2/481/7084435>.
- Israni, A. K., Zaun, D. A., Gauntt, K., Schaffhausen, C. R., Lozano, C., McKinney, W. T., Miller, J. M., and Snyder, J. J. (Feb. 2024). “OPTN/SRTR 2022 Annual Data Report: Deceased Organ Donation”. In: *American Journal of Transplantation* 24.2, S457–S488. ISSN: 16006135. DOI: [10.1016/j.ajt.2024.01.018](https://doi.org/10.1016/j.ajt.2024.01.018). URL: <https://linkinghub.elsevier.com/retrieve/pii/S1600613524000832>.
- Jia, J., Nie, Y., Li, J., Xie, H., Zhou, L., Yu, J., and Zheng, S.-S. (May 12, 2020). “A Systematic Review and Meta-Analysis of Machine Perfusion vs. Static Cold Storage of Liver Allografts on Liver Transplantation Outcomes: The Future Direction of Graft Preservation”. In: *Frontiers in Medicine* 7, p. 135. ISSN: 2296-858X. DOI: [10.3389/fmed.2020.00135](https://doi.org/10.3389/fmed.2020.00135). URL: <https://www.frontiersin.org/article/10.3389/fmed.2020.00135/full>.
- Jing, L., Yao, L., Zhao, M., Peng, L.-p., and Liu, M. (May 2018). “Organ Preservation: From the Past to the Future”. In: *Acta Pharmacologica Sinica* 39.5, pp. 845–857. ISSN: 1671-4083, 1745-7254. DOI: [10.1038/aps.2017.182](https://doi.org/10.1038/aps.2017.182). URL: <https://www.nature.com/articles/aps2017182>.

- Kalisvaart, M., Croome, K. P., Hernandez-Alejandro, R., Pirenne, J., Cortés-Cerisuelo, M., Miñambres, E., and Abt, P. L. (June 2021). “Donor Warm Ischemia Time in DCD Liver Transplantation—Working Group Report From the ILTS DCD, Liver Preservation, and Machine Perfusion Consensus Conference”. In: *Transplantation* 105.6, pp. 1156–1164. ISSN: 0041-1337. DOI: [10.1097/TP.0000000000003819](https://doi.org/10.1097/TP.0000000000003819). URL: <https://journals.lww.com/10.1097/TP.0000000000003819>.
- Kamath, P. S. and Kim, R. W. (Mar. 2007). “The Model for End-Stage Liver Disease (MELD)”. In: *Hepatology* 45.3, pp. 797–805. ISSN: 0270-9139. DOI: [10.1002/hep.21563](https://doi.org/10.1002/hep.21563). URL: <https://journals.lww.com/01515467-200703000-00027>.
- Kastoryano, S. (2022). “Regime and Treatment Effects in Duration Models: Decomposing Expectation and Transplant Effects on the Kidney Waitlist”. In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: [10.2139/ssrn.4118215](https://doi.org/10.2139/ssrn.4118215). URL: <https://www.ssrn.com/abstract=4118215>.
- Kim, H. J., Lu, B., Nehus, E. J., and Kim, M.-O. (Feb. 10, 2019). “Estimating Heterogeneous Treatment Effects for Latent Subgroups in Observational Studies”. In: *Statistics in Medicine* 38.3, pp. 339–353. ISSN: 0277-6715, 1097-0258. DOI: [10.1002/sim.7970](https://doi.org/10.1002/sim.7970). URL: <https://onlinelibrary.wiley.com/doi/10.1002/sim.7970>.
- Kingman, J. F. C. (2002). *Poisson Processes*. Repr. Oxford Studies in Probability 3. Oxford: Clarendon [u.a.] 104 pp. ISBN: 978-0-19-853693-2.
- Liu, L., Wang, Y., and Xu, Y. (Jan. 2024). “A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data”. In: *American Journal of Political Science* 68.1, pp. 160–176. ISSN: 0092-5853, 1540-5907. DOI: [10.1111/ajps.12723](https://doi.org/10.1111/ajps.12723). URL: <https://onlinelibrary.wiley.com/doi/10.1111/ajps.12723>.
- Lozanovski, V. J., Döhler, B., Weiss, K. H., Mehrabi, A., and Süsal, C. (May 12, 2020). “The Differential Influence of Cold Ischemia Time on Outcome After Liver Transplantation for Different Indications—Who Is at Risk? A Collaborative Transplant Study Report”. In: *Frontiers in Immunology* 11, p. 892. ISSN: 1664-3224. DOI: [10.3389/fimmu.2020.00892](https://doi.org/10.3389/fimmu.2020.00892). URL: <https://www.frontiersin.org/article/10.3389/fimmu.2020.00892/full>.
