Allocating Scarce Resources: Modeling and Optimization

by

Samuel Gilmour
B.S., University of Auckland (2017)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Operations Research
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2023

Author .....
Sloan School of Management
May 5, 2023

Certified by .....
Nikolaos Trichakis
Associate Professor of Operations Management
Thesis Supervisor

Certified by .....
Patrick Jaillet
Dugald C. Jackson Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by .....
Georgia Perakis
William F. Pounds Professor of Management Science
Co-Director, Operations Research Center

Allocating Scarce Resources:

Modeling and Optimization

by

Samuel Gilmour

Submitted to the Sloan School of Management
on May 5, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

There are countless settings in which an authority must choose how to allocate scarce resources among a set of recipients. Deceased-donor organs must be allocated to patients, public school spaces to students, and public housing to residents. When resources are extremely scarce, it is particularly important for the authority to build an allocation system that achieves an acceptable trade-off between efficiency and equity. This thesis contributes several models and tools that both support the design of new allocation systems and extract insight from existing ones. In both cases, efficiency and equity take center stage.

Chapter 2 considers the problem faced by an authority who allocates resources according to a scoring system. A scoring system is based on two foundations: a scoring rule, which is a function that computes scores for each recipient-resource pair based on some observable properties that relate the pair, and an allocation procedure, which determines the allocation using only the scores. We introduce a model that allocates a set of resource types among a set of patient types according to a scoring system, before presenting several optimization formulations and heuristics that directly optimize the scoring rule while scaling to practical problem sizes. We also show how a scoring rule of high quality in the type-based model can fail in a setting where individual recipients and resources exhibit within-type variation in properties, and suggest approaches that perform well when allocating individuals.

Moving away from the specific setting of a scoring system, Chapter 3 shows that the ability for recipients to choose whether to accept or decline the offer of a resource can act as a hidden source of inequity in an allocation system. We formulate several game-theoretic models based on two groups of recipients, selective and non-selective, who display different propensities to accept or decline the offer of a resource. We define the notion of an equilibrium in these models and provide numerical experiments showing that inequity can arise directly as a result of the disparity in selectiveness between recipients.

Chapter 4 studies a mass screening program for SARS-CoV-2 that was implemented in Greece during 2021, in which the Greek National Public Health Organi-

zation allocated a finite supply of mandatory self-tests among different segments of the population. We develop a novel compartmental model to describe the dynamics of the COVID-19 pandemic in Greece, placing particular focus on the testing procedures. We fit the model to detailed data to quantify the overall effectiveness of the program in reducing hospitalizations and deaths, and also to understand the effects of several operational decisions. We conclude that self-testing is an extremely important intervention to consider for pandemic preparedness.

Thesis Supervisor: Nikolaos Trichakis
Title: Associate Professor of Operations Management

Thesis Supervisor: Patrick Jaillet
Title: Dugald C. Jackson Professor of Electrical Engineering and Computer Science

Acknowledgments

The last five years have been a rollercoaster, which I consider myself extremely lucky to have ridden with some incredible people.

First of all, a sincere thanks to my advisors, Patrick Jaillet and Nikos Trichakis. Patrick and Nikos have been immensely supportive throughout my PhD. Whenever I felt stuck or found myself down a rabbit hole, they were there to fish me out and set me on a better path. Their clear-headed approaches to research have been incredibly illuminating to observe, and if I can take even a fraction of their skill and dedication forward into my professional career, then I will be doing very well indeed.

I greatly appreciate the support of the other half of my thesis committee, Alexandre Jacquillat and Kimon Drakopoulos. It has also been a pleasure to work with Kimon for a couple of years on our collaboration with the Greek National Public Health Organization, where Kimon showed his astonishing ability to break down the most complicated systems into digestable parts. I learnt from him that when I feel overwhelmed, taking a step back is usually enough. I also would like to thank Gkikas Magiorkinis and Spyros Sapounas for making this project possible.

I will never forget the friends I have made at the ORC, and in particular the cohort of 2018 legends: Kayla, Vassilis, Georgia, Josh and Galit. The first two years of the program were the most intense of my life, and spending them with these outstanding people made them some of the most fulfilling, too. The fun we had will probably be what I look back on most fondly in the decades to come, and though all of our paths are about to diverge, I hope we never fully lose touch.

I also want to thank my friends outside the ORC, in Boston and New Zealand, for the time spent unwinding at basketball games or stumbling around during the not-quite-annual disc golf expeditions. A heartfelt thanks go to Elis, John, and Dave for staving off the darkness and making me laugh when laughter was a scarce resource.

But the biggest thanks need to go to my family. It has been awesome to reconnect with my cousins Dan and Natalie, and their wonderful partners Ellie and Lou, in the last couple of years, and our video calls have consistently been massive mood boosters.

It truly means a lot to know we have each other's backs.

To Ana, who has been the most amazing partner in crime and shoulder to cry on. I am not entirely sure how I got so lucky, but your love and support have meant the world during some difficult times, and you will always be the best thing to come out of my time in Boston. Thank you, I love you, and I really cannot wait for the future.

Lastly, to my parents, Paul and Julie. As I grow older, I learn to appreciate all the sacrifices you made to give me the opportunities I have had in life. You have encouraged me, celebrated my successes, cheered me up after my failures, and most importantly, believed in me to the fullest since before I can even remember. You are the absolute best. From the bottom of my heart, thank you, and I love you more than words can say. This thesis is dedicated to you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Anyone who has studied economics, even just for a single term in high school, should be familiar with the concept of resource scarcity. The fact that many resources have only finite supply means that societies need mechanisms to coordinate how these resources are allocated among the recipients that demand them. The two most common mechanisms that operate to this end are markets, in which participants exchange resources among themselves, and planning systems, whereby resources are allocated according to some computational procedure supervised by a central authority. In essence, this thesis studies how to design effective planning systems.

Societies typically choose a mechanism to allocate a particular type of scarce resource based on ethical principles that prevail at the time. In the present day, planning systems are considered the most suitable for allocating extremely scarce resources which contribute directly to fundamental human rights such as education and healthcare, so that the utility derived from these resources is not concentrated in the hands of those with greater economic power. Deceased-donor organs [1], COVID-19 treatments [2, 3], admissions to public schools, and public housing are some examples of resources which are allocated using planning systems in most countries around the world.

Given that a planning system has been chosen to allocate a type of scarce resource, the fundamental question faced by the authority is: how should the allocation mechanism be designed to achieve their goals? Broadly speaking, the authority aims to

achieve a system which is *efficient*, maximizing some notion of utility, and *equitable*, minimizing unfairness and disparities in outcomes among recipients. Unfortunately, these goals are frequently in tension with each other, and arriving at an allocation mechanism which represents a suitable trade-off requires extreme care [4].

Operations Research (OR) can help authorities design an allocation mechanism in several ways. First, mathematical models describing the mechanism can be useful approximations to reality, so as to improve understanding of the trade-off between efficiency and equity. Second, given such models, tools from optimization can be used to help the authority tune their allocation mechanism and obtain one which lies on the frontier of the trade-off between efficiency and equity. This thesis proposes models for understanding systems and provides tools to optimize them in several settings.

Chapter 2 considers the problem of optimizing a scoring system, which is a particular type of mechanism in which a score is assigned to each recipient in the system every time a resource must be allocated. Once the scores have been assigned, the mechanism allocates the resource only on the basis of the scores. In other words, given the scores, a scoring system is fully blind to any differences in the properties of recipients. The chapter begins by proposing a general model of a scoring system that allocates heterogeneous resource types among heterogeneous recipient types, and describes some basic steady-state quantities arising in this system.

The chapter proceeds by specifically modeling a *priority system*, whereby a resource is offered to the top-scoring recipient, and formulates the optimization problem faced by the authority when deciding how to choose the scores. The formulation provides significant modeling flexibility by remaining solvable to optimality with around 30 recipient and resource types. In the case where more are required, the chapter describes and analyzes some heuristics for approximating the solution.

The chapter continues by modeling a *lottery system*, whereby a resource is allocated randomly among recipients based on their scores. Again, it formulates the optimization problem faced by an authority choosing the scores, and again, this formulation provides modeling flexibility by scaling to around 15 recipient and resource types.

In many allocation systems, recipients are allowed to decline the offer of a resource and choose to wait for future offers. Chapter 3 shows that this form of recipient choice can be an important source of inequity. It formulates a general game-theoretic queueing model based on two groups of recipients, selective and non-selective, who display different propensities to accept or decline the offer of a resource, and defines the notion of an equilibrium.

The chapter then instantiates the model with several allocation mechanisms: for example, one where each individual recipient receives the next offer with equal probability, and one where their chance is weighted by their waiting time in the system so far. Numerical testing shows that inequity can arise in these models directly as a result of the disparity in selectiveness between recipients.

Finally, Chapter 4 considers a more practical setting based around a real-world dataset. In 2021, at the height of the COVID-19 pandemic, the Greek government implemented a mass screening program to curb the spread of the virus by requiring citizens to take self-tests at home. Had the entire population been tested, the cost would have been prohibitive. Therefore, the government had to choose how to distribute scarce self-tests among the population to best slow the spread of infection. This chapter develops a novel compartmental model that describes the spread of a disease through a population, with a particular focus on the self-testing procedure.

The model was fitted to granular data provided by the Greek National Public Health Organization (NPHO) and used to quantify the overall impact of the program on hospitalizations and deaths, ultimately concluding that the program avoided a significant number of both. In addition, the fitted model was used to understand the effects of several operational decisions made by the government during the implementation of the program.

Chapters 2 and 3 describe quite general allocation settings that are relevant for many types of scarce resource, whereas Chapter 4 describes a very specific setting. In each case, it is our hope that the models and optimization formulations will be of use to practitioners who are faced with the difficult problem of designing an allocation mechanism for their particular scarce resource.

Chapter 2

Modeling and Optimizing Scoring Allocation Systems

2.1 Introduction

One type of mechanism for allocating scarce resources that is frequently implemented in practice is a *scoring system*. In this type of mechanism, for each resource, an authority assigns a score to each recipient in the system and then uses a procedure to allocate the resource to a recipient in a way that depends only on the scores. In other words, if two recipients have the same score, they are treated as identical by the allocation procedure.

Though in principle they could be assigned arbitrarily, scores are typically computed using a *scoring rule*. This is a function that acts on observable properties relating the recipient-resource pair to produce a number (a score). When a system is implemented with a scoring rule that can be easily interpreted, it has the attractive property that allocations are easily understood by its participants.

In addition to the scoring rule, a scoring system also requires a specific procedure to act on the scores and allocate each resource. Two in particular are common in practice. A *priority mechanism*, which allocates each resource to the recipient with the highest score, is one. Such a system satisfies several fundamental principles of equity identified by Young [5]. *Lottery mechanisms* are a different approach to allocation

that use randomization to appeal to an intuitive human notion of fairness. In this chapter, we study both priority and lottery systems that make use of a scoring rule.

The main question we focus on is the following: given either a priority or lottery allocation mechanism, how can the authority select a scoring rule to obtain a desirable allocation? Surprisingly few studies have aimed to answer this question. Even though scoring rule systems are widely implemented in practice and backed by sound economic principles, guidance for implementing these systems has therefore been sparse. This chapter addresses the deficit by providing models and formulations that can be used by an authority as they attempt to design an effective scoring system for allocating a scarce resource.

2.1.1 Main Results and Contributions

Section 2.2 describes a general setting for allocating scarce resources with a scoring rule and develops a fluid approximation model based on heterogeneous recipient and resource types. Section 2.3 introduces priority allocation, instantiates the optimization problem faced by the authority, and proves it to be NP-hard. We present a mixed-integer program (MIP) formulation with a simple structure and show that the formulation scales to sizes that are useful for practical modeling purposes. Section 2.4 provides two heuristics for solving this problem approximately, as well as bounds on their performance. Section 2.5 explains how scoring rules of high quality in the type-based model can fail when individual recipients and resources exhibit within-type variation in their properties, and provides an approach to deal with this situation.

In Section 2.6, we introduce lottery allocation, define the optimization problem faced by the authority, formulate it as a nonconvex quadratic problem, and provide numerical experiments for a heuristic and an approach for allocating to individuals. Section 2.7 concludes with a brief discussion on equity in the context of the two allocation mechanisms.

2.1.2 Related Literature

As previously mentioned, little effort has been dedicated to the problem of selecting a scoring rule that maximizes welfare under a given allocation mechanism. Of the few studies that address the problem, all focus on priority allocation where a resource is allocated to the recipient with the highest score. They are generally motivated by applications in organ allocation systems.

Bertsimas, Farias, and Trichakis [6] study the deceased-donor kidney allocation system with a data-driven optimization approach. They develop a heuristic for finding a scoring rule that produces an efficient allocation while approximately satisfying ex-post equity constraints under the priority mechanism. Their procedure first solves a maximum-weighted bipartite matching problem with equity constraints and uses the optimal dual variables to compute *equity-adjusted* reward coefficients for each patient-organ pair. They then solve a regression problem to describe these coefficients using a scoring rule. Though they also optimize for the scoring rule, their heuristic does not attempt to model the mechanics of the allocation mechanism as we do in this chapter.

Ding, McCormick, and Nagarajan [7] use tools from queuing theory to develop a type-based model of resources and recipients where the scoring rule is the sum of two terms: a function of the types of a recipient-resource pair, and an increasing function of the waiting time within a recipient queue. All resources of a particular type are allocated to the recipient queue with the highest score. They derive an approximation to the steady-state behaviour of the system and find a scoring rule that maximizes a specific objective function which trades off efficiency with equity.

With a discrete stochastic model, David and Yechiali [8] build on Righter [9] to consider the problem of allocating a finite set of organs among a finite set of patients. In their setting, the reward obtained by matching an organ to a patient depends only on a single attribute. While their characterization of the optimal policy allows it to be interpreted as a scoring rule, our model is concerned with types that are identified by multiple attributes rather than just one.

It is also worth highlighting some related work which either entirely sweeps away the role of the scoring rule or marginalizes its importance. Among this group is a growing body of recent work that studies how techniques from optimization can be applied to allocation systems. Shi [10] recognizes that priorities are often optimization variables for the authority to choose. He proposes and studies a model that allocates a continuum of customers (split into market segments) to resources using a priority mechanism, while optimizing over the priorities. The model is general enough to encompass allocation mechanisms such as the Gale-Shapley deferred acceptance, top trading cycles, and serial dictatorship algorithms. Ashlagi and Shi [11] and Bodoh-Creed [12] consider the allocation of heterogeneous resources to customers who have both a publicly known type and privately known vector of utilities. They optimize the allocation as a function for each customer type that maps a utility vector onto a set of allocation proportions over the resources.

In the realm of queueing theory, Afeche, Caldentey, and Gupta [13] consider a queuing system with a set of heterogeneous customer classes processed by a set of heterogeneous servers, and seek to find the matching topologies (which are not selected based on scores) in this bipartite graph which define the Pareto frontier for system reward and customer delay. Sisselman and Whitt [14] and Mehrotra et al. [15] both model allocation of calls in a call center and aim to maximize match-specific rewards.

While there is a lack of technical research on the role of scoring rules in allocation systems, real-world implementations have been studied extensively from the perspective of their operations: the methods used to choose scoring rules, debates on the ethics of allocation mechanisms, and the outcomes observed when systems are simulated or implemented in practice. Sparrow [16] reviews a historical example arising from the demobilization of United States (US) army troops after the Second World War, in which individual soldiers posted overseas were removed from duty according to points assigned by a scoring rule. The scoring rule took into account factors such as the number of dependents of the soldier and their time spent in combat. Been et al. [17] and Thakral [18] study the public housing allocation system in the US (for which some regions use a points-based system) and Greely [19] notes that jobs in the

US civil service have historically been assigned on the basis of a scoring rule. Edwards [20] describes waitlist systems that operate within the National Health Service (NHS) in the United Kingdom (UK) and the scoring systems that allocate medical services to patients on these lists. Zenios, Wein, and Chertow [21] use a simulation model and data taken from the US to compare various mechanisms for deceased-donor kidney allocation, among which they include the scoring system used at that time.

2.2 Allocating Resources with a Scoring Rule

This section describes the resource allocation setting considered in this chapter. We briefly describe the role of scoring rules in many real-world allocation systems and then develop a fluid model that leads to tractable (though NP-hard) optimization problems. This section does not introduce any specific procedures for allocating resources. It is left for Sections 2.3 and 2.6 to describe the two allocation mechanisms we study: the priority and lottery mechanisms, respectively.

Our motivation is a setting where an authority must allocate resources among recipients. In much of the literature on allocation systems, the authority is allowed complete freedom in choosing which recipient to allocate each resource. Our setting is different: when a resource arrives, the authority computes a score for each recipient using a function called a *scoring rule* which acts on a vector of properties that relates the recipient-resource pair. An *allocation mechanism* then allocates the resource to a recipient only using these scores. This approach is used frequently in practice – for instance, the kidney allocation system in the US uses a scoring rule that assigns points to patients for properties such as the number of years they have spent on the waitlist, and then allocates the next kidney to the patient with the highest score [22].

Throughout the chapter, we specifically consider the setting where resources are scarce. This statement will soon be made precise when the fluid model is introduced, but in essence it means that new recipients arrive at a much greater rate than the resources. Many real-world allocation systems have this characteristic: for instance, in 2020, patients waiting to receive a kidney transplant in the US outnumbered available kidneys by 5 to 1 [23]. In 2021, low-income families outnumbered available affordable housing units in some urban centers within the US by up to 6 to 1 [24].

When resources are scarce, it is important that they are allocated in line with certain axiomatic principles of equity. Young [5] notes that scoring rule systems provide *impartiality*: no distinctions are made between recipients except for differences in their properties. Two recipients who have the same properties will necessarily have the same scores and therefore will be treated identically by the allocation mechanism.

He goes on to identify two further principles that are satisfied specifically by the priority mechanism; we will return to these when appropriate.

2.2.1 Fluid Model

Our model is based on I heterogeneous recipient types and J heterogeneous resource types. Figure 2-1 illustrates the setup, which operates as follows. Recipients of type i arrive into the system at a rate of $\lambda_i \in \mathbb{R}_+$ and enter a queue. Resources of type $j \in [J]$ arrive at a rate of $\mu_j \in \mathbb{R}_+$ and are fractionally allocated across the recipient queues. Let $x_{ij} \in [0, 1]$ be the fraction of resource type j arrivals that are allocated to the recipient queue of type i, meaning that resources are allocated to queue i at a rate $\sum_{j=1}^{J} \mu_j x_{ij}$. Assume that $\bar{\lambda}_i \coloneqq \lambda_i - \sum_{j=1}^{J} \mu_j x_{ij} \geq 0$ (we will return to the significance of this assumption soon).

The allocation fractions in $\mathbf{x}$ must belong to a set, $\mathcal{X}$, so that no more than the total resources arriving are allocated among the recipient queues:

$$\mathcal{X} \coloneqq \left\{ \mathbf{x} \in [0,1]^{I \times J} : \sum_{i=1}^{I} x_{ij} \leq 1, \ \forall j \in [J] \right\}$$

Let $L_i(t)$ be the length of queue i at some time t and $\delta > 0$ be the next small unit of time. Within this unit of time, $\delta\lambda_i$ individuals join the queue and $\delta \sum_{j=1}^{J} \mu_j x_{ij}$ individuals leave the queue after receiving a resource. We introduce a new parameter, q_i, to model recipient attrition seen in systems such as organ allocation (where patients can die or become too sick to receive a transplant). Now, each remaining individual in the queue decides to leave on their own with probability δq_i, which depletes the queue by $\delta q_i \left(L(t) + \delta\bar{\lambda}_i \right)$. These dynamics allow us to derive an expression for the rate of change of the length of the queue with time:

$$L_i(t+\delta) - L_i(t) = \delta\bar{\lambda}_i - \delta q_i \left(L_i(t) + \delta\bar{\lambda}_i \right)$$

$$\lim_{\delta \to 0} \left[\frac{L_i(t+\delta) - L_i(t)}{\delta} \right] = \bar{\lambda}_i - q_i L_i(t)$$

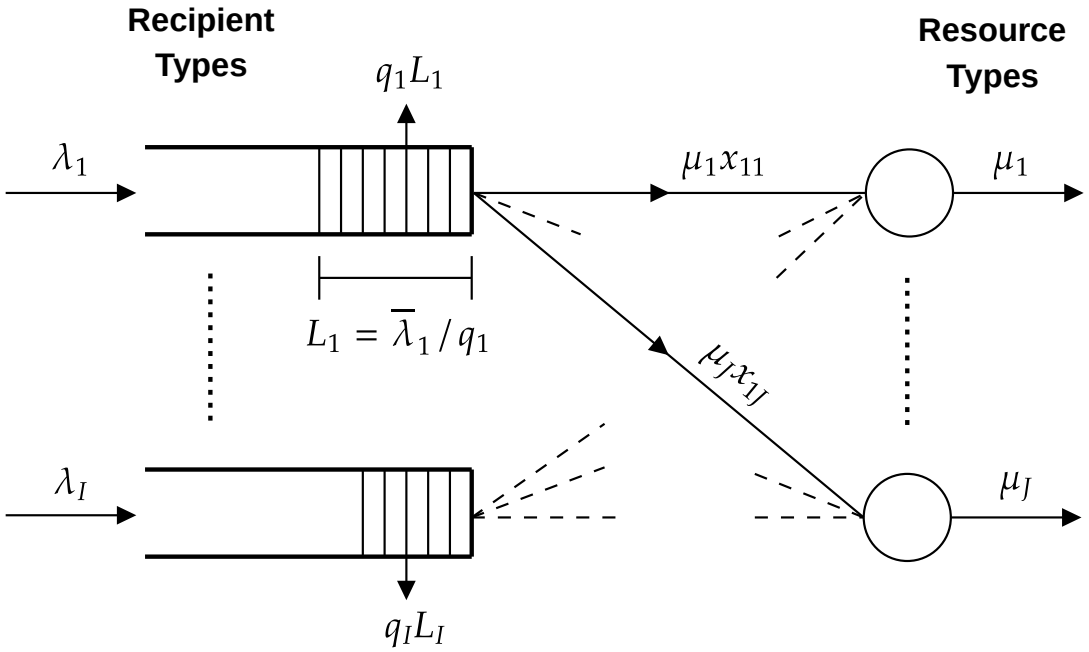

Figure 2-1: Setting for the fluid approximation model, where recipients arrive and form queues of types while being served by arriving resources (split across queues).

Solving the resulting differential equation for $L_i(t)$ and taking the limit as $t \to \infty$ leaves an expression for the steady-state length of the queue:

$$L_i = \frac{\bar{\lambda}_i}{q_i}$$

Finally, let $R_{ij} \in \mathbb{R}_+$ be the reward obtained by allocating a resource of type j to a recipient of type i, and define $r_{ij} \coloneqq \mu_j R_{ij}$. The rate of reward earned by the authority is:

$$\sum_{i=1}^{I} \sum_{j=1}^{J} r_{ij} x_{ij}$$

2.2.2 Optimizing Allocation Fractions

So far, we have described a model where fixed allocation fractions give rise to queues with lengths and an overall rate of earning reward. But our objective is to model a system that operates with a scoring rule – therefore, allocation fractions should be

computed based on the scores of each recipient-resource pair, and scores should be computed based on properties that relate these pairs.

With $K \in \mathbb{N}$ being this number of properties, let $\mathbf{f}_{ij} \in \mathbb{R}^K$ be the values relating recipient type i with resource type j. These vectors act as the inputs for the scoring rule, $s : \mathbb{R}^K \to \mathbb{R}$. As a shorthand we use $s_{ij} = s(\mathbf{f}_{ij})$ to refer to the score of recipient type i for resource type j. An *allocation function* is simply a function that maps scores onto allocation fractions:

Definition 2.1 (Allocation Function)**.** *An allocation function, defined on a set of score matrices, $\mathcal{S} \subseteq \mathbb{R}^{I \times J}$, has the signature $\mathbf{x} : \mathcal{S} \mapsto \mathcal{X}$. It maps a matrix of scores onto a matrix of allocation fractions, and must be scale-free by satisfying the property:*

$$\mathbf{x}(\theta \mathbf{s}) = \mathbf{x}(\mathbf{s}) \qquad \forall \theta > 0,\ \forall \mathbf{s} \in \mathcal{S}$$

The two allocation functions we define and study correspond to the priority and lottery allocation mechanisms, and we will define them so that the type-based model approximates the (deterministic or random) procedure that allocates individual resources to individual recipients.

An allocation function may be defined on all score matrices in $\mathbb{R}^{I \times J}$, but the authority is the one who actually generates the scores by choosing a scoring rule from a restricted class to act on the properties. This decision is described in more detail soon, but for now let $\overline{\mathcal{S}} \subseteq \mathcal{S}$ be the set of scores that the authority can produce.

Now we return to the assumption that $\bar{\lambda}_i \geq 0$ for each queue i. When resources are scarce, it can be expected that the recipient queues in our model do not deplete regardless of the scores chosen by the authority. This observation motivates the precise definition of scarcity used in the rest of this chapter:

Assumption 2.1 (Scarce Resources)**.** *For some allocation function, $\mathbf{x}$, resources are assumed to be scarce in the sense that, for every recipient type i, the following holds:*

$$\lambda_i - \sum_{j=1}^{J} \mu_j x_{ij}(\mathbf{s}) \geq 0 \qquad \forall \mathbf{s} \in \overline{\mathcal{S}}$$

With all these definitions in place, we now define the central optimization problem faced by an authority who aims to maximize the rate of reward earned under a particular allocation function, $\mathbf{x}$, and set of possible scores, $\overline{\mathcal{S}}$:

Problem 2.1 (Fluid Model Optimization)**.**

$$\max_{\mathbf{s}\in\overline{\mathcal{S}}} \quad \sum_{i=1}^{I}\sum_{j=1}^{J} r_{ij}x_{ij}(\mathbf{s}) \tag{2.1}$$

2.2.3 Scoring Rule Considerations

Recall that $\overline{\mathcal{S}}$ is produced by giving the authority a set of candidate scoring rules to choose from. We consider linear rules as the broad set of candidates, but since the priority and lottery allocation functions are defined on different sets of scores, some extra restrictions must be placed on these rules depending on the allocation function. These will be introduced in the relevant sections.

Above all else, a linear scoring rule is *interpretable*. Choosing a linear rule simply means specifying a weight vector $\mathbf{w} \in \mathbb{R}^K$ so that $s(\mathbf{f}_{ij}) = \mathbf{w}^\intercal \mathbf{f}_{ij}$ with the interpretation that w_k is the marginal score assigned for a unit increase in the value of the kth property in $\mathbf{f}_{ij}$. This information is easily digested by the recipients in an allocation system, and it contributes to explaining why linear scoring rules have been so widely adopted in real-world systems. Regardless, if we wish to assign scores based on a nonlinear function of the properties then an extra component that computes this nonlinear function can simply be appended to $\mathbf{f}_{ij}$.

Throughout this chapter, we will place almost no extra restrictions on the linear scoring rules. However, in practice, it may be desirable to impose some constraints: for example, a cardinality constraint on the number of nonzero weights. Many useful constraints can easily be incorporated within our optimization formulations without significant impact on the techniques used to solve them.

2.3 Priority Allocation

This section introduces the *priority* allocation function, which is designed to model the allocation procedure in real-world allocation systems where each resource is allocated to the top-scoring recipient. We will define the allocation function, analyze the optimization problem faced by the authority, and present some numerical results.

First, we note two observations made by Young [5]. As well as the principle of impartiality that has been described, allocating resources using a priority mechanism satisfies the principle of *consistency*. This means that two recipient types i and i' will be allocated a particular resource type in the same way regardless of the other recipients in the system – certainly a desirable property for an allocation system.

Young [5] also notes that when linear scoring rules are used in a priority mechanism, the resulting system is *separable*: scores do not exhibit complementary effects between properties. To be precise – fix a resource type j and suppose we have two arbitrary pairs of recipient types: $a, b \in [I]$ and $c, d \in [I]$. Each pair is identical in properties 3 through K (though possibly different between the pairs). Now let a match c in the first two properties ($f_{aj1} = f_{cj1}$ and $f_{aj2} = f_{cj2}$) and the same for b and d. A scoring mechanism is separable if $s_{aj} \geq s_{bj} \iff s_{cj} \geq s_{dj}$. Though some allocation settings may indeed be better suited to inseparable mechanisms, it is important to note that complementary effects can still be modelled by amending $\mathbf{f}_{ab}$ with extra properties that are computed based on a function of the original properties.

The priority allocation function, retaining notation as $\mathbf{x}$, can be applied to any matrix of scores. Accordingly, we set $\mathcal{S} = \mathbb{R}^{I \times J}$. Letting $\mathcal{I}_j = \arg\max_{i \in [I]}(s_{ij})$ be the top-ranked recipient types for resource type j, the priority allocation function is:

$$x_{ij}(\mathbf{s}) = \begin{cases} 1 & i \in \arg\max_{k \in \mathcal{I}_j}(r_{kj}) \\ 0 & \text{otherwise} \end{cases}$$

$x_{ij}(\mathbf{s})$ is equal to 1 whenever recipient type i is top-scoring for resource type j and also wins a tiebreaker by having the highest reward of all top-scoring queues. Note that this function is not well-defined when there are two top-scoring queues that are

also tied in the reward tiebreaker. We leave it up to implementation on how to deal with these secondary ties given that any method does not impact the reward earned. Without secondary ties, $\mathbf{x}$ certainly satisfies the required signature $\mathbf{x} : \mathcal{S} \mapsto \mathcal{X}$, whilst being scale-free. It is therefore a valid allocation function according to Definition 2.1.

In the fluid model, this allocation function assigns all resources to their top-ranked recipient queue. It is therefore useful for modeling systems such as the previously-described organ allocation systems under the assumption of resource scarcity.

Finally, a restriction must be placed on the candidate scoring rules from which the authority chooses. Since setting $\mathbf{w} = \mathbf{0}$ results in a score matrix of zeroes and $\mathbf{x}(\mathbf{0})$ simply picks the recipient type with the highest reward for each resource type, we ignore the trivial $\mathbf{w} = \mathbf{0}$ rule. The set of scores available to the authority is therefore:

$$\overline{\mathcal{S}} = [\mathbf{w}^\intercal \mathbf{f}_{ij} : \mathbf{w} \neq \mathbf{0}]$$

It is possible that some problem instances admit other rules that produce the score matrix of zeroes (or some other constant matrix) – we comment on how to deal with these in Section 2.3.2, but eliminating $\mathbf{w} = \mathbf{0}$ means this situation does not occur in every instance and therefore allows us to establish some complexity results.

2.3.1 Complexity of Optimization

This definition of $\mathbf{x}$ and $\overline{\mathcal{S}}$ combined with Problem 2.1 gives rise to a well-defined optimization problem. This subsection characterizes the complexity of the problem both intuitively (with an example) and formally (with a result showing it is NP-hard).

First, though, some new notation is required. The priority allocation function ensures that each $x_{ij}(\mathbf{s}) \in \{0, 1\}$ and it is therefore convenient to introduce a tuple $\mathbf{y} = (i_1, \ldots, i_J)$ when referring to an allocation where $x_{i_j j} = 1$ for each resource type j. This notation allows a *feasible allocation* to be easily defined:

Definition 2.2. $\mathbf{y} = (i_1, \ldots, i_J)$ *is feasible under the priority allocation function if:*

$$\exists \mathbf{s} \in \overline{\mathcal{S}} : x_{i_j j}(\mathbf{s}) = 1, \ \forall j \in [J]$$

Much of the difficulty in solving the optimization problem is due to the fact that only a subset of all allocations are feasible for any given instance of the problem. In fact, Proposition 2.1 provides a basic characterization of the feasible allocations in terms of the polyhedral structure of the property vectors. A proof is included in Appendix A.1.

Proposition 2.1 (Characterization of Feasibility)**.** *Let* $\mathcal{P}_j = conv(\{\mathbf{f}_{1j}, \ldots, \mathbf{f}_{Ij}\})$ *be a polytope defined for each resource type,* $j \in [J]$*. Let also* $\mathbf{y} = (i_1, \ldots, i_J)$ *be some allocation. Then* $\mathbf{y}$ *is feasible if and only if* $\sum_{j=1}^{J} \mathbf{f}_{i_j j}$ *is an extreme point of* $\sum_{j=1}^{J} \mathcal{P}_j$*.*

Unfortunately, this restricted set of feasible allocations can contain local optima, in the sense that a feasible allocation may have an objective value strictly greater than all its neighbours in $\sum_{j=1}^{J} \mathcal{P}_j$ without being the global optimum. Example 2.1 gives an illustration of such local optima – which, at an intuitive level, are the main sources of complexity in the problem. When combined with the observation that the number of extreme points in the sum of J polytopes may increase exponentially in J [25], we are left with a complex search space of exponential size that contains local optima.

Proposition 2.2 establishes that the problem is, in fact, NP-hard. A proof using a reduction from the Maximum Feasible Linear System (MAX-FLS) problem studied by Amaldi and Kann [26] is included in Appendix A.2.

Proposition 2.2 (NP-Hardness of Optimization)**.** *Problem 2.1, paired with the priority allocation function, is NP-hard when its input size is measured in the number of resource types,* J*.*

Amaldi and Kann [26] provide a simple algorithm that approximates MAX-FLS within a constant factor of 2 and also show that the problem cannot be approximated arbitrarily well. Unfortunately, these results do not easily extend to our problem. In Section 2.4 we describe two heuristics which provide good practical performance and establish some elementary performance bounds – but results with similar strength as those available for the MAX-FLS problem remain elusive.

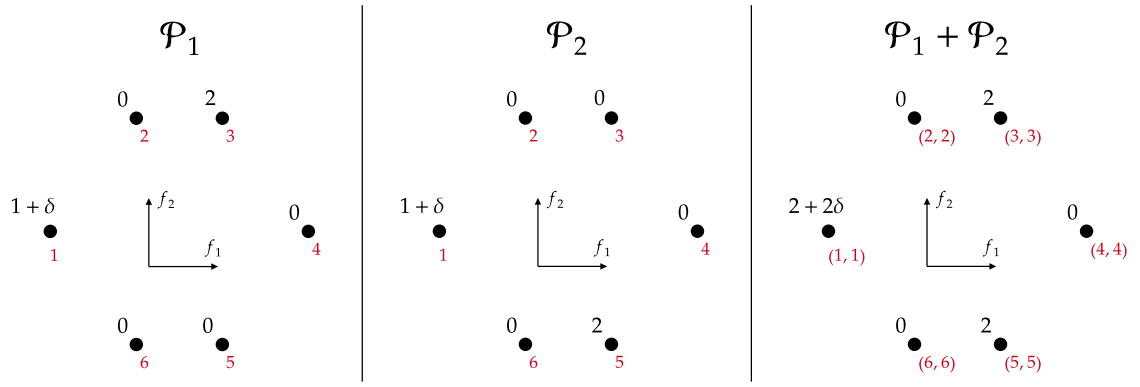

Figure 2-2: Example demonstrating a local optimum in the optimization problem with the priority allocation function.

Example 2.1. Consider a case where $I = 6$, $J = 2$, $K = 2$. The problem data take the structure shown in Figure 2-2. In this illustration, the left-hand and center groups of points correspond to properties for the first and second resource types respectively, and the right-hand group of points is the set of extreme points of the sum $\mathcal{P}_1 + \mathcal{P}_2$. Red labels are the recipient type indices and black labels are the reward coefficients for the pair. $\delta > 0$ is a small positive constant.

It is immediately clear that, for example, the allocation $\mathbf{y} = (3, 5)$ is infeasible: if recipient type 3 is selected for the first resource type, recipient type 5 cannot be selected for the second resource type. This allocation does not correspond to an extreme point of $\mathcal{P}_1 + \mathcal{P}_2$.

A greedy algorithm that begins by choosing either recipient type 3 for resource type 1 or recipient type 5 for resource type 2 will find itself stuck in a local optimum within $\mathcal{P}_1 + \mathcal{P}_2$, because these choices prevent any reward from being obtained for the other resource type. The optimal decision is to pick $\mathbf{w} = (-1, 0)$ which leads to an allocation $\mathbf{y} = (1, 1)$ and reward of $2 + 2\delta$.

2.3.2 Optimization Using a MIP

Even though the problem is NP-hard, this subsection shows that it admits formulation as a linear MIP. It is useful to first introduce a group of sets that divide the set of linear scoring rules the authority may choose from:

$$\mathcal{W}_{ij} = \left\{\mathbf{w} \in \mathbb{R}^K : \mathbf{w}^\intercal(\mathbf{f}_{ij} - \mathbf{f}_{i'j}) \geq 0,\ \forall i' \neq i\right\} \setminus \{\mathbf{0}\}$$

$\mathcal{W}_{ij}$ is the set of scoring rules for which recipient type i is top-scoring for resource type j. Our formulation in Problem 2.2 makes use of the binary variables $x_{ij} \in \{0, 1\}$ for $i \in [I]$ and $j \in [J]$ and constrains them so that $x_{ij} = 1$ if and only if the chosen linear scoring rule $\mathbf{w}$ is contained in $\mathcal{W}_{ij}$.

Problem 2.2 (Priority Allocation Optimization)**.**

$$\max_{\mathbf{w},\,\mathbf{x}} \quad \sum_{i=1}^{I}\sum_{j=1}^{J} r_{ij}x_{ij} \tag{2.2a}$$

$$\text{subject to} \quad \mathbf{e}^\intercal\mathbf{w} = 1 \tag{2.2b}$$

$$s_{ij} = \mathbf{w}^\intercal\mathbf{f}_{ij} \qquad \forall i \in [I],\ j \in [J] \tag{2.2c}$$

$$s_{ij} - s_{i'j} \geq M(x_{ij} - 1) \qquad \forall i \in [I],\ j \in [J],\ i' \neq i \tag{2.2d}$$

$$\sum_{i=1}^{I} x_{ij} \leq 1 \qquad \forall j \in [J] \tag{2.2e}$$

$$\mathbf{x} \in \{0, 1\}^{I \times J} \tag{2.2f}$$

To confirm that this formulation is correct, suppose we have computed scores using a scoring rule $\mathbf{w} \in \mathcal{W}_{ij}$. It follows that $s_{ij} - s_{i'j} \geq 0$ for all $i' \neq i$, allowing the constraint in (2.2d) to be satisfied when setting $x_{ij} = 1$. On the other hand, if $\mathbf{w} \notin \mathcal{W}_{ij}$ then there is some i' for which $s_{ij} - s_{i'j} < 0$ and setting $x_{ij} = 0$ is required to ensure the constraint holds. Note also that (2.2e) ensures at most one recipient type is matched with each resource type in case there is a tie for the top-ranked queue, and

the objective ensures that this tie is broken by selecting the match with maximum reward.

It is impossible to exactly model the requirement that $\mathbf{w} \neq \mathbf{0}$ in a MIP, but (2.2b) at least makes this solution infeasible. Since the allocation function is scale-free, the search space effectively covers all scoring rules with $\mathbf{e}^\intercal \mathbf{w} > 0$, while eliminating those with $\mathbf{e}^\intercal \mathbf{w} \leq 0$. The MIP can be solved a second time with $\mathbf{e}^\intercal \mathbf{w} = -1$ to cover $\mathbf{e}^\intercal \mathbf{w} < 0$, and covering $\mathbf{e}^\intercal \mathbf{w} = 0$ can be achieved by solving with a small perturbation of the properties.

An instance of this problem can clearly have multiple optimal solutions. Suppose that $\mathbf{w}^*$ is an optimal scoring rule to 2.2 which produces an allocation $\mathbf{y}^* = (i_1^*, \ldots, i_J^*)$. If $\mathbf{y}^*$ is the unique optimal allocation, then the set of optimal scoring rules is given exactly by (2.3):

$$\mathcal{W}^* = \left(\bigcap_{j=1}^{J} \mathcal{W}_{i_j^* j} \right) \tag{2.3}$$

On the other hand, if $\mathbf{y}^*$ is not the unique optimal allocation, then $\mathcal{W}^*$ is a subset of the optimal rules.

Finally, a remark on the structure of the properties. Let $\mathbf{F}$ refer to the matrix formed by stacking values of corresponding properties across all recipient and resource types:

$$\mathbf{F} \coloneqq \begin{bmatrix} f_{111} & \cdots & f_{11K} \\ \vdots & \ddots & \vdots \\ f_{I11} & \cdots & f_{I1K} \\ f_{121} & \cdots & f_{12K} \\ \vdots & \ddots & \vdots \\ f_{IJ1} & \cdots & f_{IJK} \end{bmatrix}$$

With this notation, $\mathbf{Fw} \in \mathbb{R}^{IJ}$ gives the scores for each recipient-resource type pair in vector form. If the columns of $\mathbf{F}$ are linearly dependent, then there exists some $\mathbf{w} \neq \mathbf{0}$ so that $\mathbf{Fw} = \mathbf{0}$. It is also possible that there is some $\mathbf{w}$ which results in all scores being identical and nonzero. Both cases are unlikely to happen in practice if it is assumed that $K << IJ$. But if Problem 2.2 is solved and all scores are returned

identical, a small $\epsilon > 0$ can be added to (2.2d) to require that the top-scoring recipient type has strictly the highest score.

2.3.3 Computational Experiments

The scalability of the formulation was tested on randomly generated data. We used $I = J$ (the same number of recipient and resource types) with $K \in \{5, 10\}$. All elements in $\mathbf{F}$ and $\mathbf{r}$ were generated independently and uniformly at random on the unit interval $[0, 1]$. For each configuration of parameters, we generated 5 independent problem instances and solved the formulation in Problem 2.2 with a time limit of 5 hours. All tests were conducted on an Intel Xeon 2.1 GHz quad-core CPU with 32 GB of RAM and solved with Gurobi 9.1.1.

Figure 2-3 shows how the range of solution times changed as the number of recipient and resource types in the instance increase. Problems with up to 30 types for $K = 5$ and up to 25 types for $K = 10$ could be solved to provable optimality within the time limit.

Figure 2-4 shows how the range of optimality gaps changed as the number of recipient and resource types varied. Optimality gaps grew quickly once the problem could not be solved within 5 hours. When $K = 10$ and the number of types grew to 50, optimality gaps were between 35% and 50%.

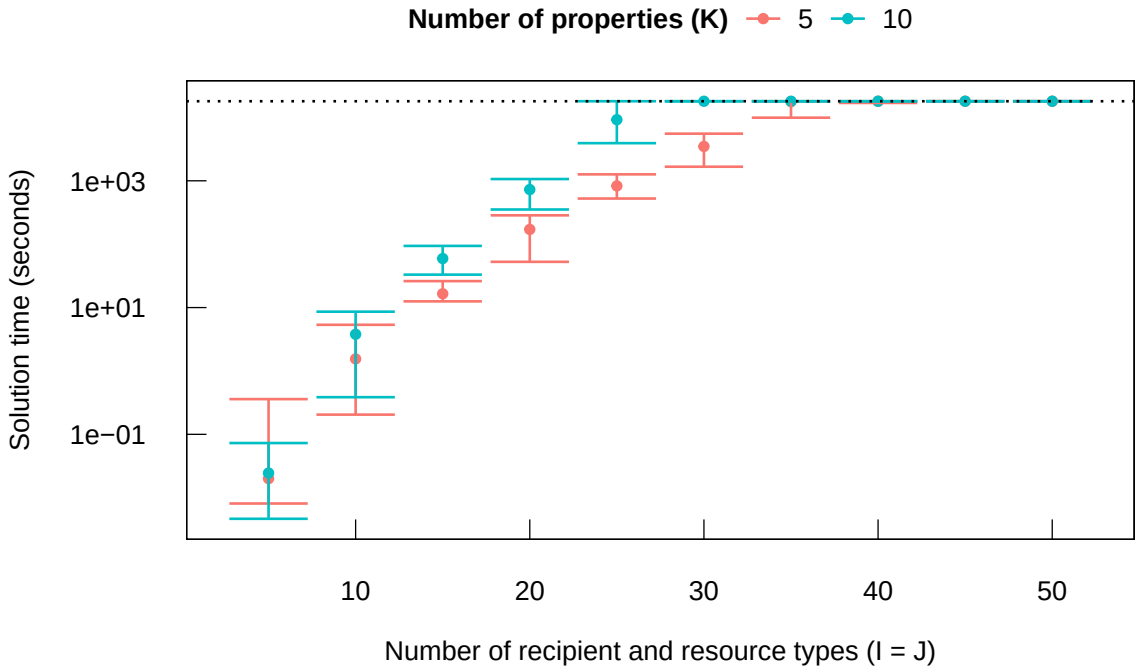

Figure 2-3: Ranges of solution times for Problem 2.2 when instances were solved to optimality. The dashed line indicates the 5 hour time limit.

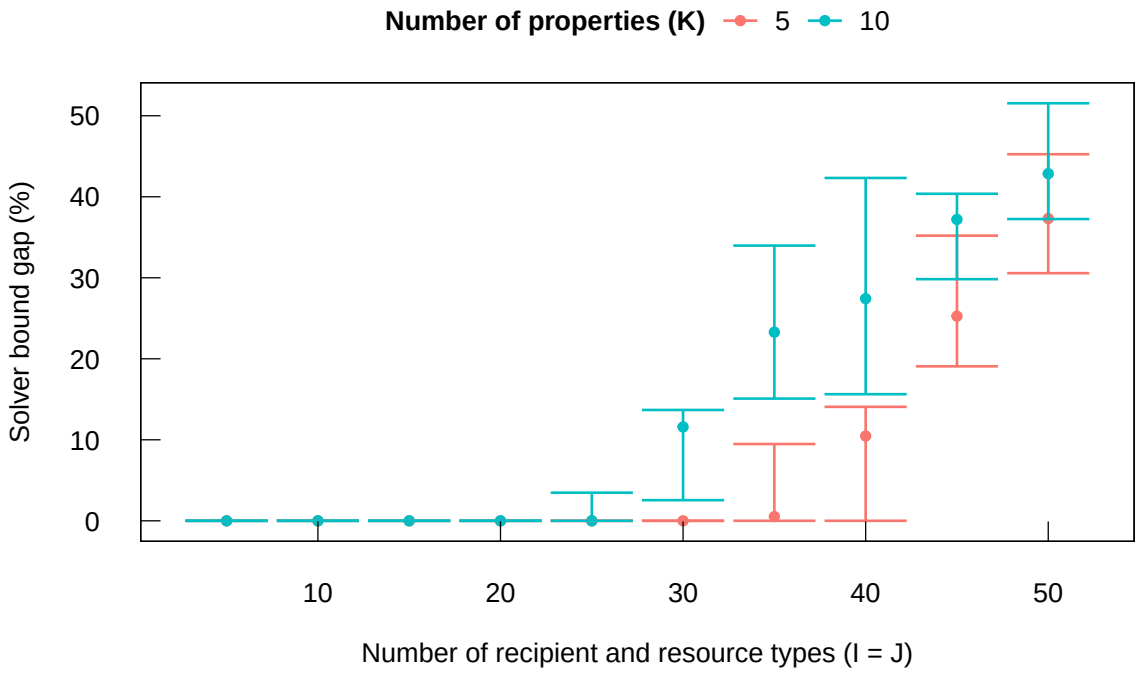

Figure 2-4: Ranges of the optimality gap for Problem 2.2 after the 5 hour time limit.