- Marchese, A., de Ferrante, H., Berrevoets, J., and Verboven, S. (2024). “Improving Causal Transplant Outcomes through Dynamic Organ Offer Estimation”. In: – (Dec. 15, 2025). “DynamITE: Optimal Time-Sensitive Organ Offers Using ITE”. In: *Proceedings of the 4th Machine Learning for Health Symposium*. Ed. by S. Hegselmann, H. Zhou, E. Healey, T. Chang, C. Ellington, V. Mhasawade, S. Tonekaboni, P. Argaw, and H. Zhang. Vol. 259. Proceedings of Machine Learning Research. PMLR, pp. 696–713. URL: <https://proceedings.mlr.press/v259/marchese25a.html>.
- Moers, C., van Gelder, F., Napieralski, B. P., van Kasterop-Kutz, M., Squifflet, J.-P., Rahmel, A., and Ploeg, R. J. (2009). “Machine Perfusion or Cold Storage in Deceased-Donor Kidney Transplantation”. In: *The new england journal of medicine*.
- Nasralla, D. et al. (May 2018). “A Randomized Trial of Normothermic Preservation in Liver Transplantation”. In: *Nature* 557.7703, pp. 50–56. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/s41586-018-0047-9](https://doi.org/10.1038/s41586-018-0047-9). URL: <https://www.nature.com/articles/s41586-018-0047-9>.
- Neçki, M. et al. (Sept. 2020). “Impact of Cold Ischemia Time on Frequency of Airway Complications Among Lung Transplant Recipients”. In: *Transplantation Proceedings* 52.7, pp. 2160–2164. ISSN: 00411345. DOI: [10.1016/j.transproceed.2020.03.047](https://doi.org/10.1016/j.transproceed.2020.03.047). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0041134520301986>.
- O’Neill, S. et al. (Sept. 2020). “Novel Organ Perfusion and Preservation Strategies in Transplantation – Where Are We Going in the United Kingdom?” In: *Transplantation* 104.9, pp. 1813–1824. ISSN: 0041-1337. DOI: [10.1097/TP.0000000000003106](https://doi.org/10.1097/TP.0000000000003106). URL: <https://journals.lww.com/10.1097/TP.0000000000003106>.
- Papalexopoulos, T., Alcorn, J., Bertsimas, D., Goff, R., Stewart, D., and Trichakis, N. (July 2024). “Reshaping National Organ Allocation Policy”. In: *Operations Research* 72.4, pp. 1475–

1486. ISSN: 0030-364X, 1526-5463. DOI: [10.1287/opre.2022.0035](https://doi.org/10.1287/opre.2022.0035). URL: <https://pubsonline.informs.org/doi/10.1287/opre.2022.0035>.
- Prosperi, M., Guo, Y., Sperrin, M., Koopman, J. S., Min, J. S., He, X., Rich, S., Wang, M., Buchan, I. E., and Bian, J. (July 13, 2020). "Causal Inference and Counterfactual Prediction in Machine Learning for Actionable Healthcare". In: *Nature Machine Intelligence* 2.7, pp. 369–375. ISSN: 2522-5839. DOI: [10.1038/s42256-020-0197-y](https://doi.org/10.1038/s42256-020-0197-y). URL: <https://www.nature.com/articles/s42256-020-0197-y>.
- Rose, J., Gunsalus, P. R., Lehr, C. J., Swiler, M. F., Dalton, J. E., and Valapour, M. (Aug. 2024). "A Modular Simulation Framework for Organ Allocation". In: *The Journal of Heart and Lung Transplantation* 43.8, pp. 1326–1335. ISSN: 10532498. DOI: [10.1016/j.healun.2024.04.063](https://doi.org/10.1016/j.healun.2024.04.063). URL: <https://linkinghub.elsevier.com/retrieve/pii/S1053249824016450>.
- Ruf, A., Dirchwolf, M., and Freeman, R. B. (Jan. 2022). "From Child-Pugh to MELD Score and beyond: Taking a Walk down Memory Lane". In: *Annals of Hepatology* 27.1, p. 100535. ISSN: 1665-2681. DOI: [10.1016/j.aohep.2021.100535](https://doi.org/10.1016/j.aohep.2021.100535). URL: <https://linkinghub.elsevier.com/retrieve/pii/S1665268121002349>.
- Schaapherder, A. F. et al. (Aug. 2022). "Donor Characteristics and Their Impact on Kidney Transplantation Outcomes: Results from Two Nationwide Instrumental Variable Analyses Based on Outcomes of Donor Kidney Pairs Accepted for Transplantation". In: *eClinicalMedicine* 50, p. 101516. ISSN: 25895370. DOI: [10.1016/j.eclim.2022.101516](https://doi.org/10.1016/j.eclim.2022.101516). URL: <https://linkinghub.elsevier.com/retrieve/pii/S2589537022002462>.