2.4 Heuristics for Priority Allocation

Though the MIP for priority allocation can be solved optimally for moderately-sized instances with a reasonable number of properties, it is useful to have heuristics that provide good solutions for larger instances. This section describes two, and provides performance bounds and experiments showing when each should be applied.

2.4.1 Projection Heuristic

The first heuristic finds a scoring rule that produces scores which are *close* to the rewards by projecting $\mathbf{r}$ onto $\overline{\mathcal{S}}$. It is similar to the approach of Bertsimas, Farias, and Trichakis [6], who solve a least-squares problem to project their equity-adjusted rewards onto the scores that can be produced by the authority in their setup.

In our setting, the motivation for this heuristic is the following: suppose there is a scoring rule $\hat{\mathbf{w}} \in \mathbb{R}^K$ for which $\mathbf{f}_{ij}^\intercal \hat{\mathbf{w}} = r_{ij}$ for all $i \in [I]$ and $j \in [J]$. Then, since it guarantees the top-ranking recipient type has the highest reward for each resource type, this scoring rule is optimal. If we instead find scores that are close to the rewards, they may be good (rather than optimal).

In this section it is more convenient to represent rewards and properties by stacking their components. Recall the previous notation used for the properties, $\mathbf{F}$, and let similar notation for the rewards be:

$$\mathbf{r} \coloneqq (r_{11}, \ldots, r_{I1}, r_{12}, \ldots, r_{IJ})$$

The *projection solution* for our setup is defined when $\mathbf{F}$ has linearly independent columns (a mild assumption) and is given in Definition 2.3. Note that the projection solution is not defined when $\mathbf{F}^\intercal \mathbf{r} = \mathbf{0}$, or equivalently $\mathbf{r} \perp \text{colspace}(\mathbf{F})$.

Definition 2.3 (Projection Solution)**.** *The projection solution,* $\hat{\mathbf{w}}(\mathbf{F}, \mathbf{r}) \in \mathbb{R}^K \setminus \{\mathbf{0}\}$*, is obtained by solving:*

$$\hat{\mathbf{w}}(\mathbf{F}, \mathbf{r}) \coloneqq \underset{\mathbf{w} \in \mathbb{R}^K}{\arg\min} \frac{1}{2} ||\mathbf{F}\mathbf{w} - \mathbf{r}||_2^2 = (\mathbf{F}^\intercal \mathbf{F})^{-1} \mathbf{F}^\intercal \mathbf{r}$$

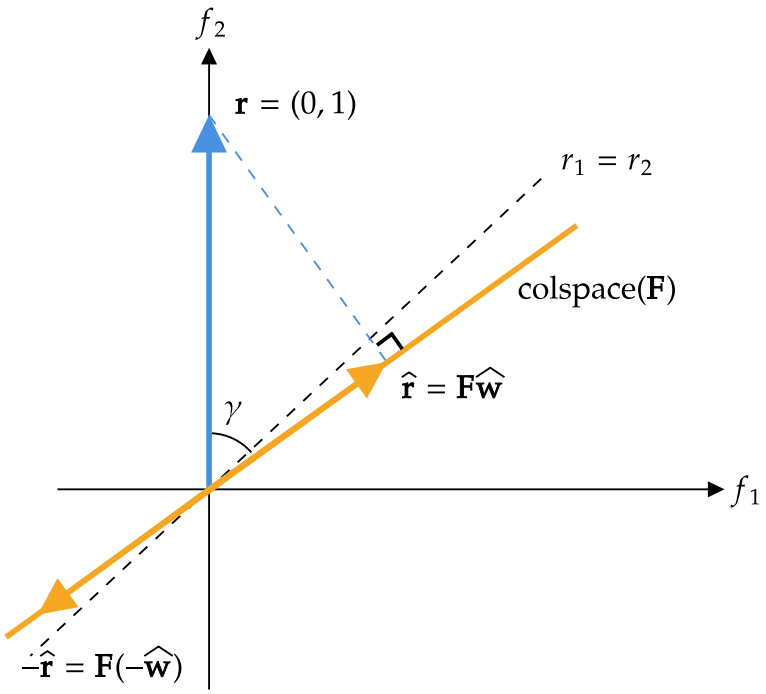

Figure 2-5: Example where projection fails. $\hat{r}_1 > \hat{r}_2$, so projection generates reward $r_1 = 0$. However, there is a scoring rule $-\hat{\mathbf{w}}$ with $-\hat{r}_1 < -\hat{r}_2$, so it is possible for the authority to earn $r_2 = 1$.

In general the projection solution can be arbitrarily far from optimality. Figure 2-5 illustrates a problem instance where projection fails for $I = 2$, $J = 1$, $K = 1$.

In Figure 2-5, the angle γ between $\mathbf{r}$ and $\overline{\mathcal{S}}$ is large and contributes to how poorly the projection solution performs. γ is a measure of how well the properties *describe* the rewards. If $\gamma = 0$, then the properties can perfectly reproduce the rewards, whereas if $\gamma = \frac{\pi}{2}$, then the properties provide no information about the rewards. We bound the performance of the projection heuristic in terms of γ, which is defined for the angle between a vector $\mathbf{r}$ and score set as:

$$\gamma = \cos^{-1}\left(\max_{\mathbf{s}\in\overline{\mathcal{S}}} \frac{\mathbf{r}^\intercal \mathbf{s}}{||\mathbf{r}|| \cdot ||\mathbf{s}||}\right) \tag{2.4}$$

The bounds consider a single resource type ($J = 1$) and drop the corresponding subscript. The bound in Proposition 2.3 (with a proof in Appendix A.3) extends easily to the case $J > 1$ when the result is interpreted as a *per-resource-type* bound.

Proposition 2.3. *Let $J = 1$, and the projected reward coefficients be given by $\hat{\mathbf{r}} \in \mathbb{R}^I$ with $\hat{r}_1 \geq \hat{r}_2 \geq \ldots \geq \hat{r}_I$. Let γ be the angle between $\mathbf{r}$ and $\overline{\mathcal{S}}$ as in (2.4).*

Let z^* be the optimal objective value of Problem 2.2, and let z be the objective value of the projection solution. The following bound holds for any angle $0 \leq \gamma < \pi/2$:

$$\frac{z^* - z}{||\mathbf{r}||} \leq \max\left(0, \frac{\hat{r}_2 - \hat{r}_1}{||\hat{\mathbf{r}}||} + \sqrt{2}\sin(\gamma)\right)$$

A second proposition writes the RHS in terms of $\mathbf{r}$ only:

Proposition 2.4. *Let $J = 1$, and the projected reward coefficients be given by $\hat{\mathbf{r}} \in \mathbb{R}^I$ with $\hat{r}_1 \geq \hat{r}_2 \geq \ldots \geq \hat{r}_I$. Let γ be the angle between $\mathbf{r}$ and $\overline{\mathcal{S}}$ as in (2.4).*

Let z^* be the optimal objective value of Problem 2.2, and let z be the objective value of the projection solution. The following bound holds for any angle $0 \leq \gamma < \pi/2$:

$$\frac{z^* - z}{||\mathbf{r}||} \leq \max\left(0, \frac{r_2 - r_1}{||\mathbf{r}||} + 2\sqrt{2}\sin(\gamma)\right)$$

The left-hand side of the bound does not have the most intuitive interpretation. Ideally, we would like to produce a bound on $(z^* - z)/z^*$ that measures the relative optimality gap. But in the absence of such a result, $||\mathbf{r}||$ replaces the denominator and acts as a proxy for the optimal value of Problem 2.2. The key observation to make is that the right-hand side increases with a known function of γ, and decreases with larger separation in the scores of the two top-ranked types.

The bound we have derived is tight for any fixed $0 \leq \gamma \leq \pi/4$, in the sense that an example can be constructed with property vectors $\mathbf{F}$ so that the left-hand side is arbitrarily close to the right-hand side in Proposition 2.3. Example 2.2 illustrates this observation. For $\gamma > \pi/4$, the bound is not necessarily tight.

Example 2.2. For a fixed $\gamma > 0$, a single resource type ($J = 1$), four recipient types ($I = 2$) and one property ($K = 1$), consider a property vector and $\hat{\mathbf{r}}$ given by:

$$\hat{\mathbf{r}}^\intercal = \mathbf{F}^\intercal = \left(\frac{1}{\sqrt{2}} + \epsilon, \frac{1}{\sqrt{2}}\right)$$

where $\epsilon > 0$ is a small positive number. Note that a basis for the subspace orthogonal to $\mathcal{S}$ is:

$$\mathbf{G} \approx \begin{bmatrix} \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

The projection solution $\hat{\mathbf{r}}$ ranks recipient type 1 highest, so the reward it obtains is given by r_1. On the other hand, there is a scoring rule ($\mathbf{w} = (0, 1)$) under which type 2 is ranked highest. So for any $\mathbf{r}$, its absolute suboptimality is $\max(0, r_2 - r_1)$.

Now consider an adversary who constructs $\mathbf{r}$ to make $\hat{\mathbf{r}}$ perform as poorly as possible, and with the angle between $\mathbf{r}$ and $\hat{\mathbf{r}}$ no more than γ. Therefore they aim to solve:

$$\max_{\delta \leq \tan(\gamma)} \quad \left(\hat{r}_1 - \delta\frac{1}{\sqrt{2}}\right) + \left(\hat{r}_2 + \delta\frac{1}{\sqrt{2}}\right)$$

For $\gamma \leq \frac{\pi}{2}$, they will choose $\delta = 1$ and $\mathbf{r} \approx (0, \sqrt{2})$. The absolute suboptimality is $r_2 - r_1 = \sqrt{2}$, and we have:

$$\frac{z^* - z}{||\mathbf{r}||} \approx \frac{\sqrt{2}}{\sqrt{2}} \approx \max\left(0, \frac{\hat{r}_2 - \hat{r}_1}{||\hat{\mathbf{r}}||} + \sqrt{2}\sin(\gamma)\right)$$

2.4.2 Lookahead Heuristic

The projection heuristic is a geometric approach to solving the problem, but we can also frame the problem as a sequence of individual allocations and make use of a combinatorial approach. This section describes such an approach and calls it the *lookahead* heuristic. It requires some new notation and a definition of the problem as a sequence of decisions – ultimately leading to formulation with a Bellman equation.

Note that a complete allocation can be constructed in J steps, where at each step a resource is selected and allocated to a recipient. At the conclusion of this process we must be left with a feasible allocation, which can be ensured by maintaining feasibility of the partial allocation as it is built up.

To this end, we let a partial allocation that has been produced in D steps be denoted by $\mathbf{z}^D \coloneqq \{(i_1, j_1), \dots (i_D, j_D)\}$. Here, j_d is the d^{th} resource to be allocated and i_d is the recipient which was allocated this resource. The set of scoring rules which are consistent with this partial allocation is:

$$\mathcal{W}^D = \bigcap_{d=1}^{D} \mathcal{W}_{i_d j_d}$$

and the set of recipients (say $\bar{\mathcal{I}}_j$) which can therefore be selected for some $j \in \overline{\mathcal{J}}$ to extend the partial allocation and maintain feasibility is given by:

$$\bar{\mathcal{I}}_j = \{i : \mathcal{W}_{ij} \cap \mathcal{W}^D \neq \emptyset\}$$

Now let the value function $f(\mathbf{z}^D)$ be the maximum reward that can be earned when resources $\overline{\mathcal{J}} \coloneqq [J] \setminus \{j_1, \dots, j_D\}$ must be allocated to complete the partial allocation $\mathbf{z}^D$. The Bellman equation describing the optimal structure of the problem is:

$$f(\mathbf{z}^D) = \max_{j \in \overline{\mathcal{J}},\, i \in \bar{\mathcal{I}}_j} \left[r_{ij} + f(\mathbf{z}^D \cup \{(i,j)\}) \right]$$

This equation says the following: at each step, the authority must choose a resource that has not been allocated ($j \in \overline{\mathcal{J}}$) and a recipient who can be matched to

this resource when restricted to scoring rules that are consistent with the partial allocation ($i \in \overline{\mathcal{I}}_j$). Their choice should maximize the sum of the reward obtained due to the pair and the reward that can be obtained by completing the partial allocation.

While $f(\mathbf{z}^D)$ cannot be computed directly, it is possible to compute an upper bound that helps to define the lookahead heuristic. This upper bound, denoted by $\overline{f}(\mathbf{z}^D)$, is simple: the sum of the maximum rewards that may be obtained for each remaining resource in $\overline{\mathcal{J}}$ while choosing only from scoring rules that are consistent with the partial allocation $\mathbf{z}^D$:

$$\overline{f}(\mathbf{z}^D) = \sum_{j \in \overline{\mathcal{J}}} \max_{i \in \overline{\mathcal{I}}_j} r_{ij}$$

The lookahead heuristic uses this upper bound to repeatedly extend a partial solution into a full solution by selecting, at each step, the pair which maximimizes the following expression:

$$\max_{j \in \overline{\mathcal{J}},\, i \in \overline{\mathcal{I}}_j} \left[r_{ij} + \overline{f}(\mathbf{z}^D \cup \{(i,j)\}) \right]$$

Proposition 2.5. *The lookahead heuristic approximates the optimal solution within a factor of $J - 1$. This bound is (arbitrarily) tight.*

A proof for the first part of Proposition 2.5 is included in Appendix A.4. Example 2.3 provides a tight example for $J = 3$ (which is easily extended for $J > 3$) where the lookahead heuristic approximates the optimal solution within a factor arbitrarily close to 2.

It is also worth noting that the lookahead heuristic finds the optimal solution when the data are generated with $\gamma = 0$. The argument is straightforward: when $\gamma = 0$, $\overline{f}(\mathbf{z}^D) = f(\mathbf{z}^D)$ for any partial solution $\mathbf{z}^D$, so the lookahead heuristic will pick pairs in decreasing order of their reward. It is difficult to find a bound on suboptimality for the lookahead heuristic in terms of γ, but this observation is the reason for testing the two heuristics on problem instances parameterized by γ in the next section.

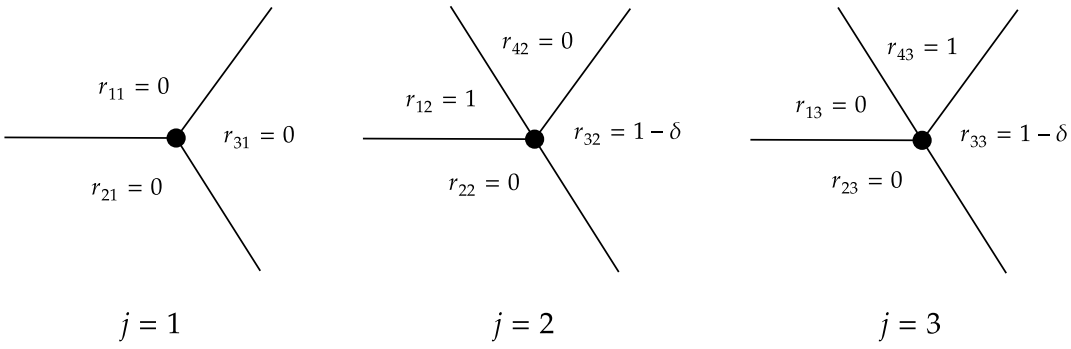

Figure 2-6: Problem data for the lookahead example. Each separate diagram shows the reward coefficients and structure of the $\mathcal{W}_{ij}$ sets for a resource type j.

Example 2.3. Suppose $I = 4$ and $J = 3$. Figure 2-6 provides an alternative representation of the problem data: instead of being in terms of the property vectors, $\mathbf{f}_{ij}$, it shows the sets $\mathcal{W}_{ij}$ where recipient i is the top-scoring type for resource j. The corresponding pair (i, j) is denoted by the reward coefficient for this match.

The lookahead heuristic begins with $\mathbf{z}^0 = \emptyset$. The unique choice for the first match is $\mathbf{z}^1 = \{(1, 1)\}$ since this produces $r_{11} + \overline{f}(\{(1, 1)\}) = 2$ and it can be readily observed that no other selection is able to match this value.

But the partial solution $\mathbf{z}^1$ can only gain a reward of one unit (from either $(1, 2)$ or $(4, 3)$) when completed into a full solution. Compared with the optimal solution $\mathbf{z} = \{(3, 1), (3, 2), (3, 3)\}$ with an objective value of $2 - 2\delta$, the approximatoin ratio of the lookahead solution is $1/(2 - 2\delta)$.

Essentially, this effect occurs when the first selection made by the lookahead solution is maximally naive. The example can be easily extended for any $J > 3$ to obtain an optimality bound arbitrarily close to $1/(J - 1)$.

2.4.3 Numerical Results

We tested the performance of both heuristics on randomly generated data. We used $I = J = 10$ and $K = 5$ so that the MIP formulation could be solved to optimality and the exact optimality gap of each heuristic could be measured. All elements in $\mathbf{F}$ were generated independently and uniformly at random on the unit interval $[0, 1]$. A random point $\hat{\mathbf{r}} \in \overline{\mathcal{S}}$ was generated along with $\boldsymbol{\delta} \in \overline{\mathcal{S}}^{\perp}$ such that $||\boldsymbol{\delta}||_2 = 1$. Then, γ was selected uniformly at random from the interval $[0, \frac{\pi}{2}]$ and the reward coefficients were set to be $\mathbf{r} = \hat{\mathbf{r}} + ||\hat{\mathbf{r}}|| \tan(\gamma)\boldsymbol{\delta}$ so that the angle between $\hat{\mathbf{r}}$ and $\mathbf{r}$ was γ. We generated 500 of these problem instances and solved the formulation in Problem 2.2 to optimality before comparing with the projection solution.

Figure 2-7 shows the optimality gap achieved by the two heuristics on each of the problem instances. The fitted line is from a Loess regression model and shaded band is the single standard error range. When $\gamma \approx 0$, both heuristics are optimal. When $\gamma = \frac{\pi}{8} = 22.5°$, the optimality gaps within a single standard error of a Loess regression model are between 2% and 13% for the lookahead heuristic and between 12% and 30% for the projection heuristic. At $\gamma = \frac{\pi}{4} = 45°$ they are between 10% and 33%, and 45% and 75% respectively. Though a bound on performance in terms of γ could not be found for the lookahead heuristic, it is the better performer in these experiments.

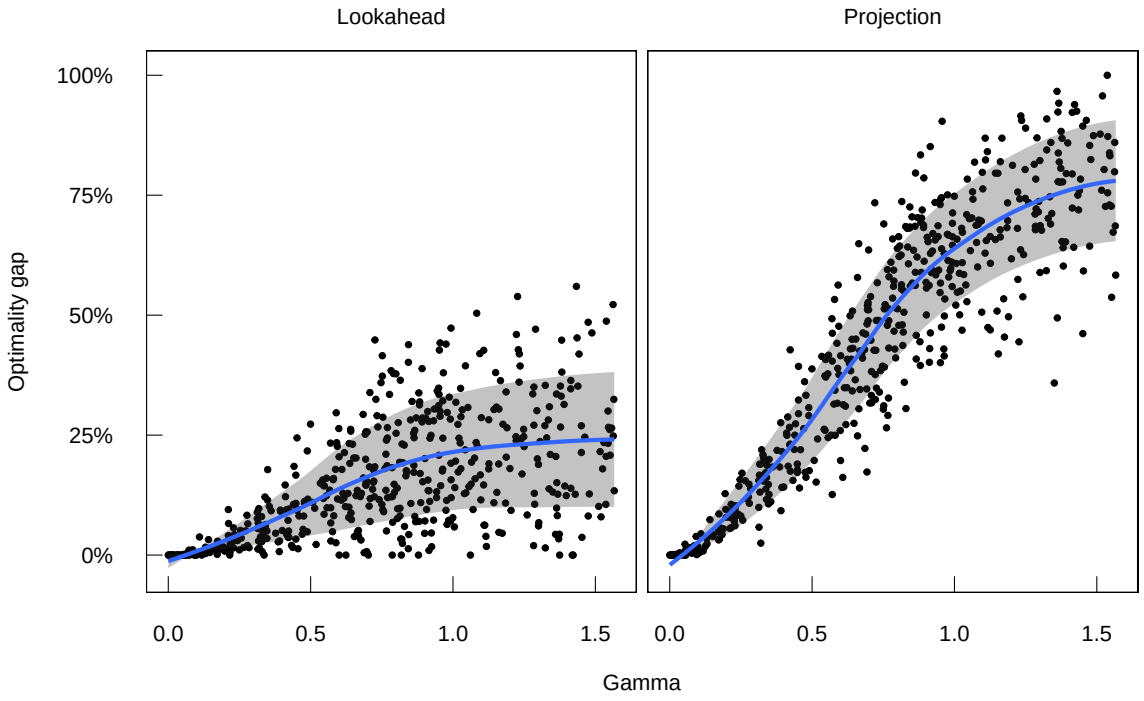

Figure 2-7: Optimality gap for the two heuristics applied to Problem 2.2 as the value of γ varied.

2.5 Scoring Rules for Individual Allocation

In practice, scoring systems allocate individual resources to individual recipients. Our type-based model is an accurate one in systems where every individual can be assigned to a types and individuals with the same type share identical properties. On the other hand, there are systems where this is not possible and individuals which share a type only have *similar* properties. The kidney allocation system in the US includes time spent on the waiting list as a property, and since this variable is continuous, individuals cannot be divided into a small set of types which share identical properties. This section explains how the MIP and heuristics for priority allocation optimization may fail when used in this type of system, and suggests an approach to find a more appropriate scoring rule.

2.5.1 Modeling Individual Allocation

We first establish a model which allocates individual resources to individual recipients and use it to make the idea of a discrepancy relative to the previous type-based model more precise. The model is very simple, allocating a finite set of resources to a finite set of recipients, but the ideas it illustrates are relevant even when the model is extended to a more complicated setting with stochasticity and infinite streams of recipient and resource arrivals.

Individual resources are indexed by $\mathcal{B}$ and individual recipients by $\mathcal{A}$. The properties relating a pair of individuals, $(a, b) \in \mathcal{A} \times \mathcal{B}$, are denoted $\bar{\mathbf{f}}_{ab} \in \mathbb{R}^K$. For some fixed scoring rule $s : \mathbb{R}^K \to \mathbb{R}$, the score relating two individuals is $\bar{s}_{ab} = s(\bar{\mathbf{f}}_{ab})$ and the reward obtained by a match is $\bar{r}_{ab} \in \mathbb{R}_+$.

The model retains the notion of types: I recipient types and J resource types. Let the type of recipient a be $\alpha_a \in [I]$ and the type of resource b be $\beta_b \in [J]$. Assume that it is possible to find a *representative* property vector that relates a pair of types, $(i, j) \in [I] \times [J]$, and call this $\mathbf{f}_{ij} \in \mathbb{R}^K$. In keeping with this notation, a score relating two types is $s_{ab} = s(\mathbf{f}_{ab})$ and a representative reward for matching two types is $r_{ij} \in \mathbb{R}_+$. Let the rate λ_i be the fraction of recipients in $\mathcal{A}$ who are type i (and let

μ_j be computed similarly for the resources).

The procedure for constructing an allocation of individuals is a natural one. Let $\mathcal{M} \subseteq \mathcal{A}$ be the set of recipients who have already been matched, starting with $\mathcal{M} = \emptyset$. The resources arrive in order. When $b \in \mathcal{B}$ arrives, scores are computed for all recipients who remain, $\mathcal{A} \setminus \mathcal{M}$, and an allocation mechanism is applied to these scores to select a recipient a and update $\mathcal{M} \leftarrow \mathcal{M} \cup \{a\}$. In this context, the allocation mechanism is a set of functions:

Definition 2.4 (Allocation Mechanism)**.** *An allocation mechanism is a family of (possibly randomized) functions,* $\{g_m\}_{m=1}^{|\mathcal{A}|}$*. Each of the functions has the signature:*

$$g_m : \mathbb{R}^m \mapsto [m]$$

g_m takes takes as input the vector of m scores from the recipients who are waiting for a resource, and selects one of these recipients to allocate the next resource.

Let $\bar{x}_{ij}(\mathbf{w})$ be the fraction of all individual type j resources allocated to individual type i recipients, and note that we have parameterized the input with the scoring rule rather than scores, since a scoring rule generates an allocation in both the individual and type-based models. Now, by specifying some allocation function $x_{ij}(\mathbf{w})$ in the type-based model and using the properties $\mathbf{F}$, rewards $\mathbf{r}$, and rates $\boldsymbol{\lambda}$, $\boldsymbol{\mu}$ previously defined, we obtain a corresponding type-based model. The discrepancy maps a scoring rule onto a measure of difference between the individual allocations and the underlying type-based allocations:

$$d_p(\mathbf{w}) = ||\bar{\mathbf{x}}(\mathbf{w}) - \mathbf{x}(\mathbf{w})||_p$$

2.5.2 Where Do Discrepancies Arise?

Under the priority allocation mechanism, a large discrepancy occurs when within-type variation in properties (and therefore scores) causes the type of top-ranked individual recipients to differ from the type-based model. Example 2.4 shows that even arbitrarily small within-type variation can lead to a scoring rule that is optimal for the underlying type-based model performing poorly when used to allocate individuals.

If a scoring rule leads to a large discrepancy, we must accept that the type-based model is not a good approximation of the individual allocation under this scoring rule. On the other hand, if it leads to a small discrepancy then the approximation is likely a good one. If it is known that the scoring rule is optimal in the type-based model and that the approximation is good, then we can hope it also performs well when used to allocate individuals. This chain of reasoning is the basis for the heuristic presented in the next subsection.

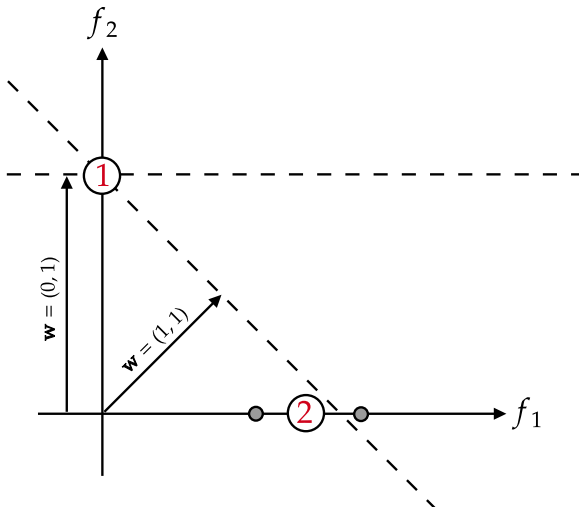

Figure 2-8: Illustration of the setup in Example 2.4.

Example 2.4. This example preserves notation from 2.2 and uses a set of recipients $\mathcal{A} = \{1, 2, 3\}$. The individual data are:

$$\bar{\mathbf{f}}_1 = (0, 1),\ \bar{r}_1 = 2 \qquad \bar{\mathbf{f}}_2 = (1 - 3\delta, 0),\ \bar{r}_2 = 1 - \delta \qquad \bar{\mathbf{f}}_3 = (1 + \delta, 0),\ \bar{r}_3 = 1 + \delta$$

Red labels indicate property vectors in the approximate dataset, and the small grey points indicate variation of individuals. There are two recipient types, with $\alpha_1 = 1$ and $\alpha_2 = \alpha_3 = 2$. The representative data are taken as the within-type averages:

$$\mathbf{f}_1 = (0, 1),\ r_1 = 2 \qquad \mathbf{f}_2 = (1 - \delta, 0),\ r_2 = 1$$

The scoring rule $\mathbf{w}_1 = (1, 1)$ is optimal when Problem 2.2 is solved on the approximate dataset as it ranks recipient 1 (with the maximum reward) highest. But, if $\mathbf{w}_1$ is used to allocate resources in the original dataset, recipient 3 has the highest score and will be allocated the resource (though they do not have the highest reward).

A better scoring rule is $\mathbf{w}_2 = (0, 1)$. This choice places more weight on the property that differentiates the two recipient types in the fluid model and separates the scores of individuals from each type further. Even when the properties of type 2 recipients vary about $\mathbf{f}_2$, the resource will be allocated to a recipient of type 1.

2.5.3 Minimizing Discrepancy

It is quite natural to form an optimization problem that balances maximizing the objective of the type-based model with minimizing the discrepancy between the individual and type-based models. By letting $\eta \geq 0$ determine this trade-off, we obtain:

$$\max_{\mathbf{w}} \quad \sum_{i=1}^{I} \sum_{j=1}^{J} r_{ij} x_{ij}(\mathbf{w}) - \eta d_p(\mathbf{w})$$

Clearly, when $\eta = 0$, the problem reduces to Problem 2.2 and there is a tractable MIP formulation available. For any other value $\eta > 0$, the problem is harder to solve given that $\bar{\mathbf{x}}(\mathbf{w})$ is a complex function that depends on potentially many individual allocations. After all, if we had a tractable representation, we could simply optimize $\sum_{i=1}^{I} \sum_{j=1}^{J} r_{ij} \bar{x}_{ij}(\mathbf{w})$.

However, a tractable problem can be obtained by replacing the discrepancy term with a surrogate which measures the minimum difference between the top-ranked and second-ranked scores of the recipient-resource type pairs. Maximizing this quantity is likely to reduce the size of the discrepancy (by our previous discussion). More precisely, define:

$$\bar{d}(\mathbf{w}) = \min_{j \in [J]} \left(\mathbf{s}_j(\mathbf{w})^{(1)} - \mathbf{s}_j(\mathbf{w})^{(2)} \right)$$

where $\mathbf{s}_j(\mathbf{w})^{(i)}$ is the ith highest-ranked recipient type for resource j. Our new problem becomes:

$$\max_{\mathbf{w}} \quad \sum_{i=1}^{I} \sum_{j=1}^{J} r_{ij} x_{ij}(\mathbf{w}) + \eta \bar{d}(\mathbf{w})$$

$\bar{d}(\mathbf{w})$ remains a nontrivial function to model. The remainder of this section shows how to solve the problem efficiently for the regime where $\eta \to 0$, in which case we obtain a linear program. Appendix A.5 provides a linear MIP formulation for the general case where $\eta > 0$.

In particular, when $\eta \to 0$, the problem becomes to maximize $\bar{d}(\mathbf{s})$ over the set of optimal solutions to Problem 2.2, denoted $\mathcal{W}^*$ and obtained from (2.3). A formulation to do this is presented in Problem 2.3.

Problem 2.3 (Margin Formulation)**.**

$$\max_{\mathbf{w},\,\mathbf{s},\,\boldsymbol{\gamma},\,z} \quad z \tag{2.5a}$$

$$\text{subject to} \quad z \leq s_{i_j^* j} - \gamma_j \qquad \forall j \in [J] \tag{2.5b}$$

$$\gamma_j \geq s_{ij} \qquad \forall j \in [J],\ i \neq i_j^* \tag{2.5c}$$

$$s_{ij} = \mathbf{f}_{ij}^{\intercal}\mathbf{w} \qquad \forall i \in [I],\ j \in [J] \tag{2.5d}$$

$$\mathbf{w} \in \mathcal{W}^* \tag{2.5e}$$

Since the objective is to maximize z, the γ_j terms in (2.5b) are set as small as the constraints allow. (2.5c) therefore ensures that γ_j takes on the score of the second-ranked recipient type for resource type j, and so each right-hand side term in (2.5b) measures a score difference between the optimal recipient type and the second-ranked type for that resource type j. The objective maximizes the minimum difference.

Problem 2.3 does not explicitly take into account the individual variation of recipients and resources around the representative types – it only aims to maximize score differences according to the representative types themselves. A more useful approach would take into account the details of the individual variation, whilst also remaining tractable to solve.

Problem 2.4 is a formulation which addresses this issue. It relies on defining two well-known functions (for fixed values of L_{ij}), where $\mathcal{D}_{ij} = \{(a,b) \in \mathcal{A} \times \mathcal{B} : \sigma(a) = i,\ \sigma(b) = j\}$ is the set of pairs of individual recipients and resources with the corresponding types.

Definition 2.5 (Mean Top-L Function)**.**

$$g_{ij}(\mathbf{s}) = \max_{\mathbf{t}} \quad \frac{1}{L_{ij}} \sum_{(a,b)\in\mathcal{D}_{ij}} s_{ab} t_{ab}$$

$$\text{subject to} \quad \mathbf{e}^{\intercal}\mathbf{t} = L_{ij}$$

$$\mathbf{t} \in \{0,1\}^{|\mathcal{D}_{ij}|}$$

Definition 2.6 (Mean Bottom-L Function)**.**

$$h_{ij}(\mathbf{s}) = \min_{\mathbf{t}} \quad \frac{1}{L_{ij}} \sum_{(a,b)\in\mathcal{D}_{ij}} s_{ab}t_{ab}$$

$$\text{subject to} \quad \mathbf{e}^\intercal\mathbf{t} = L_{ij}$$

$$\mathbf{t} \in \{0,1\}^{|\mathcal{D}_{ij}|}$$

These functions represent the average of the L_{ij} largest and smallest scores respectively, over all property vectors in the original instance associated with a particular recipient type i and resource type j.

g_{ij} and h_{ij} can be formulated as piecewise linear convex and concave functions, respectively [27], and we use them to define the convex optimization in Problem 2.4 for some fixed values of L_{ij}. Note that when $L_{ij} = |\mathcal{D}_{ij}|$, the functions return the mean score of the pairing with recipient type i and resource type j. If the mean function had been used to compute the approximate instance, then we recover the formulation in Problem 2.3.

Problem 2.4 (Top/Bottom-L Formulation)**.**

$$\max_{\mathbf{w},\,\mathbf{s},\,\boldsymbol{\gamma},\,z} \quad z \tag{2.6a}$$

$$\text{subject to} \quad z \le h_{i_j^* n}(\mathbf{s}) - \gamma_j \qquad \forall j \in [J] \tag{2.6b}$$

$$\gamma_j \ge g_{ij}(\mathbf{s}) \qquad \forall j \in [J],\ i \ne i_j^* \tag{2.6c}$$

$$s_{ij} = \mathbf{f}_{ij}^\intercal\mathbf{w} \qquad \forall i \in [I],\ j \in [J] \tag{2.6d}$$

$$\mathbf{w} \in \mathcal{W}^* \tag{2.6e}$$

This formulation modifies Problem 2.3 to take into account only the most *influential* individuals for a particular choice of the scoring rule $\mathbf{w}$. The individuals from suboptimal recipient types who have the highest scores and the individuals from optimal recipient types who have the lowest scores are the most likely candidates for

misclassifications – and it is these which the formulation aims to separate. Since g_{ij} and h_{ij} are piecewise linear functions, the problem may be easily reformulated as a linear program.

2.5.4 Numerical Experiments

We tested this approach on synthetic data for $I = J = 10$ and $K = 5$. Property vectors for each recipient-resource type pairing were first selected independently and uniformly at random on $[0, 1]^K$. Then, the dataset of individuals was generated by adding zero-mean Gaussian noise with a diagonal covariance matrix whose entries were generated independently and uniformly at random on $[0, \eta]^K$ for some *variability* parameter $\eta > 0$. The rewards were generated as if they were the $(K+1)$th property.

A baseline scoring rule was obtained by solving Problem 2.2 for the underlying type-based instance, and this rule was compared to the modified rule returned by the solution to Problem 2.4. Comparisons were made by using both rules to simulate allocations on the original data instance with individuals under priority allocation where each resource was allocated to the top-scoring recipient in the system at that point. The values of L_{ij} varied according to a percentage of $|\mathcal{D}_{ij}|$.

We tracked the discrepancy in allocation, $d_1(\mathbf{w})$, under both the baseline and modified scoring rules and report the ratio of these values across data instances in Figure 2-9. The modified solution gives consistently smaller discrepancies than the baseline solution, but as variability in the data increases this effect becomes less pronounced. At a variability of 0.02 the median reduction in discrepancy is 40% and at a variability of 0.1 the median reduction is 10%.

We see a similar, though smaller, effect on reward. The modified solution leads to an improved reward for all variabilities, though this becomes less as variability increases. The median improvement in reward earned is 5% for a variability of 0.02 and 2% for a variability of 0.1. On both metrics, there was little to distinguish the effect of the different fractions used to compute each L_{ij}.

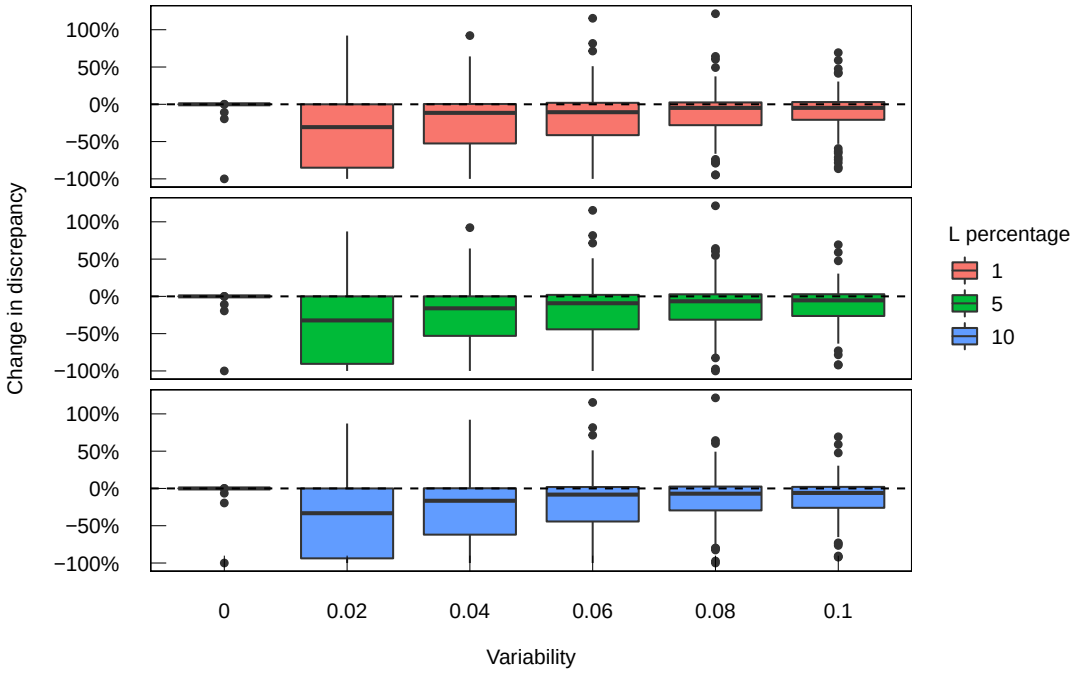

Figure 2-9: Changes in allocation discrepancy of the modified rule relative to the baseline rule, as the variability in the synthetic data increases.

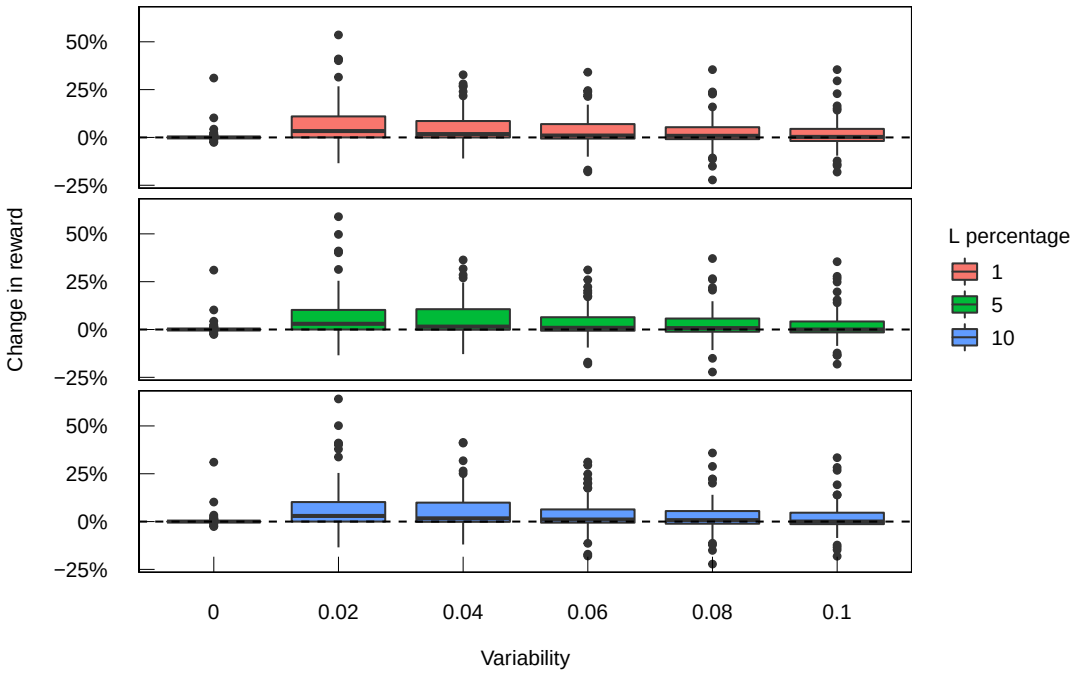

Figure 2-10: Changes in reward earned by the modified rule relative to the baseline rule, as data variability increases.

2.6 Lottery Allocation

The priority allocation mechanism may not always be the appropriate choice for a scoring system. It has the fundamental property that when two recipients have nearly identical scores, the one with the larger score receives the resource ahead of the other – but it is easy to imagine a setting where two recipients with nearly identical scores should have a similar random chance of receiving the resource. Systems which implement a random procedure such as this are referred to as *lottery* systems. Leaving the outcome of an allocation up to chance appeals to an intuitive human notion of fairness, which helps to explain why these systems are proposed time and again – including, most topically, for distributing scarce medical resources during the COVID-19 pandemic [28, 29].

This section focuses on a variation of a lottery system called a *weighted lottery* [30] in which the probabilities of receiving a resource may be different for different recipients. More precisely, the lottery mechanism we study computes scores for each recipient and then randomly allocates a resource among them with probabilities proportional to their scores: if $s_{ab} = 1$ and $s_{a'b} = 2$, then recipient a' has twice the chance of receiving resource b as recipient a.

There are many real-world examples of lottery systems that make use of scoring rules. One system was recently implemented in Philadelphia to distribute scarce medications for treating COVID-19 patients [31, 32, 33]. Another example is the Dutch system for allocating admissions to medical schools, which ran until 2017 as a weighted lottery with student scores equal to their high-school GPAs, and is proposed to be reintroduced in 2023 [34].

2.6.1 Defining the Allocation Function

The first step in modeling a lottery allocation mechanism is to define the corresponding allocation function in the fluid model and the scores on which it is defined. We start with the scores, which must simply be nonnegative and nonzero (so that the

probabilities of allocation can be proportional to them):

$$\mathcal{S} = \left\{\mathbf{s} \in \mathbb{R}^{I \times J} : \mathbf{s} \geq \mathbf{0}\right\} \setminus \{\mathbf{0}\}$$

The scores chosen by the authority must also satisfy this property:

$$\overline{\mathcal{S}} = \{\mathbf{F}\mathbf{w} \geq \mathbf{0} : \mathbf{w} \neq 0\} \setminus \{\mathbf{0}\}$$

Next, we fix some some scores, and attempt to define the allocation function $x_{ij}(\mathbf{s})$ so that it represents the fraction of resources of type j that are allocated to recipients of type i. Note that the total *weight* held collectively by the queue of recipient type i is $s_{ij}L_{ij} = \frac{s_{ij}}{q_i}\left(\lambda_i - \sum_{k=1}^{J} \mu_k x_{ik}\right)$. The allocation fractions must therefore satisfy the following expression for each recipient type i and resource type j (written in terms of the operator T_{ij}):

$$x_{ij} = \frac{s_{ij}L_{ij}}{\sum_{l=1}^{I} s_{lj}L_{lj}} = \frac{\frac{s_{ij}}{q_i}\left(\lambda_i - \sum_{k=1}^{J} \mu_k x_{ik}\right)}{\sum_{l=1}^{I} \frac{s_{lj}}{q_l}\left(\lambda_l - \sum_{k=1}^{J} \mu_k x_{lk}\right)} \coloneqq T_{ij}(\mathbf{x}) \tag{2.7}$$

This expression does not give the allocation fractions in closed form. In fact, computing the fractions amounts to finding a fixed point of the operator $T(\mathbf{x})$. It is not immediately clear this operator has a unique fixed point, or that it has a fixed point at all. But if the fractions $\mathbf{x}$ are to represent a valid allocation *function* from the set of scores onto $\mathcal{X}$, then a fixed point must both exist and be unique.

Let us first establish the *existence* of a fixed point. It is clear that $T(\mathbf{x}) \in \mathcal{X}$ when $\mathbf{x} \in \mathcal{X}$ due to the scarcity assumption (the numerators in $T_{ij}(\mathbf{x})$ are nonnegative) and the normalization terms (which ensure $\sum_{i=1}^{I} T_{ij}(\mathbf{x}) = 1$). Since $\mathcal{X}$ is compact and convex, and $T(\mathbf{x})$ is continuous, the Brouwer fixed-point theorem applies and ensures that T has at least one fixed point.

Next we turn to *uniqueness* of the fixed point, though a proof remains elusive. T is not a contraction mapping and therefore the Banach fixed-point theorem cannot be applied. On the other hand, Appendix A.6 provides the results of extensive numerical

testing showing that Banach-Picard iterations converge to a unique fixed point in all instances and with generally very few iterations. We strongly suspect that T has a unique fixed point, and note that this point is easy to compute.

2.6.2 Optimization Formulation

Though an allocation function satisfying the ratios in (2.7) is likely to exist, we do not have this function in closed form. This subsection formulates the optimization problem faced by the authority as a nonconvex bilinear program. This class of problems are known to be NP-hard (and generally hard to solve in practice) but modern solvers make use of a MIP reformulation that can solve instances of reasonable size either to optimality or with a well-quantified optimality gap. We include numerical results of our formulation on the same synthetic data that were used in Section 2.3.