- Schold, J. D. (June 2019). "Assessing the Outcome of Organ Allocation: Choosing Metrics". In: *Current Opinion in Organ Transplantation* 24.3, pp. 332–336. ISSN: 1087-2418, 1531-7013. DOI: [10.1097/mot.0000000000000637](https://doi.org/10.1097/mot.0000000000000637). URL: <https://journals.lww.com/00075200-201906000-00021>.
- Schwab, S., Elmer, A., Sidler, D., Straumann, L., Stürzinger, U., and Immer, F. (Sept. 10, 2024). "Selection Bias in Reporting of Median Waiting Times in Organ Transplantation". In: *JAMA Network Open* 7.9, e2432415. ISSN: 2574-3805. DOI: [10.1001/jamanetworkopen.2024.32415](https://doi.org/10.1001/jamanetworkopen.2024.32415). URL: <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2823322>.
- Shalit, U., Johansson, F. D., and Sontag, D. (Aug. 6, 2017). "Estimating Individual Treatment Effect: Generalization Bounds and Algorithms". In: *Proceedings of the 34th International Conference on Machine Learning*. Ed. by D. Precup and Y. W. Teh. Vol. 70. Proceedings of Machine Learning Research. PMLR, pp. 3076–3085. URL: <https://proceedings.mlr.press/v70/shalit17a.html>.
- Tim Weaver, Josh Pyke (2023). "Background and Methodology in Simulation Analysis: OASim". In.
- "Treatment Effects" (2005). In: Reichardt, C. S. *Encyclopedia of Social Measurement*. Elsevier, pp. 875–882. ISBN: 978-0-12-369398-3. DOI: [10.1016/b0-12-369398-5/00186-9](https://doi.org/10.1016/b0-12-369398-5/00186-9). URL: <https://linkinghub.elsevier.com/retrieve/pii/B0123693985001869>.
- Trivedi, H. D. (May 2022). "The Evolution of the MELD Score and Its Implications in Liver Transplant Allocation: A Beginner's Guide for Trainees". In: *ACG Case Reports Journal* 9.5, e00763. ISSN: 2326-3253. DOI: [10.14309/crj.0000000000000763](https://doi.org/10.14309/crj.0000000000000763). URL: <https://journals.lww.com/10.14309/crj.0000000000000763>.
- Weeder, P. D., Van Rijn, R., and Porte, R. J. (July 2015). "Machine Perfusion in Liver Transplantation as a Tool to Prevent Non-Anastomotic Biliary Strictures: Rationale, Current Evidence and Future Directions". In: *Journal of Hepatology* 63.1, pp. 265–275. ISSN: 01688278. DOI: [10.1016/j.jhep.2015.03.008](https://doi.org/10.1016/j.jhep.2015.03.008). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168827815001737>.
- Wood, N. L., Weaver, T., Kwong, A. J., and Gentry, S. E. (Nov. 2024). "Principles for Simulating the Organ Allocation System". In: *Liver Transplantation* 30.11, pp. 1101–1103. ISSN: 1527-6465. DOI: [10.1097/LVT.0000000000000440](https://doi.org/10.1097/LVT.0000000000000440). URL: <https://journals.lww.com/10.1097/LVT.0000000000000440>.

- Yoon, J., Jordon, J., and Schaar, M. van der (2018). “GANITE: Estimation of Individualized Treatment Effects Using Generative Adversarial Nets”. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=ByKWUeWA->.
- Zhang, S., Suen, S.-C., Sundaram, V., and Gong, C. L. (May 4, 2025). “Quantifying the Benefits of Increasing Decision-Making Frequency for Health Applications with Regular Decision Epochs”. In: *IJSE Transactions* 57.5, pp. 469–483. ISSN: 2472-5854, 2472-5862. DOI: [10.1080/24725854.2024.2321492](https://doi.org/10.1080/24725854.2024.2321492). URL: <https://www.tandfonline.com/doi/full/10.1080/24725854.2024.2321492>.
- Zhang, W., Zhang, Z., and Liu, A. (Dec. 1, 2023). “Optimizing Treatment Allocation in Randomized Clinical Trials by Leveraging Baseline Covariates”. In: *Biometrics* 79.4, pp. 2815–2829. ISSN: 0006-341X, 1541-0420. DOI: [10.1111/biom.13914](https://doi.org/10.1111/biom.13914). URL: <https://academic.oup.com/biometrics/article/79/4/2815/7587583>.
- Zhang, Y., Hu, A., Lin, Y., Cao, Y., Muller, S., Wong, G., and Yang, J. Y. H. (Sept. 29, 2023). “simKAP: Simulation Framework for the Kidney Allocation Process with Decision Making Model”. In: *Scientific Reports* 13.1, p. 16367. ISSN: 2045-2322. DOI: [10.1038/s41598-023-41162-w](https://doi.org/10.1038/s41598-023-41162-w). URL: <https://www.nature.com/articles/s41598-023-41162-w>.