The key to the formulation is to introduce variables $\boldsymbol{\sigma}$ and constrain them to represent the normalization terms appearing in (2.7). The formulation is provided in Problem 2.5 and followed with a justification of its correctness:

Problem 2.5 (Lottery Allocation Optimization)**.**

$$\max_{\mathbf{w},\,\mathbf{s},\,\mathbf{x},\,\boldsymbol{\sigma}} \quad \sum_{j=1}^{J}\sum_{i=1}^{I} r_{ij}x_{ij} \tag{2.8a}$$

$$\text{subject to} \quad \sigma_j x_{ij} = \frac{s_{ij}}{q_i}\left(\lambda_i - \sum_{j'=1}^{J}\mu_{j'}x_{ij'}\right) \qquad \forall i \in [I],\ j \in [J] \tag{2.8b}$$

$$\sum_{i=1}^{I} x_{ij} = 1 \qquad \forall j \in [J] \tag{2.8c}$$

$$s_{ij} = \mathbf{w}^\intercal \mathbf{f}_{ij} \qquad \forall i \in [I],\ j \in [J] \tag{2.8d}$$

$$\mathbf{e}^\intercal \mathbf{w} = 1 \tag{2.8e}$$

$$\mathbf{s}, \mathbf{x}, \boldsymbol{\sigma} \geq \mathbf{0} \tag{2.8f}$$

(2.8b) ensures that the allocation fractions are constrained correctly provided that

each σ_j takes on the correct value of the normalization term. To confirm that the σ_j variables are indeed defined correctly, note that for each $j \in [J]$:

$$\sigma_j = \left(\sum_{i=1}^{I} x_{ij} \right) \sigma_j = \sum_{i=1}^{I} x_{ij} \sigma_j = \sum_{i=1}^{I} \frac{s_{ij}}{q_i} \left(\lambda_i - \sum_{j'=1}^{J} \mu_{j'} x_{ij'} \right)$$

where the first equality follows from (2.8c) and the third from (2.8b).

The normalization constraint in (2.8e) is again included to ensure that the resulting scores lie in $\mathcal{S}$. The same comment as for the priority allocation optimization problem applies: the problem can be solved again for $\mathbf{e}^\intercal \mathbf{w} = -1$ and perturbations of the properties to cover the entire search space of $\mathbf{w}$.

We tested the scalability of this formulation on data generated randomly in the same way as in Section 2.3.3, with the addition of random $\boldsymbol{\mu}$ values from $[0, 1]$ and random $\boldsymbol{\lambda}$ values from $[0, 1]$ which were then scaled to satisfy Assumption 2.1. All tests were conducted with a time limit of 5 hours on an Intel Xeon 2.1 GHz quad-core CPU with 32 GB of RAM and solved with the nonconvex bilinear procedure in Gurobi 9.1.1.

Figure 2-11 shows how the ranges of solution times varied as the number of recipient and resource types in the instance increased. There was little to distinguish between $K = 5$ and $K = 10$, and problems with up to approximately 15 types could be solved to optimality within the time limit. While this is fewer than the 30 types that could be solved to optimality for priority allocation with $K = 5$, it still provides a reasonable amount of modeling flexibility.

Figure 2-12 shows how the ranges of optimality gaps varied as the number of types increased. Once again, gaps were similar for both $K = 5$ and $K = 10$. However, these gaps grew quickly once the problem could not be solved within 5 hours. When $K = 5$ and the number of types was greater than 20, the optimality gaps were between 75% and 100%.

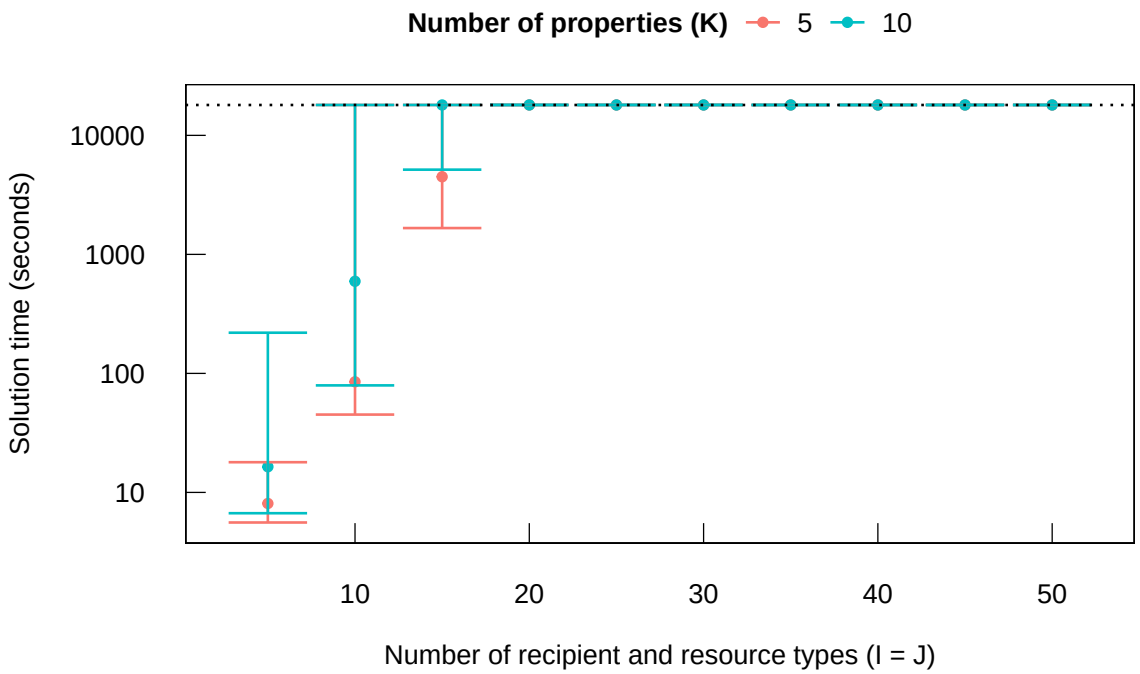

Figure 2-11: Ranges of solution times for Problem 2.5. Dashed line indicates the 5 hour time limit.

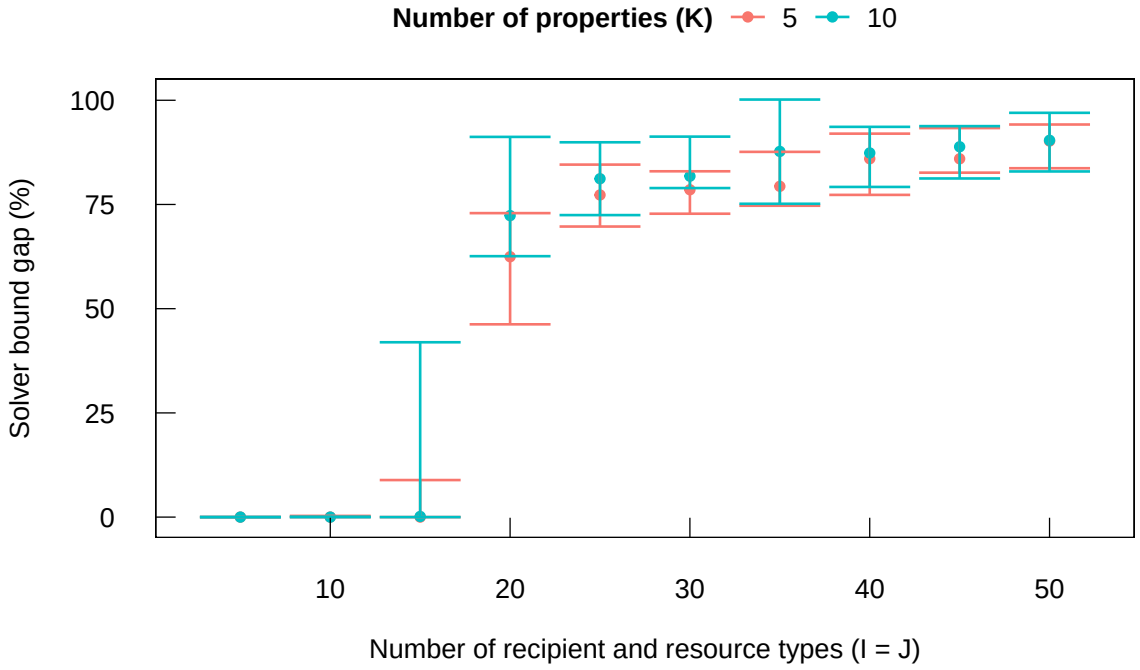

Figure 2-12: Ranges of the optimality gap for Problem 2.5 after the 5 hour time limit.

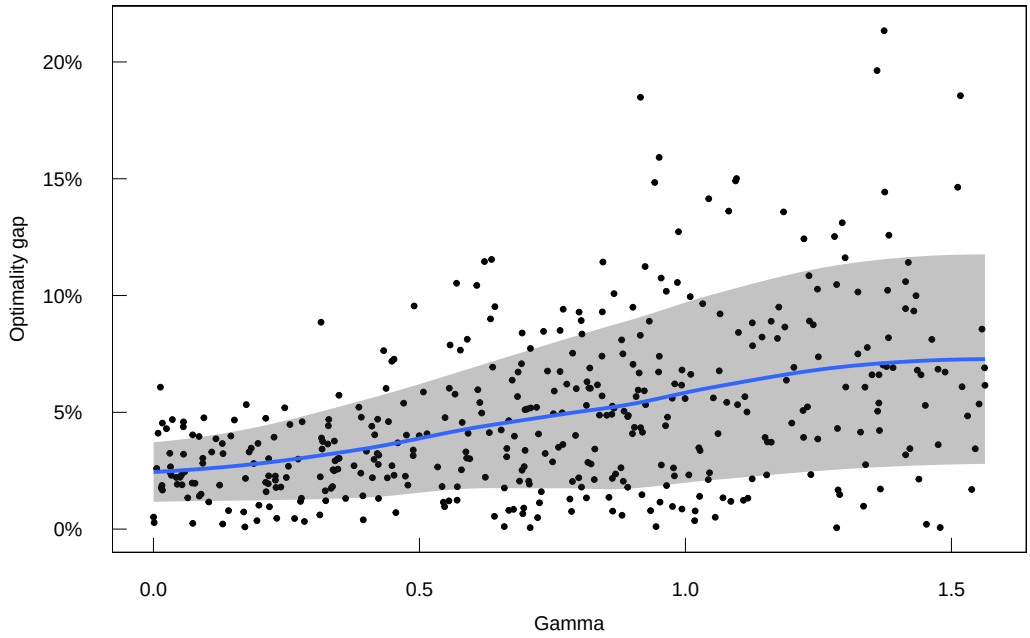

Figure 2-13: Optimality gap for the projection solution, applied to Problem 2.5.

2.6.3 Heuristic Performance

Problem 2.5 is difficult to solve, and it is therefore useful to study heuristics for solving it. While the lookahead heuristic that has been previously introduced cannot be easily re-purposed for this problem, the projection heuristic can.

The motivation for applying the projection heuristic to lottery allocation is similar, though not identical, to before. Suppose that we are able to reproduce the reward coefficients with a scoring rule $\hat{\mathbf{w}} \in \mathcal{W}$. Then, for any given resource type j, a recipient type that earns a high reward will tend to receive a larger share of the resource and a recipient type that earns a low reward will tend to receive a smaller share (which should also likely be true in the optimal solution).

The techniques previously used to bound performance of the projection heuristic do not apply to the lottery allocation problem, and we turn instead to the computational analysis in Figure 2-13. Results are shown for the same instances that were generated in Section 2.4.3. When $\gamma = \frac{\pi}{8} = 22.5°$ the solutions obtained from the projection heuristic have optimality gaps within a single standard error of between 2% and 6%. When $\gamma = \frac{\pi}{4} = 45°$ they are between 2% and 8%.

2.6.4 Modifying for Individual Allocation

In this subsection, we briefly consider how to approach the problem of allocating individuals under the lottery system by modifying a solution from the full optimization problem. Section 2.5.3 is a good starting point, but the surrogate discrepancy function we included in the objective for priority allocation no longer applies.

For the lottery system, discrepancy is minimized when the scores of individual (a, b) pairs are similar to the scores of their corresponding type pair (α_a, β_b). This observation motivates us to select the surrogate discrepancy function as a least-squares term that compares the scores of the individuals and their representative types, in order to get the problem:

$$\max_{\mathbf{w}} \quad \sum_{i=1}^{I} \sum_{j=1}^{J} r_{ij} x_{ij}(\mathbf{w}) - \eta \sum_{a \in A} \sum_{b \in \mathcal{B}} \left(\bar{s}_{ab}(\mathbf{w}) - s_{\alpha_a \beta_b}(\mathbf{w})\right)^2$$

For general $\eta > 0$, this is a mixed-integer second order cone program (SOCP) and is solvable by several leading commercial solvers, including Gurobi. Our numerical tests focused on the regime $\eta \to 0$, where the problem reduces to a SOCP. We use the same data instances that were used when testing the priority allocation function in Section 2.5. Figure 2-14 shows ratios of the discrepancy under the baseline solution and the modified solution, and Figure 2-15 shows similar for the reward. The heuristic did not improve either discrepancy or reward earned.

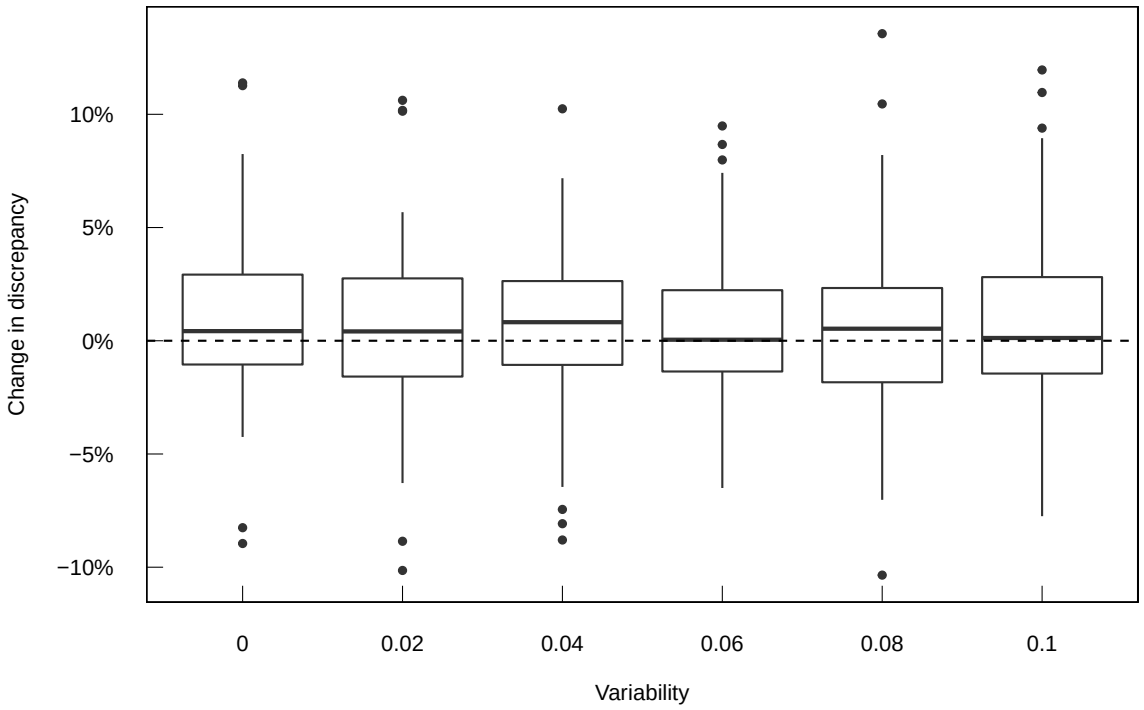

Figure 2-14: Ratios of the discrepancy under the modified rule relative to the baseline rule from the lottery model, as data variability changes.

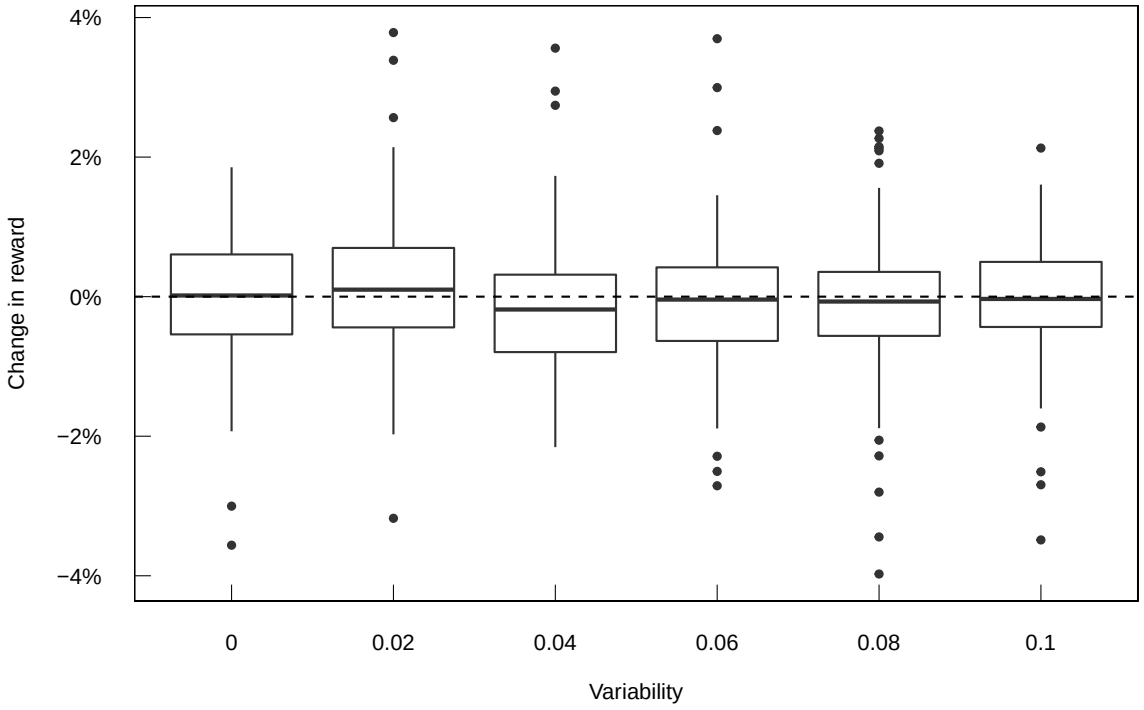

Figure 2-15: Objective ratio of the modified rule relative to the optimal rule for the lottery model, as data variability changes.

2.7 Ex-Post Equity in Allocations

At various points in this chapter, we have highlighted some desirable aspects of a system for allocating scarce resources – for example, the three principles of equity identified by Young [5] and the role of chance in a fair mechanism. Each of these properties relate to the *process* followed by the system (ex-ante) rather than the outcome of the allocation (ex-post). This section provides some comments on modifying our approaches to achieve ex-post equity.

In our context, an ex-post equity constraint requires the allocation fractions to satisfy some property. For example: the authority may wish to constrain the proportion of recipients who receive a resource and have some attribute to lie within a certain range. If the first and second recipient types are the only ones with this attribute, then lower and upper bounds can be enforced with a simple linear constraint on the allocation variables:

$$L \leq \sum_{j=1}^{J} \frac{1}{J}(x_{1j} + x_{2j}) \leq U$$

The class of linear constraints is large and covers several flavors of bound on the proportion or average value of an attribute in both the recipient types which are allocated resources and those which are not. These constraints are also interpretable and intuitive – Bertsimas, Farias, and Trichakis [6] note that it is a type favoured by policymakers in practice.

To be more precise, we may associate an additional set of N attributes with each recipient type. Let $a_{in} \in \mathbb{R}$ be the nth attribute associated with the ith recipient type. These attributes may be distinct from the properties in each $\mathbf{f}_{ij}$ or they may overlap. The key distinction is that the properties in $\mathbf{f}_{ij}$ are *unprotected* – the system may compute priority scores on their basis. Some of the attributes in each $\mathbf{a}_i$ may be *protected*, meaning that although they can be used to define equity constraints, they cannot be used as a basis for computing priority scores. Race and gender are two examples of attributes generally considered to be protected in most practical systems.

Enforcing linear ex-post equity constraints within our optimization problems simply amounts to adding linear constraints on the allocation fractions. This does not alter the problem class of the formulations. The priority allocation function still gives rise to a linear MIP and the lottery allocation function still gives rise to a nonconvex bilinear program, though the additional constraints presumably increase the time required to solve instances to optimality. However, it is important to note that satisfying even a single ex-post equity constraint is as hard as the complete optimization problem. This makes enforcing such a constraint in either of the two heuristics we have discussed impossible.

An approximate way of dealing with this trouble is provided by Bertsimas, Farias, and Trichakis [6], where it was used successfully. Here, an offline matching problem with ex-post equity constraints is solved to obtain optimal dual variables. The reward coefficients are then modified with these dual variables to obtain *equity-adjusted* coefficients under which the optimal solution to a standard bipartite matching problem satisfies the constraints. Using these equity-adjusted coefficients as inputs into the projection and lookahead heuristics may provide a rough method for obtaining allocations that approximately satisfy the equity constraints.

2.8 Discussion and Conclusions

This chapter introduced a general model for allocating a set of heterogeneous resource types among a set of heterogeneous recipient types on the basis of scores computed from observable properties. We used the model to define an optimization problem faced by the authority when tasked with choosing the scores from a restricted set, and analysed specific instances of the optimization problem associated with two real-world allocation mechanisms: priority allocation and lottery allocation.

Both optimization problems are hard to solve, but tractable with modern optimization techniques and a moderate number of recipient and resource types. The priority model was solved to optimality within five hours with 30 resource and recipient types, and the lottery model with 15 resource and recipient types. Our formulations also allow for equity constraints to be included without increasing the overall complexity of the models.

For problems larger than these, or in a setting where a solution must be obtained quickly, we presented some heuristics. The projection heuristic is available in closed form and approximates the optimal solution to the priority model well when the reward coefficients are well explained by the properties. It also approximates the optimal solution to the lottery model even when reward coefficients are poorly explained by the properties, rarely achieving an optimality gap less than 10%. Our lookahead heuristic was more effective at solving the priority model than the projection heuristic when the reward coefficients are poorly explained by the properties, with a maximum optimality gap around 25%.

We also studied the setting where properties of individual recipients and resources vary within their assigned types, and suggested some strategies for choosing a scoring rule. The optimization problem becomes more complex for both priority and lottery models, but we demonstrated that introducing a surrogate objective function leads to tractable problems. On the lottery model, we showed that the techniques had no benefit, however, the techniques significantly reduced discrepancy and improved the reward earned in the priority model.

Our models are attractive for their focus on describing the underlying allocation mechanism. In settings where available data is sparse or sophisticated tools such as simulation models are not available, this approach is essential. On the other hand, it is worth noting that more data-driven approaches to modelling and optimization may work better in settings where granular data and simulation tools are available.

There are several limitations to the models we presented. First, the queueing model that underlies the optimization formulations is extremely simple; it assumes that recipients keep the same type throughout their time in the system. While this may be a reasonable approximation for many settings, others (including organ allocation) are heavily influenced by recipient properties changing over time. Future research could generalize our underlying model to account for transitions between types and provide more modeling flexibility for settings where it is required.

Second, our model does not specifically take into account recipient choice; specifically, the potential for recipients to decline the offer of a resource. This is an extremely influential factor in many allocation systems – particularly those with perishable resources, where it can lead to resource wastage if offers are not made sensibly by the authority. Chapter 3 looks at the effect that recipient choice may have on equity in an allocation system, but another interesting direction for future research could be to study how recipient choice may be incorporated into optimization models for picking the scoring rule (perhaps within the framework of robust optimization).

Chapter 3

Inequity Due to Recipient Choice in Allocation Systems

3.1 Introduction

Almost all real-world allocation systems operated by central authorities have *equity* as one of their goals. Equity is both context dependent, being informed by ethical principles relevant for that particular allocation setting, and highly subjective, generally requiring a trade-off to be made along several competing dimensions. For example, the Organ Procurement and Transplantation Network (OPTN) lists patient utility, justice, and respect for personhood as the three guiding ethical principles for assessing their organ allocation systems, while recognizing that these factors are often at odds with each other [35].

One component of equity is *equality*, which is more concrete in nature and refers to the right held by participants in a system for equal treatment. Equality in an allocation system is sometimes sacrificed as part of a well-understood trade-off; for example, giving the youngest and healthiest patients priority access to organs maximizes utility, but does not treat older and sicker patients as equals [35].

On the other hand, inequality can creep into a system in a way that is not obvious or well-understood; for example, giving priority points to a blood type that is more common in the white population than the non-white population, skewing offers toward

the white population. Clearly, understanding these insidious threats to equality is highly important when designing an allocation system that has equity as a central goal. The aim of this chapter is to model a source of inequality in allocation systems that has not previously been studied: recipient choice.

There is an important distinction to be made between *offers* of resources to recipients, and the actual *allocations* that are produced, since recipients can typically choose to decline the offer of a resource and remain in the system to wait for a future offer. For the rest of this chapter, we fix the setting to be an organ allocation system where a patient may choose to decline the offer of an organ if they believe they would be better served by waiting for a future offer. However, it is easy to see how the same ideas apply to other settings in which scarce resources must be allocated and recipients can choose whether to accept or decline offers.

A patient choosing to decline an organ is exceptionally common in the deceased-donor organ allocation systems overseen by the OPTN. In the kidney program, between July 1, 2015, and June 30, 2016, there were 1 512 496 offers declined by patients and only 11 922 offers accepted [36]. Patient choice is a significant feature of the system, with a pronounced effect on the allocations that are ultimately produced.

It is also important to note that patient choice behavior varies widely across transplant centers, which have the power to decline an organ offered to one of their patients before the patient even sees the offer. This behavior is driven by the methods used for evaluating the performance of transplant centers, which often include metrics such as post-transplant patient survival. Depending on the risk tolerance of the center, it may be prudent for them to decline an offer before it ever reaches the patient [37].

Given that some patients (and their transplant centers) are more selective than others, the question we address in this chapter is the following. Can a disparity in selectiveness when it comes to patient choice lead to inequality in outcomes? If so, how significant can this inequality be?

We seek to answer this question within the framework of some commonly-used approaches to offering organs. One approach is the *waitlist* allocation model, in which each arriving organ is offered uniformly at random among all patients on a

waitlist. Another is the *waiting time* allocation model, in which organ offers are biased towards patients who have been waiting in the system for longer. Both methods seem appealingly equitable in nature, but as we will show, recipient choice can lead to inequity.

3.1.1 Main Results and Contributions

Section 3.2 introduces a general steady-state queueing model under which organs of varying quality (good and bad) are distributed among a population of selective and non-selective patients whose health declines over time. Selective patients are strategic, setting a threshold time before which they choose to decline any offers of bad organs. The model defines a utility function, the concept of a Nash equilibrium, and a measure of inequality between the selective and non-selective patients.

In Sections 3.3, 3.4, and 3.5, we describe instances of the model with specific approaches to making offers: the waitlist, waiting time, and rate allocation models. For the first two models we provide algorithms to compute steady-states of the system, as well as conjectures regarding the existence and uniqueness of equilibria. The third model is introduced as a reference point and shown to be immune to strategic behaviour by selective patients.

Section 3.6 describes and presents the results of numerical studies on the waitlist and waiting time allocation models. We report findings relating to monotonicity of the inequality measure with respect to certain parameters, and attempt to explain the driving forces behind these results.

3.1.2 Related Literature

Very little literature has studied the effects of patient choice on equity in allocation systems. Perhaps the closest related study is by Su and Zenios [38], who note that in the early 2000s, an explosion of patients enrolling on organ transplant waitlists in the US coincided with a continued shortage of organs that left the system resembling a first come, first served (FCFS) queue. Their model studies patients who join a

FCFS M/M/1 queue with exponential reneging and must choose whether to accept or decline an organ when it is offered to them. In contrast to our model, a major assumption is full patient homogeneity, meaning that all patients are strategic. They place only brief focus on studying the inequality that may arise in the system.

A less closely-related study is by Su and Zenios [39], who propose a model in which an arriving stream of organs with different qualities are split between several types of patient queue. Patients, who are assigned a type that is only observed by them, have private information about their propensity for good outcomes after receiving an organ of varying quality. They strategically report a type (not necessarily their own) to join a queue that maximizes their utility at steady state. In contrast to this study, ours attempts to explicitly model patients who choose between accepting and declining the offer of an organ.

There are many other models that study equilibria in queueing systems, but which are not closely related to the setting we study. Hassin and Haviv [40] collect and summarise many of these models.

3.2 Modeling Recipient Choice

This section introduces a general model under which organs of varying quality are offered to patients. Some of these patients are selective, choosing to accept or decline offers after observing the quality of the organ. Other patients are non-selective, accepting all organ offers. Our objective is to study whether inequity may arise from this disparity in strategic behavior. We do not instantiate a specific model in this section; this is left to Sections 3.3, 3.4, and 3.5, which introduce three models with slightly different approaches to offering the organs among patients.

A fundamental assumption that frames the general model is the following: a selective patient declines an organ under the belief that they will obtain a better outcome by waiting. Their strategic behavior therefore trades off the positive utility that would be obtained by accepting an existing offer with the possibility of greater utility from waiting for a better organ, and negative utility from deteriorating health. It is important to note that this assumption encompasses behaviors from both patients and transplant centers in real-life organ allocation systems.

3.2.1 Model Description

The general model is based on two series of m queues running parallel to each other. There is one series for each of two patient types: selective, and non-selective, which are labelled as S and N, respectively. Selective patients arrive into the first selective queue (denoted $S1$) and non-selective patients into the first non-selective queue (denoted $N1$) as a pair of independent Poisson processes with rates λ_S and λ_N, respectively. For notational purposes, we let $\lambda_{\text{tot}} \coloneqq \lambda_S + \lambda_N$ be the total patient arrivals and $\theta_{\text{S}} \coloneqq \lambda_S/(\lambda_S + \lambda_N)$ be the fraction of patient arrivals which are selective. We also let $\mathcal{G} = \{S1, \ldots, Sm, N1, \ldots, Nm\}$ be the set of all $2m$ queues in the system.

The m queues in each series represent different *health stages* of disease. At one end, $i = 1$ is the earliest stage with the healthiest patients, while at the other, $i = m$ is the final stage with patients closest to death. Once a patient arrives into a health stage queue, they progress to the next stage after an exponentially distributed time

with rate parameter $m\gamma$. A patient dies if they progress out of the final health stage. The expected time to progress through each health stage is $1/(m\gamma)$, so the expected time to progress through all stages is $m/(m\gamma) = 1/\gamma$, and therefore $1/\gamma$ represents the average time to death for a patient entering the system.

Each queue in the system has a length, denoted L_q for queue $q \in \mathcal{G}$. We will soon dig into precisely how these lengths are defined, but for the moment it suffices to say that they exist, and that they represent a steady state that the system settles into based on the arrival rates of patients and the service rates of organs consumed by each queue. Note that we also require that each $L_q > 0$.

Two types of organs arrive into the system: good, and bad, which are labelled G and B, respectively. Good organs arrive with total rate μ_G, and bad organs with total rate μ_B, as a pair of independent Poisson processes. Analogous to the patient arrivals, we let $\mu_{\text{tot}} \coloneqq \mu_G + \mu_B$ be the total organ arrivals and $\theta_{\text{G}} \coloneqq \mu_G/(\mu_G + \mu_B)$ be the fraction of total organ arrivals corresponding to good organs.

The main question now is: how are organs consumed by the queues? The answer requires two ingredients to be introduced. The first is the notion of a *threshold* health stage, denoted $y \in \{1, \ldots, m+1\}$. When selective patients are in any stage $i < y$, they decline a bad organ whenever it is offered to them. On the other hand, good organs are always accepted by all patients. Choosing the threshold is the only strategic decision available to the selective patients, and though we will return to it when defining equilibria, for now it should be considered as fixed. For notational purposes, we let $\mathcal{B} = \{Sy, \ldots, Sm, N1, \ldots, Nm\}$ be the set of queues accepting bad organs.

The second ingredient is a collection of scores. One is assigned to each queue, and these are denoted $\{T_q\}_{q \in \mathcal{G}}$. The scores represent the number of *lottery tickets* held by each individual in that queue, where the next arriving organ is offered at random among all patients by selecting one of these lottery tickets. Scores therefore determine the proportion of organs allocated to each queue.

For some fixed scores, the rate at which good organs are consumed by a queue is the product of the total arrival rate of good organs and the fraction of total lottery

tickets held by that queue. With $U_G \coloneqq \sum_{r \in \mathcal{G}} T_r L_r$ representing the total number of lottery tickets in all queues, the rate for queue $q \in \mathcal{G}$ is:

$$\mu_G \left(\frac{T_q L_q}{U_G} \right)$$

On the other hand, bad organs are consumed at a rate of zero by any queue in $\mathcal{G} \setminus \mathcal{B}$. With $U_B \coloneqq \sum_{r \in \mathcal{B}} T_r L_r$ representing the total number of lottery tickets in queues that accept bad organs, the rate that these are consumed by queue $q \in \mathcal{B}$ is:

$$\mu_B \left(\frac{T_q L_q}{U_B} \right)$$

By defining the rates of consumption in this way, the model ensures that any time a bad organ is offered to a queue in $\mathcal{G} \setminus \mathcal{B}$, it is declined and eventually ends up being offered to a queue in $\mathcal{B}$. It is therefore implicitly assumed that an organ never spoils and becomes unsuitable for transplant (which unfortunately does happen in real life [41]).

The scores are what differentiate the three models we study in this chapter. The waitlist model sets scores to $T_q = 1$, the waiting time model sets scores to the cumulative waiting time of each queue, and the rate model uses an approach that deviates slightly from the scoring framework. It is important to note that the scores can depend on properties of the system itself at steady state (such as in the waiting time model), which leaves considerable modeling flexibility to consider extensions that are not studied in this chapter.

It remains to describe how the lengths are calculated in our model. Even when the scores are fixed to constants, the rates at which organs are consumed by each queue depend on the lengths of the queues, which also depend on the consumption rates. We would expect the system to settle into a steady state in which the flows into and out of each queue are balanced. Figure 3-1 illustrates the balances required; we also write them out explicitly in the following paragraphs.

For the selective group, equations only need to balance the good organs received

in all queues before the threshold:

$$\lambda_S = \mu_G \left(\frac{T_{S1} L_{S1}}{U_G} \right) + m\gamma L_{S1}$$

$$m\gamma L_{S(i-1)} = \mu_G \left(\frac{T_{Si} L_{Si}}{U_G} \right) + m\gamma L_{Si} \qquad \forall 2 \leq i < y$$

Whereas after the threshold, equations need to balance good and bad organs received:

$$m\gamma L_{S(i-1)} = \mu_G \left(\frac{T_{Si} L_{Si}}{U_G} \right) + \mu_B \left(\frac{T_{Si} L_{Si}}{U_B} \right) + m\gamma L_{Si} \qquad \forall y \leq i \leq m$$

The non-selective group simply accepts all good and bad organs, so their equations do not change at the threshold:

$$\lambda_N = \mu_G \left(\frac{T_{N1} L_{N1}}{U_G} \right) + \mu_B \left(\frac{T_{N1} L_{N1}}{U_B} \right) + m\gamma L_{N1}$$

$$m\gamma L_{N(i-1)} = \mu_G \left(\frac{T_{Ni} L_{Ni}}{U_G} \right) + \mu_B \left(\frac{T_{Ni} L_{Ni}}{U_B} \right) + m\gamma L_{Ni} \qquad \forall 2 \leq i \leq m$$

Recalling that U_G and U_B are weighted sums involving the queue lengths, there are $2m$ equations and $2m$ unknowns (the queue lengths). Even if the scores are fixed as constants, it is certainly not clear that the system has a solution. We return to this point when defining the waitlist, waiting time, and rate models, and show that there is a unique solution with positive queue lengths in each case.

Before proceeding, it is useful to define some aggregations of the queue lengths that arise in the models. We let the aggregate length of all queues accepting good organs (i.e. all queues) be $L_G \coloneqq \sum_{q \in \mathcal{G}} L_q$, and the aggregate length of all queues accepting bad organs be $L_B \coloneqq \sum_{q \in \mathcal{B}} L_q$. Also, we let $L_S \coloneqq \sum_{i=1}^{m} L_{Si}$ and $L_N \coloneqq \sum_{i=1}^{m} L_{Ni}$ be the aggregate lengths of each patient type, and $L_{SB} \coloneqq \sum_{i=y}^{m} L_{Si}$ be the aggregate length of selective patient queues which accept bad organs.

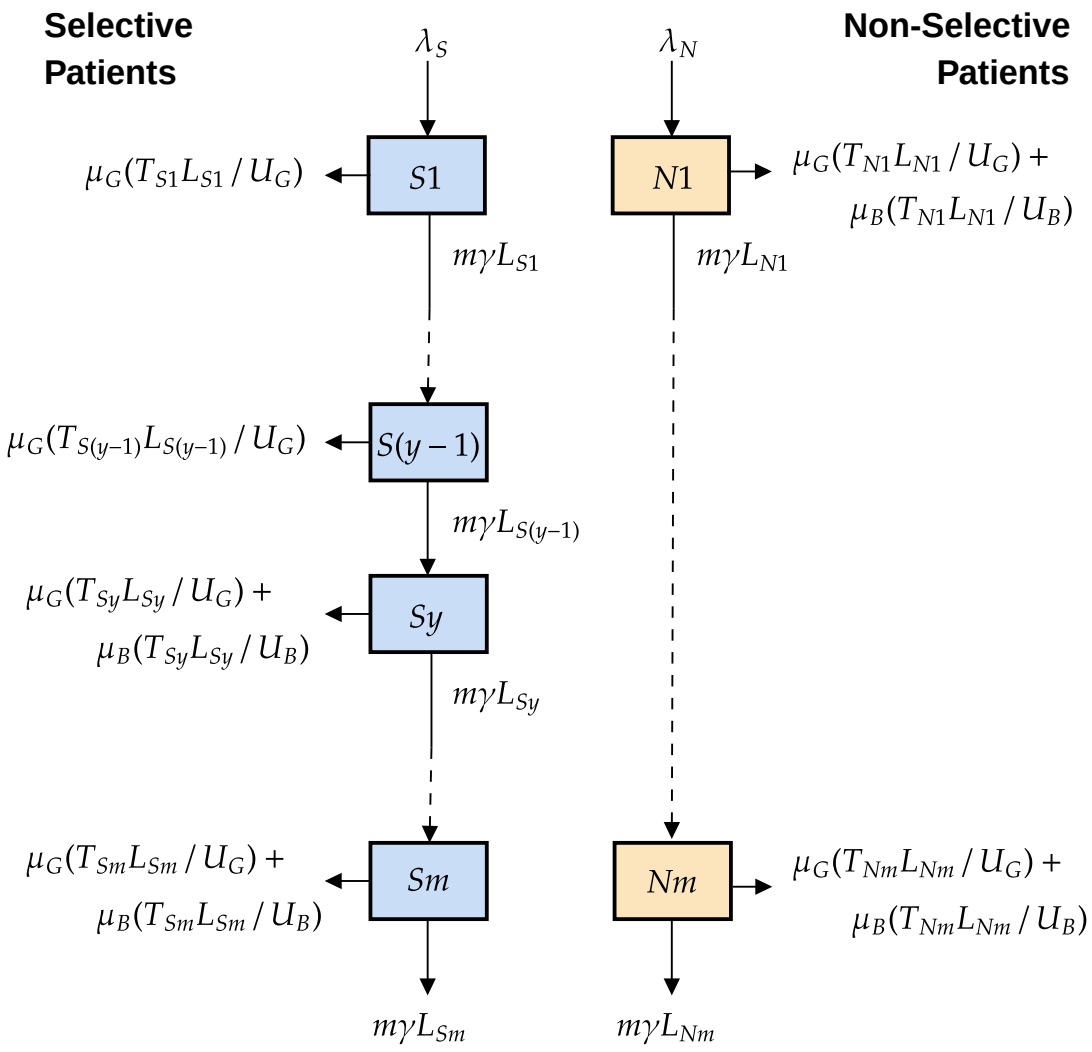

Figure 3-1: Illustration of the queues and flows between them. Balance equations are derived by setting the inward and outward flows for each queue equal to each other.

3.2.2 Outcome Probabilities

Ultimately, the selective group needs a utility function to drive their behavior, and the non-selective group needs a utility function so that inequality can be measured. This subsection defines some quantities of interest: the probabilities of patients observing each of the three possible outcomes (receiving a good organ, a bad organ, and death).

A patient who has arrived into a particular queue is subject to several distinct Poisson processes. If they are accepting bad organs in this queue, they are subject to three processes: one for receiving a good organ, one for a bad organ, and one for moving to the next health stage. On the other hand, if they are declining bad organs, they are only subject to two. These three processes occur with the following rates in each queue, $q \in \mathcal{G}$, where dependence on the threshold y (through the effect y has on the queue lengths) is made explicit:

$$G_q(y) \coloneqq \mu_G \left(\frac{T_q}{U_G}\right) \qquad B_q(y) \coloneqq \mu_B \left(\frac{T_q}{U_B}\right) \qquad D \coloneqq m\gamma$$

Because we aim to define an equilibrium, the probabilities associated with a patient choosing to deviate from population behavior must be quantified. We make the standard assumption that an individual deviating from the population does not alter steady-state queue lengths, allowing the probabilities for an individual choosing threshold y when the population chooses threshold z to be written as follows.

For the selective patients, the probability of receiving a good organ is:

$$p_S^{\text{good}}(y,z) = \sum_{i=1}^{z-1} \frac{G_{Si}(y)}{G_{Si}(y)+D} \left(\prod_{j=1}^{i-1} \frac{D}{G_{Sj}(y)+D}\right) + \left(\prod_{i=1}^{z-1} \frac{D}{G_{Si}(y)+D}\right) \sum_{i=z}^{m} \frac{G_{Si}(y)}{G_{Si}(y)+B_{Si}(y)+D} \left(\prod_{j=z}^{i-1} \frac{D}{G_{Sj}(y)+B_{Sj}(y)+D}\right)$$

where the first term sums the probabilities of receiving a good organ before the individual threshold, z, and the second term sums the probabilities of receiving a good organ after the individual threshold.

Remaining with the selective patients, the probability of receiving a bad organ is:

$$p_S^{\text{bad}}(y,z) = \left(\prod_{i=1}^{z-1} \frac{D}{G_{Si}(y)+D}\right) \times \sum_{i=z}^{m} \frac{B_{Si}(y)}{G_{Si}(y)+B_{Si}(y)+D} \left(\prod_{j=z}^{i-1} \frac{D}{G_{Sj}(y)+B_{Sj}(y)+D}\right)$$

and the probability of death is:

$$p_S^{\text{die}}(y,z) = \left(\prod_{i=1}^{z-1} \frac{D}{G_{Si}(y)+D}\right) \left(\prod_{i=z}^{m} \frac{D}{G_{Si}(y)+B_{Si}(y)+D}\right)$$

On the other hand, because they do not depend on the individual threshold, the non-selective group has expressions for these probabilities which are simpler:

$$p_N^{\text{good}}(y) = \sum_{i=1}^{m} \frac{G_{Ni}(y)}{G_{Ni}(y)+B_{Ni}(y)+D} \left(\prod_{j=1}^{i-1} \frac{D}{G_{Nj}(y)+B_{Nj}(y)+D}\right)$$

$$p_N^{\text{bad}}(y) = \sum_{i=1}^{m} \frac{B_{Ni}(y)}{G_{Ni}(y)+B_{Ni}(y)+D} \left(\prod_{j=1}^{i-1} \frac{D}{G_{Nj}(y)+B_{Nj}(y)+D}\right)$$

$$p_N^{\text{die}}(y) = \prod_{i=1}^{m} \frac{D}{G_{Ni}(y)+B_{Ni}(y)+D}$$

3.2.3 Utility and Equilibrium

Given these expressions for the outcome probabilities, we can introduce a utility function for the selective patients using a parameter, $\eta \in [1/2, 1]$:

$$u_S(y,z) = \eta p_S^{\text{good}}(y,z) + (1-\eta) p_S^{\text{bad}}(y,z)$$

The value η trades off the importance of receiving a good organ against the importance of receiving a bad organ. When $\eta = 1/2$, organs are valued equally, and when $\eta = 1$, a bad organ is considered equivalent to death. For any value of η, death contributes no utility at all.

Within this framework, a population threshold y^* is a Nash equilibrium whenever an individual cannot improve their utility by deviating from y^*. In other words, an equilibrium threshold satisfies the relationship in (3.1):

$$y^* \in \underset{z \in \{1, m+1\}}{\arg\max}\ u_S(y^*, z) \tag{3.1}$$

For each of the three models we study, there turns out to be a unique equilibrium threshold for any value of η. Therefore, given any problem data and some value of η, it is assumed that the system settles into this equilibrium, with utilities for both the selective and non-selective groups given by:

$$u_S^* = \eta p_S^{\text{good}}(y^*, y^*) + (1-\eta) p_S^{\text{bad}}(y^*, y^*)$$

$$u_N^* = \eta p_N^{\text{good}}(y^*) + (1-\eta) p_N^{\text{bad}}(y^*)$$

Throughout the rest of this chapter we consider the ratio u_S^*/u_N^* as the measure of inequality. As this ratio grows, selective patients receive greater utility compared to non-selective individuals. When the ratio is close to one, all patients in the system share similar utility.

3.2.4 Conditions on Data

Before moving on to introduce the specific models, two small restrictions must be imposed on the data. First, because resources are assumed to be scarce, the data must have $\mu_{\text{tot}} \leq \lambda_{\text{tot}}$. If this condition is not satisfied, then supply outpaces demand, all queues effectively have zero length, and the model is not well defined.

Second, we also require that $\mu_B \leq \lambda_N$. This specifically deals with the case $y = m+1$, where the selective patients accept no bad organs, and all of these organs are redirected to the non-selective patients. Without this restriction, when $y = m+1$, all non-selective queues have zero length and the model is not well defined.

3.3 Waitlist Allocation Model

The simplest model for offering organs to queues in the system is the waitlist model. This model is named as such because every patient on the waitlist has the same chance of being offered the next arriving organ. Though it appears on the surface to be an equitable offer scheme, we will ultimately show that inequity may be introduced into this system by recipient choice. In this section, we describe the model, provide an efficient method for calculating its steady-state queue lengths, and also show numerically that it has a unique equilibrium threshold.

3.3.1 Model Definition

In the waitlist model, all scores are set to $T_q = 1$, meaning that each individual patient has an equal chance of being offered the next arriving organ. Queues therefore consume organs in proportion with their lengths. The waitlist model has a particular set of balance equations which are written out for clarity below. For selective patients, the equations are:

$$\lambda_S = \mu_G \left(\frac{L_{S1}}{L_G} \right) + m\gamma L_{S1} \tag{3.2a}$$

$$m\gamma L_{S(i-1)} = \mu_G \left(\frac{L_{Si}}{L_G} \right) + m\gamma L_{Si} \qquad \forall 2 \leq i < y \tag{3.2b}$$

$$m\gamma L_{S(i-1)} = \mu_G \left(\frac{L_{Si}}{L_G} \right) + \mu_B \left(\frac{L_{Si}}{L_B} \right) + m\gamma L_{Si} \qquad \forall y \leq i \leq m \tag{3.2c}$$

whereas for the non-selective patients, they are:

$$\lambda_N = \mu_G \left(\frac{L_{N1}}{L_G} \right) + \mu_B \left(\frac{L_{N1}}{L_B} \right) + m\gamma L_{N1} \tag{3.3a}$$

$$m\gamma L_{N(i-1)} = \mu_G \left(\frac{L_{Ni}}{L_G} \right) + \mu_B \left(\frac{L_{Ni}}{L_B} \right) + m\gamma L_{Ni} \qquad \forall 2 \leq i \leq m \tag{3.3b}$$

As mentioned previously, it is unclear whether these systems have a unique solution, a closed-form solution, or even a solution at all. However, there is a procedure

that exploits the structure of the equations to compute a collection of positive lengths that satisfy the equations. This procedure is based on finding the roots of a function of a single variable, which we call $f(\overline{L}_G)$.

The function $f(\overline{L}_G)$ does not have a simple form, but it does have a simple interpretation, and we now describe how it is computed. The input, $\overline{L}_G > 0$, can be interpreted as a *guess* at a correct value for L_G in the system formed by (3.2) and (3.3). Given this guess, the corresponding guess at a correct value of L_{S1} (called $\overline{L}_{S1}$) can be computed from equation (3.2a):

$$\overline{L}_{S1} = \lambda_S \left(\frac{1}{\mu_G / \overline{L}_G + m\gamma} \right)$$

Next, given $\overline{L}_G$ and $\overline{L}_{S1}$, a guess at the correct value of L_{S2} (called $\overline{L}_{S2}$) can be computed from equation (3.2b):

$$\overline{L}_{S2} = m\gamma \overline{L}_{S1} \left(\frac{1}{\mu_G / \overline{L}_G + m\gamma} \right)$$

and so on, until the guesses $\{\overline{L}_{Si}\}_{i=1}^{y-1}$ representing the lengths of all queues before the threshold y have been computed. A guess at a correct value of L_B (called $\overline{L}_B$) can then be computed as:

$$\overline{L}_B = \overline{L}_G - \sum_{i=1}^{y-1} \overline{L}_{Si}$$

From $\overline{L}_G$ and $\overline{L}_B$, all remaining guesses at lengths can be computed by utilizing the equations in (3.2) and (3.3). The value of $f(\overline{L}_G)$ is the sum of all these guesses:

$$f(\overline{L}_G) = \sum_{q \in \mathcal{G}} \overline{L}_q$$

It is clear that if $f(\overline{L}_G) = \overline{L}_G$, and all $\overline{L}_q > 0$, then by construction we are left with lengths that are positive and satisfy all relationships in (3.2) and (3.3). Conjecture 3.1, which is supported by extensive numerical evidence, would establish that $f(\overline{L}_G)$ has a unique fixed point, and therefore that the system has a unique solution with positive lengths.

Conjecture 3.1. *$g(L) \coloneqq f(L) - L$ is a function with the following properties:*

(i) $g(L)$ is well-defined and finite for all $L > L^*$, where:

$$L^* = \left(\frac{\mu_G}{m\gamma}\right)\left(\frac{(1-\mu_G/\lambda_S)^{1/(y-1)}}{1-(1-\mu_G/\lambda_S)^{1/(y-1)}}\right)$$

(ii) $g(L)$ has exactly one root on the interval (L^*, ∞).

Conjecture 3.1 suggests an algorithm for computing the unique solution to the system in (3.2) and (3.3). First, start with $L_1 = L^* + \delta$ (where $\delta > 0$ is a small perturbation) and L_2 so that $g(L_1) > 0$ and $g(L_2) < 0$. Then, use a bisection search beginning with the interval $[L_1, L_2]$ to find the unique root of g.

3.3.2 Model Behavior and Equilibrium

Figures 3-2 and 3-3 plot the three outcome probabilities for two problem instances as the population threshold varies between 1 and $m + 1$ and individuals choose not to deviate. Most of the behavior observed is intuitive: the selective group can increase their chance of receiving a good organ by delaying the threshold, while reducing their chance of receiving a bad organ.

However, it is also worth noting that the model displays some more complex behavior that depends on the data instance. For example, in Figure 3-2, where fewer patients are selective, increasing the threshold provides diminishing returns for the selective queue on the probability of receiving a good organ. On the other hand, in Figure 3-3, where more patients are selective, increasing the threshold provides increasing returns.

The final questions to answer relate to equilibria. With some value of η defining the utility function, do equilibria exist? Is there a unique equilibrium? Conjecture 3.2 is supported by numerical testing and would answer both questions in the affirmative.

Conjecture 3.2. *For all problem instances, and all values of $\eta \in [1/2, 1]$, the wait-list allocation model has a unique equilibrium threshold satisfying the relationship in equation* (3.1).

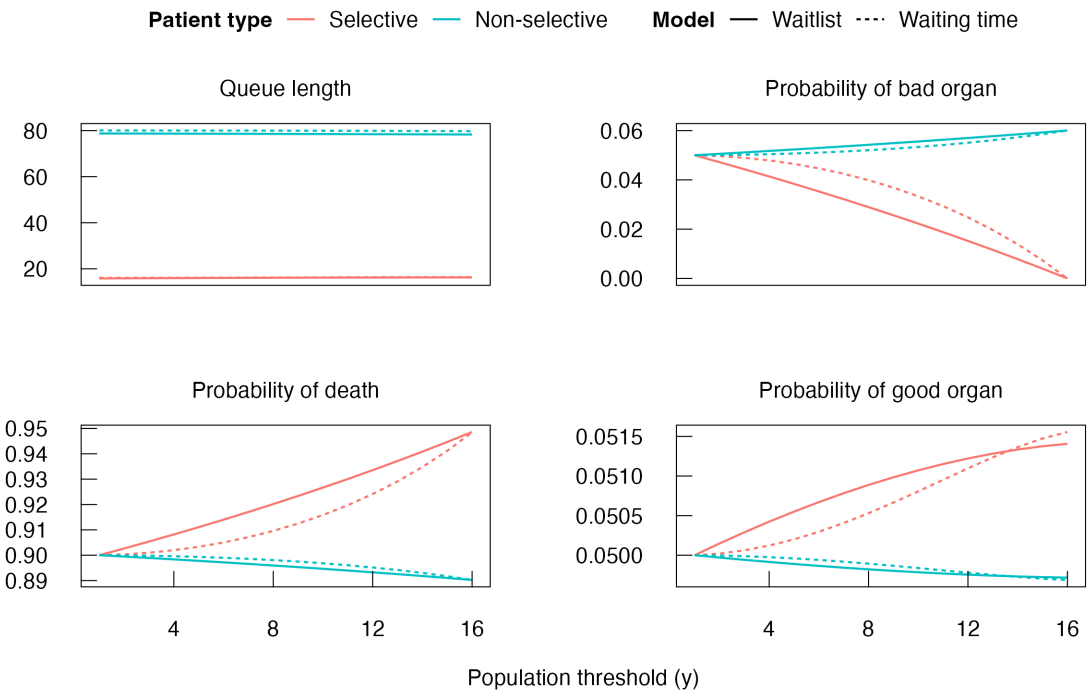

Figure 3-2: Queue length and outcome probabilities for waitlist and waiting time models on problem instance with $\lambda_{\text{tot}} = 100$, $\mu_{\text{tot}} = 10$, $\theta_S = 0.17$, and $\theta_G = 0.5$.

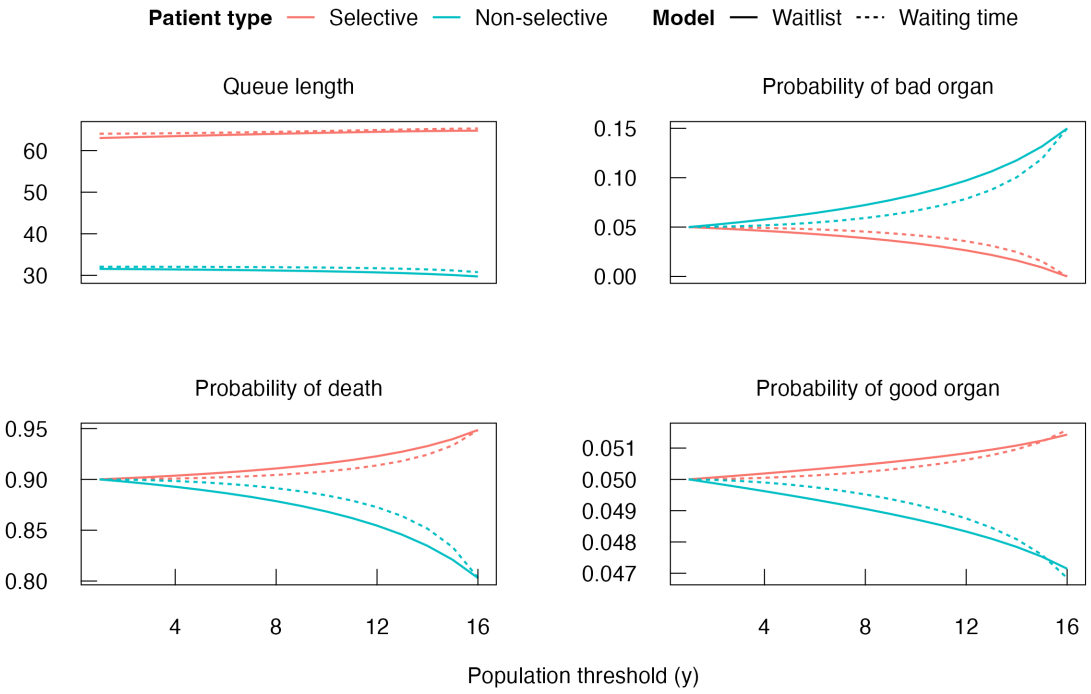

Figure 3-3: Queue length and outcome probabilities for waitlist and waiting time models on problem instance with $\lambda_{\text{tot}} = 100$, $\mu_{\text{tot}} = 10$, $\theta_S = 0.67$, and $\theta_G = 0.5$.

As an example, the details of several problem instances with $\eta = 0.95$ are shown in Figure 3-4. Each panel represents a problem instance, and the two series plotted correspond to the utility of an individual adhering to a population threshold, and the best response for an individual choosing to deviate. The curves intersect at a unique threshold, which is the unique equilibrium.

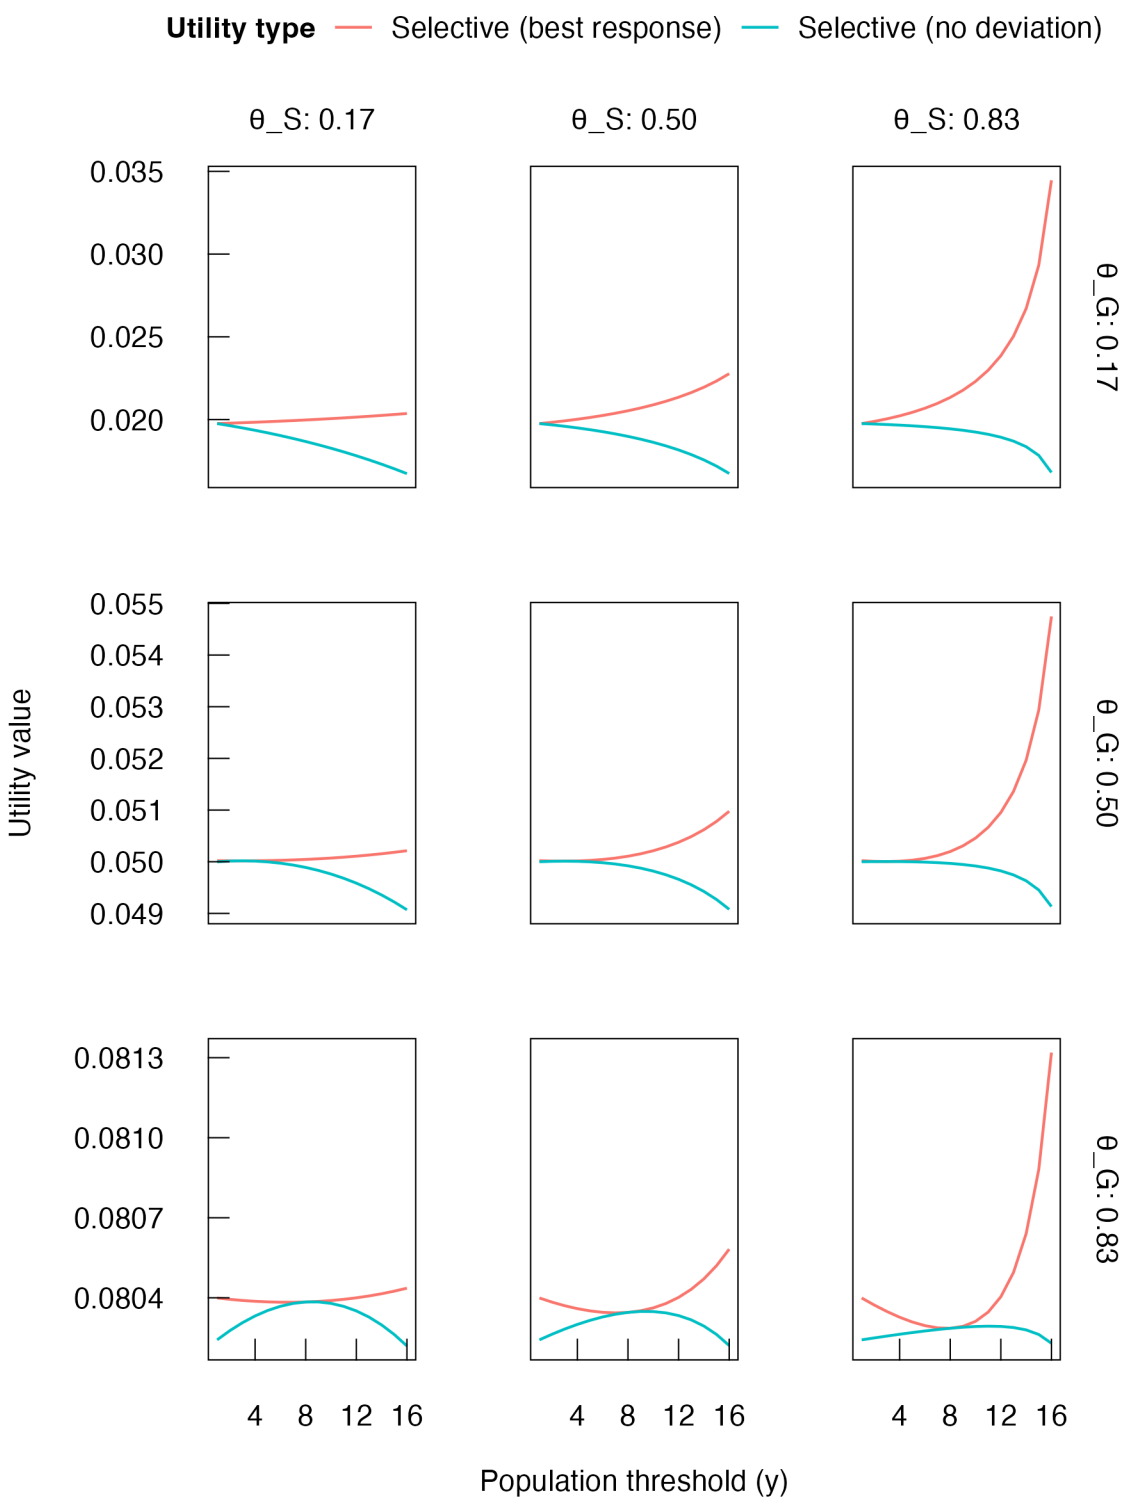

Figure 3-4: Utility for an individual who chooses to deviate (best response) or not from a population threshold in the waitlist model, for various problem instances with $\eta = 0.95$, $\lambda_{\text{tot}} = 100$, and $\mu_{\text{tot}} = 10$.

3.4 Waiting Time Allocation Model

The waiting time model is a natural extension of the waitlist model which appears in many real-world allocation systems. In this model, scores are set in proportion with the waiting time a patient has spent in the system. Waiting time allocation appeals to an intuitive notion of equity, but it will ultimately be shown that inequality can be induced by recipient choice. This section provides a similar exposition to what the waitlist model received in the previous section.

3.4.1 Model Description

Though these quantities have not been explicitly defined previously, each of the queues in our model has a waiting time associated with it by simple application of Little's Law. Letting W_q represent the waiting time in queue $q \in \mathcal{G}$, we have:

$$W_{S1} = \frac{L_{S1}}{\lambda_S} \qquad \text{and} \qquad W_{N1} = \frac{L_{N1}}{\lambda_N}$$

for the first queues in each series, and:

$$W_{Si} = \frac{L_{Si}}{m\gamma L_{S(i-1)}} \qquad \text{and} \qquad W_{Ni} = \frac{L_{Ni}}{m\gamma L_{N(i-1)}}$$

for the remaining queues.

The cumulative time waited by a patient who has arrived in queue L_{Si} can be modelled as $V_{Si} = \sum_{j \leq i} W_{Sj}$ (and similar for the non-selective queues). The waiting time model sets the scores for $q \in \mathcal{G}$ to $T_q = V_q$. This gives rise to the following balance equations.

For the selective patients, they are:

$$\lambda_S = \mu_G \left(\frac{V_{S1} L_{S1}}{\sum_{q \in \mathcal{G}} V_q L_q} \right) + m\gamma L_{S1} \tag{3.4a}$$

$$m\gamma L_{S(i-1)} = \mu_G \left(\frac{V_{Si} L_{Si}}{\sum_{q \in \mathcal{G}} V_q L_q} \right) + m\gamma L_{Si} \qquad \forall 2 \leq i < y \tag{3.4b}$$

$$m\gamma L_{S(i-1)} = \mu_G \left(\frac{V_{Si} L_{Si}}{\sum_{q \in \mathcal{G}} V_q L_q} \right) + \mu_B \left(\frac{V_{Si} L_{Si}}{\sum_{q \in \mathcal{B}} V_q L_q} \right) + m\gamma L_{Si} \qquad \forall y \leq i \leq m \tag{3.4c}$$

Whereas for the non-selective group, they are:

$$\lambda_N = \mu_G \left(\frac{V_{N1} L_{N1}}{\sum_{q \in \mathcal{G}} V_q L_q} \right) + \mu_B \left(\frac{V_{N1} L_{N1}}{\sum_{q \in \mathcal{B}} V_q L_q} \right) + m\gamma L_{N1} \tag{3.5a}$$

$$m\gamma L_{N(i-1)} = \mu_G \left(\frac{V_{Ni} L_{Ni}}{\sum_{q \in \mathcal{G}} V_q L_q} \right) + \mu_B \left(\frac{V_{Ni} L_{Ni}}{\sum_{q \in \mathcal{B}} V_q L_q} \right) + m\gamma L_{Ni} \qquad \forall 2 \leq i \leq m \tag{3.5b}$$

It is once again unclear whether the equations have a unique solution. However, a similar approach to the one used for the waitlist model allows queue lengths that satisfy the equations to be computed with minimal computational burden. The exact procedure is slightly more complex but once again relies on defining a function $f(\overline{U}_G)$, which takes as input a *guess* at the total weight across all queues, $\sum_{q \in \mathcal{G}} V_q L_q$. The value of $f(\overline{U}_G)$ is then computed as follows.

Given $\overline{U}_G$, a guess at L_{S1} (called $\overline{L}_{S1}$) can be computed by substituting the expression for waiting time into (3.4a) and solving the resulting quadratic equation:

$$\lambda_S = \mu_G \left(\frac{V_{S1} \overline{L}_{S1}}{\overline{U}_G} \right) + m\gamma \overline{L}_{S1} = \mu_G \left(\frac{\overline{L}_{S1}^2}{\lambda_S \overline{U}_G} \right) + m\gamma \overline{L}_{S1}$$

which has solutions:

$$\overline{L}_{S1} = \frac{-m\gamma \pm \sqrt{(m\gamma)^2 + (4\mu_G/\overline{U}_G)}}{(2\mu_G)/(\lambda_S \overline{U}_G)}$$

Exactly one of these solutions for $\overline{L}_{S1}$ is positive, and this is the one we select. We are also left with a guess at V_{S1} (called $\overline{V}_{S1}$). Next, given these guesses, we obtain a guess at L_{S2} (called $\overline{L}_{S2}$) by solving the next quadratic balance equation in (3.4b):

$$\begin{aligned} m\gamma \overline{L}_{S1} &= \mu_G \left(\frac{\overline{V}_{S2} \overline{L}_{S2}}{\overline{U}_G} \right) + m\gamma \overline{L}_{S2} \\ &= \mu_G \left(\frac{(\overline{V}_{S1} + \overline{W}_{S2}) \overline{L}_{S2}}{\overline{U}_G} \right) + m\gamma \overline{L}_{S2} \\ &= \mu_G \left(\frac{\overline{V}_{S1} + \overline{L}_{S2}^2/(m\gamma \overline{L}_{S1})}{\overline{U}_G} \right) + m\gamma \overline{L}_{S2} \end{aligned}$$

which again has a single positive solution. This procedure is repeated until reaching the threshold, at which point $\overline{U}_B = \overline{U}_G - \sum_{i=1}^{y-1} V_{Si} L_{Si}$ can be computed. In the same way as for the waitlist model, we build up a full set of guesses at lengths and waiting times, before finally computing the value of $f(\overline{U}_G)$ as:

$$f(\overline{U}_G) = \sum_{q \in \mathcal{G}} \overline{V}_q \overline{L}_q$$

Once again, if $f(\overline{U}_G) = \overline{U}_G$, and all $\overline{L}_q > 0$, then we have lengths that are positive and satisfy all balance equations. Conjecture 3.3 is also supported by extensive numerical evidence and would establish that $f(\overline{U}_G)$ has a unique fixed point.

Conjecture 3.3. *$g(L) \coloneqq f(L) - L$ is a function with the following properties:*

(i) There exists an $L^* > 0$, so that $g(L)$ is well-defined and finite for all $L > L^*$.

(ii) $g(L)$ has exactly one root on the interval (L^*, ∞).

Conjecture 3.3 suggests an algorithm for computing the unique solution to the system in (3.4) and (3.5) that is nearly identical to the one for the waitlist model.

3.4.2 Model Behavior and Equilibrium

Figures 3-2 and 3-3 plot the three outcome probabilities observed in the waiting time model as the population threshold varies between 1 and $m+1$ and individuals choose not to deviate. Observed behavior is similar to the waitlist model, where the selective group can increase their chance of receiving a good organ by delaying the threshold. Figure 3-2, with fewer selective patients, again shows some complex behavior, such as an interval of increasing returns on the probability of receiving a good organ being followed by an interval of diminishing returns as the threshold grows larger.

Just like the waitlist model, the waiting time model appears to have a unique equilibrium threshold for every data instance (as stated in Conjecture 3.4). This claim is backed up by extensive numerical testing and illustrated for a range of data instances in Figure 3-5.

Conjecture 3.4. *For all problem instances, and all values of* $\eta \in [1/2, 1]$*, the waiting time allocation model has a unique equilibrium threshold satisfying the relationship in equation* (3.1)*.*

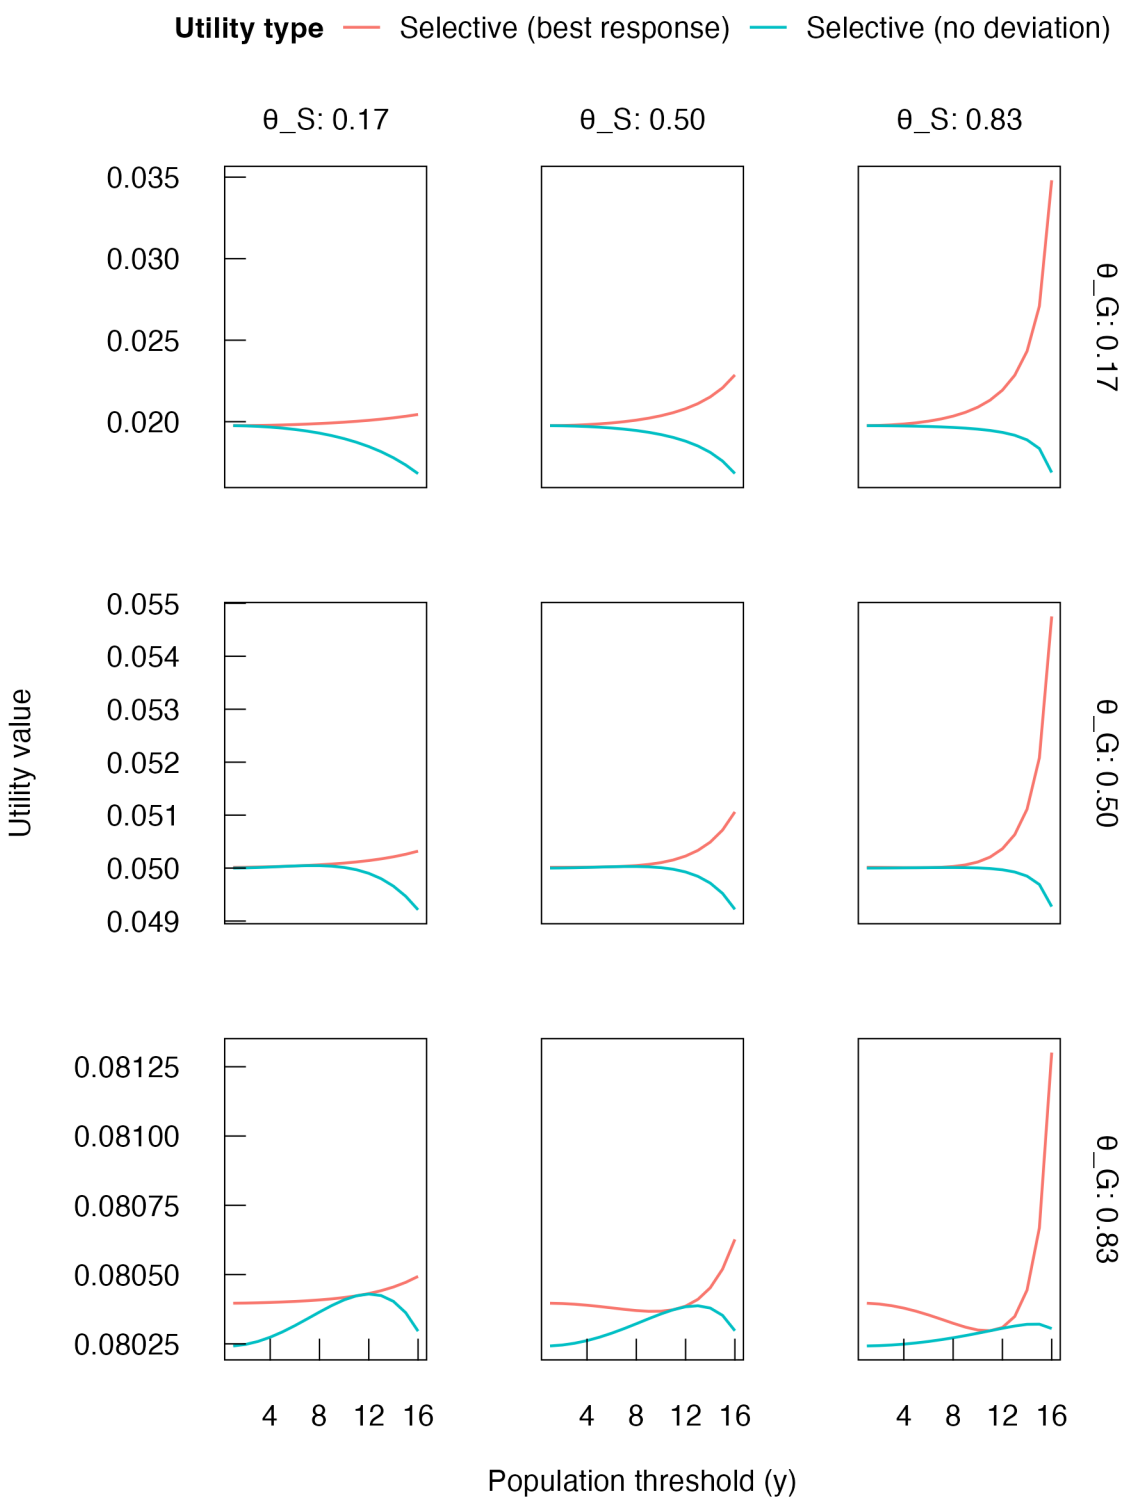

Figure 3-5: Utility for an individual who chooses to deviate (best response) or not from a population threshold in the waiting time model, for various problem instances with $\eta = 0.95$, $\lambda_{\text{tot}} = 100$, and $\mu_{\text{tot}} = 10$.

3.5 Rate Allocation Model

Having introduced the first two models and noted (without showing) that they are susceptible to inequality in the presence of recipient choice, it is reasonable to ask: is there a model in which equality is guaranteed? The model which answers this question in the affirmative is the rate allocation model.

The rate model deviates from the general model we have discussed so far in the sense that it is not defined by assigning scores to queues. However, it has a very intuitive interpretation. First, all arriving organs are split between the two patient types in proportion with the arrival rates. For example, if $\lambda_S = 2\lambda_N$, then two thirds of all arriving organs (both good and bad) are offered to the selective group and one third to the non-selective group. Within the types, allocations are made proportional to the queue lengths.

The corresponding balance equations for the selective patients are:

$$\lambda_S = \theta_S \mu_G \left(\frac{L_{S1}}{L_S} \right) + m\gamma L_{S1} \tag{3.6a}$$

$$m\gamma L_{S(i-1)} = \theta_S \mu_G \left(\frac{L_{Si}}{L_S} \right) + m\gamma L_{Si} \qquad \forall 2 \leq i < y \tag{3.6b}$$

$$m\gamma L_{S(i-1)} = \theta_S \mu_G \left(\frac{L_{Si}}{L_S} \right) + \theta_S \mu_B \left(\frac{L_{Si}}{L_{SB}} \right) + m\gamma L_{Si} \qquad \forall y \leq i \leq m \tag{3.6c}$$

whereas for the non-selective patients they are:

$$\lambda_N = (1-\theta_S) \mu_G \left(\frac{L_{N1}}{L_N} \right) + (1-\theta_S) \mu_B \left(\frac{L_{N1}}{L_N} \right) + m\gamma L_{N1} \tag{3.7a}$$

$$m\gamma L_{N(i-1)} = (1-\theta_S) \mu_G \left(\frac{L_{Ni}}{L_N} \right) + (1-\theta_S) \mu_B \left(\frac{L_{Ni}}{L_N} \right) + m\gamma L_{Ni} \qquad \forall 2 \leq i \leq m \tag{3.7b}$$

These equations have a unique solution of positive queue lengths, and an algorithm for computing them that is almost identical to the ones for the waitlist and waiting time models. We omit details of the algorithm to avoid repetition.

It is more interesting in this case to look at the equilibrium behavior. First, Conjecture 3.5, which is supported by extensive numerical evidence, claims the existence of a unique equilibrium.

Conjecture 3.5. *For all problem instances, and all values of* $\eta \in [1/2, 1]$*, the rate allocation model has a unique equilibrium threshold satisfying the relationship in equation* (3.1).

Next, Proposition 3.1 establishes that if no individual chooses to deviate from the population threshold, then outcomes are perfectly equitable.

Proposition 3.1. *Under the rate allocation model, the following relationships hold:*

$$p_S^{\text{good}}(y, y) = p_N^{\text{good}}(y) \qquad \textit{and} \qquad p_S^{\text{bad}}(y, y) = p_N^{\text{bad}}(y) \qquad \forall y \in \{1, \ldots, m+1\}$$

Proof. As a slight modification of the quantities introduced in Section 3.2.2, the three Poisson process rates which selective patients are subject to are:

$$G_q(y) = \frac{\theta_S \mu_G}{L_S} \qquad B_q(y) = \frac{\theta_S \mu_B}{L_{SB}} \qquad D = m\gamma$$

whereas for non-selective patients these rates are:

$$G_q(y) = \frac{(1-\theta_S)\mu_G}{L_N} \qquad B_q(y) = \frac{(1-\theta_S)\mu_B}{L_N} \qquad D = m\gamma$$

Substituting the relationships for selective patients into the expressions for outcome probabilities and simplifying leads to:

$$p_S^{\text{good}}(y, y) = \frac{\theta_S \mu_G}{\lambda_S} \qquad \text{and} \qquad p_S^{\text{bad}}(y, y) = \frac{\theta_S \mu_B}{\lambda_S}$$

which are simply statements that the outcome probabilities correspond to the overall fractions of arriving patients which leave due to consuming either good or bad organs. Similarly, for the non-selective patients:

$$p_N^{\text{good}}(y) = \frac{(1-\theta_S)\mu_G}{\lambda_N} \qquad \text{and} \qquad p_N^{\text{bad}}(y) = \frac{(1-\theta_S)\mu_B}{\lambda_N}$$

Finally, for any value of η, these expressions can be substituted into the expression for the inequality measure to get:

$$\begin{aligned}\frac{u_S}{u_N} &= \frac{(\theta_S/\lambda_S)\left(\eta\mu_G + (1-\eta)\mu_B\right)}{((1-\theta_S)/\lambda_N)\left(\eta\mu_G + (1-\eta)\mu_B\right)} \\ &= \left(\frac{\lambda_N}{\lambda_S}\right)\left(\frac{\theta_S}{1-\theta_S}\right) \\ &= 1\end{aligned}$$

□

A simple corollary of Proposition 3.1 (assuming Conjecture 3.5 holds) is that the unique equilibrium must satisfy $u_S^* = u_N^*$ and therefore achieve perfect equality in utility.

3.6 Experiments and Results

We performed extensive numerical experiments on both the waitlist and waiting time models across a range of parameters to study whether inequality was introduced by strategic behavior. Problem data were generated using the following procedure, and exact parameter values for the experiments are shown in Table 3.1.

(i) Fix $\gamma = 1$, which sets the scale for the model. Fix $m = 15$.

(ii) Select λ_{tot} as the total arrival rate of patients.

(iii) Select ρ as the *scarcity* parameter, and calculate the total arrival rate of organs as $\mu_{\text{tot}} = \rho\lambda_{\text{tot}}$.

(iv) Select θ_S and θ_G as the fractions of selective patients and good organs, respectively. Calculate $\lambda_S = \theta_S\lambda_{\text{tot}}$ and $\mu_G = \theta_G\mu_{\text{tot}}$.

(v) Select η.

For each combination of parameters, the unique equilibrium threshold was computed, as well as utilities for both patient groups, inequality ratio, and queue lengths. The results show that under certain combinations of parameters, strategic behavior induced inequality in both the waitlist and waiting time models. This section highlights the main insights from the experiments and attempts to explain the mechanisms behind them.

Parameter	**Set**	**Number of Values**
λ_{tot}	$\{100, 10\,000\}$	2
ρ	$\{0.01, 0.1, 0.5\}$	3
λ	$[0.001, 0.999]$	7
μ	$[0.001, 0.999]$	7
η	$[0.5, 1.0]$	12

Table 3.1: Values of parameters used to generate problem instances for studying the waitlist and waiting time models.

Inequality increases with the value of good organs

The inequality ratio was increasing in η in both the waitlist and waiting time models. In other words, as more utility was derived from receiving a good organ compared to a bad organ, inequality between selective and non-selective patients increased. Figure 3-6 shows this effect across multiple parameter combinations with η on the horizontal axis. This experiment also showed the inequality ratio to be greater for the waiting time model than the waitlist model, but only marginally so.

For both models, the effect can be explained by $u_S^*(\eta)$ being convex and $u_N^*(\eta)$ being concave. Appendix B.2 proves that this is a sufficient condition for u_S^*/u_N^* to be increasing. A rough decomposition using the chain rule, where the equilibrium threshold depends on η, provides some insight as to why $u_S^*(\eta)$ is convex. Similar reasoning can be applied to $u_N^*(\eta)$.

$$\left(\frac{\mathrm{d}u_S^*}{\mathrm{d}\eta}\right)_\eta = p_S^{\text{good}}(y^*(\eta)) - p_S^{\text{bad}}(y^*(\eta)) + \left(\frac{\mathrm{d}y^*}{\mathrm{d}\eta}\right)_\eta \left(\eta \left(\frac{\mathrm{d}p_S^{\text{good}}}{\mathrm{d}y}\right)_{y^*(\eta)} + (1-\eta)\left(\frac{\mathrm{d}p_S^{\text{bad}}}{\mathrm{d}y}\right)_{y^*(\eta)}\right)$$

The first term gives the increase in utility due to increasing η: as η increases, weight shifts from p_S^{bad} to p_S^{good}. Since $y^*(\eta)$ is increasing in η, and Figures 3-2 and 3-3 showed that $p_S^{\text{good}}(y,y)$ increases and $p_S^{\text{bad}}(y,y)$ decreases with y, this term is increasing in η.

The second term represents the change in utility due to a changing equilibrium threshold. $\frac{\mathrm{d}y^*}{\mathrm{d}\eta}$ has two regimes: $\frac{\mathrm{d}y^*}{\mathrm{d}\eta} = 0$, when η is small enough that $y^*(\eta) = 1$, and $\frac{\mathrm{d}y^*}{\mathrm{d}\eta} > 0$, when η is large enough that $y^*(\eta) > 1$. The second part of the term represents the change in utility due to an increasing threshold. This whole term is more opaque, but numerical testing shows that the two parts of the term combine to increase in η. Hence, $u_S^*(\eta)$ is convex, and similar testing shows that $u_N^*(\eta)$ is concave.

Furthermore, Figure 3-6 includes plots for $\rho \in \{0.1, 0.5\}$. It is clear that the effect is less pronounced when organs are more scarce.

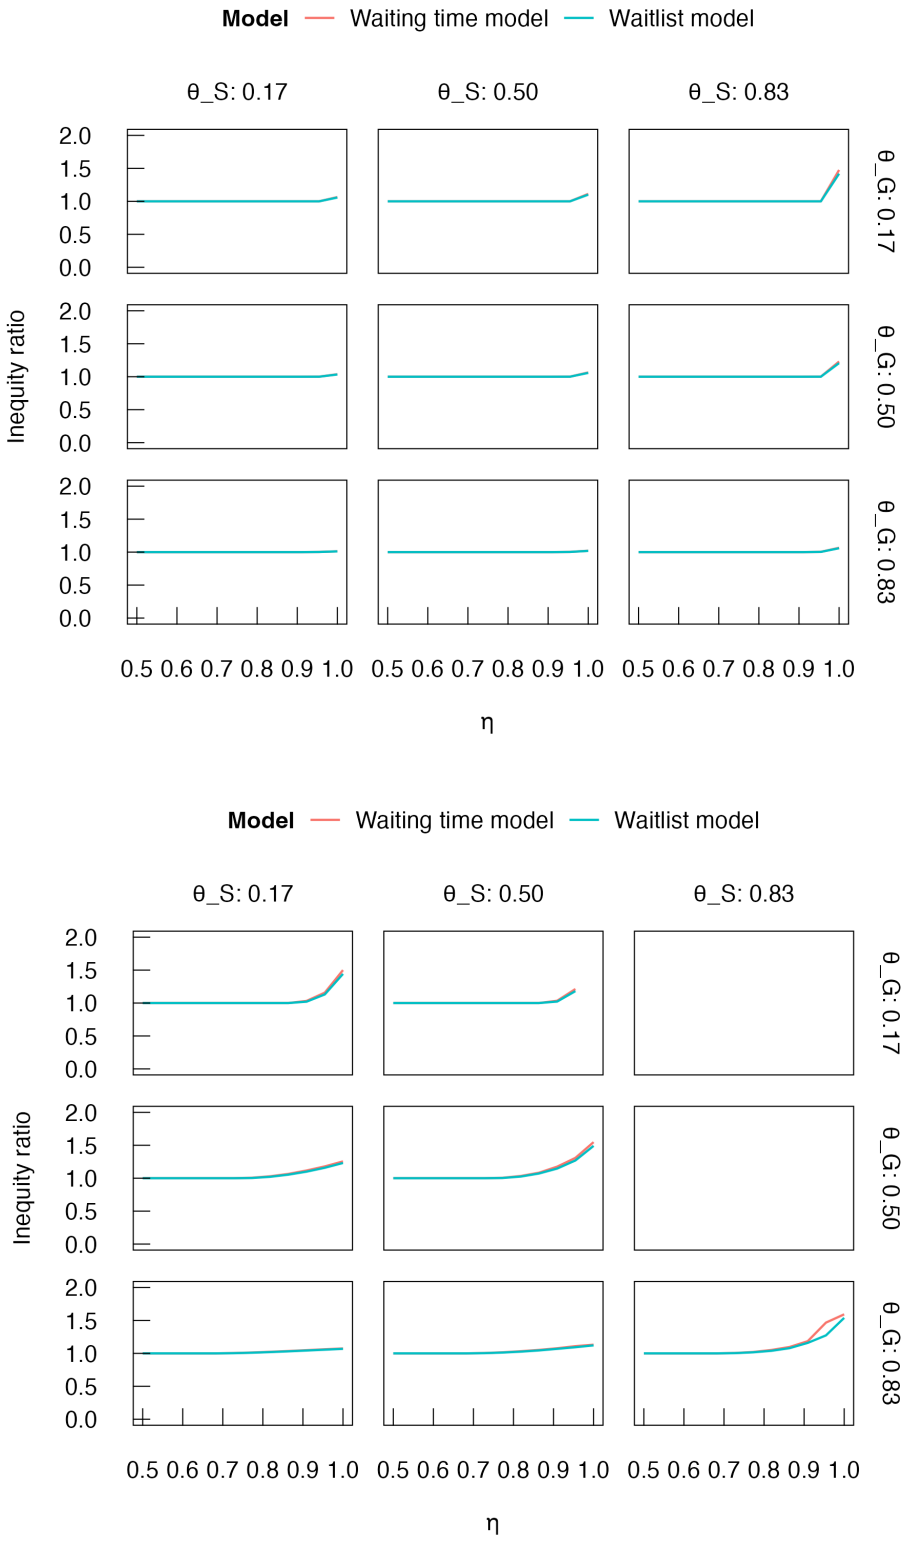

Figure 3-6: Inequality ratio (varying with η) across multiple combinations of θ_S (columns) and θ_G (rows), for $\lambda_{\text{tot}} = 100$, with $\rho = 0.1$ (top) and $\rho = 0.5$ (bottom).

Inequality increases with proportion of selective patients

The inequality ratio was increasing in θ_S in both the waitlist and waiting time models. In other words, as the fraction of patients who are selective increases, inequality between these patients and the non-selective patients increases. Figure 3-7 shows this effect across multiple parameter combinations with θ_S on the horizontal axis.

For both models, the effect is explained by the fact that $\frac{\mathrm{d}u_S^*}{\mathrm{d}\theta_S} \geq \frac{\mathrm{d}u_N^*}{\mathrm{d}\theta_S}$ and $\frac{\mathrm{d}u_N^*}{\mathrm{d}\theta_S} \leq 0$. Appendix B.2 proves that this is a sufficient condition for u_S^*/u_N^* to be increasing.

Again, a rough decomposition using the chain rule sheds light on the mechanism behind this effect. Note that utility in this case depends on the equilibrium threshold, y^*, and also on the value of θ_S itself. In other words, θ_S determines y^*, which itself has an effect on utilities, but for a fixed threshold, θ_S also modifies utilities through its effect on queue lengths. Therefore the rate of change of utility for the selective patients can be written:

$$\frac{\mathrm{d}u_S^*}{\mathrm{d}\theta_S} = \eta \left(\frac{\partial p_G}{\partial \theta_S}\right)_{(\theta_S, y^*(\theta_S))} + \quad (1-\eta)\left(\frac{\partial p_B}{\partial \theta_S}\right)_{(\theta_S, y^*(\theta_S))} + \left(\frac{\mathrm{d}y^*}{\mathrm{d}\theta_S}\right)_{\theta_S} \left(\eta \left(\frac{\partial p_G}{\partial y}\right)_{(\theta_S, y^*(\theta_S))} + (1-\eta)\left(\frac{\partial p_B}{\partial y}\right)_{(\theta_S, y^*(\theta_S))}\right)$$

An interesting observation which holds true in the experiments is the following: as the proportion of selective patients varies, the equilibrium threshold is unaffected. Therefore, the gradient of interest, $\frac{\mathrm{d}u_S^*}{\mathrm{d}\theta_S} - \frac{\mathrm{d}u_N^*}{\mathrm{d}\theta_S}$, reduces to:

$$\frac{\mathrm{d}u_S^*}{\mathrm{d}\theta_S} - \frac{\mathrm{d}u_N^*}{\mathrm{d}\theta_S} = \eta \left(\left(\frac{\partial p_G}{\partial \theta_S}\right)_{(\theta_S, y^*(\theta_S))} - \left(\frac{\partial q_G}{\partial \theta_S}\right)_{(\theta_S, y^*(\theta_S))}\right) + (1-\eta)\left(\left(\frac{\partial p_B}{\partial \theta_S}\right)_{(\theta_S, y^*(\theta_S))} - \left(\frac{\partial q_B}{\partial \theta_S}\right)_{(\theta_S, y^*(\theta_S))}\right)$$

When η is larger, the first term (which is always positive) carries more weight and dominates the second term. When η is smaller, we have y^* closer to 1, and so the first term still dominates the second.

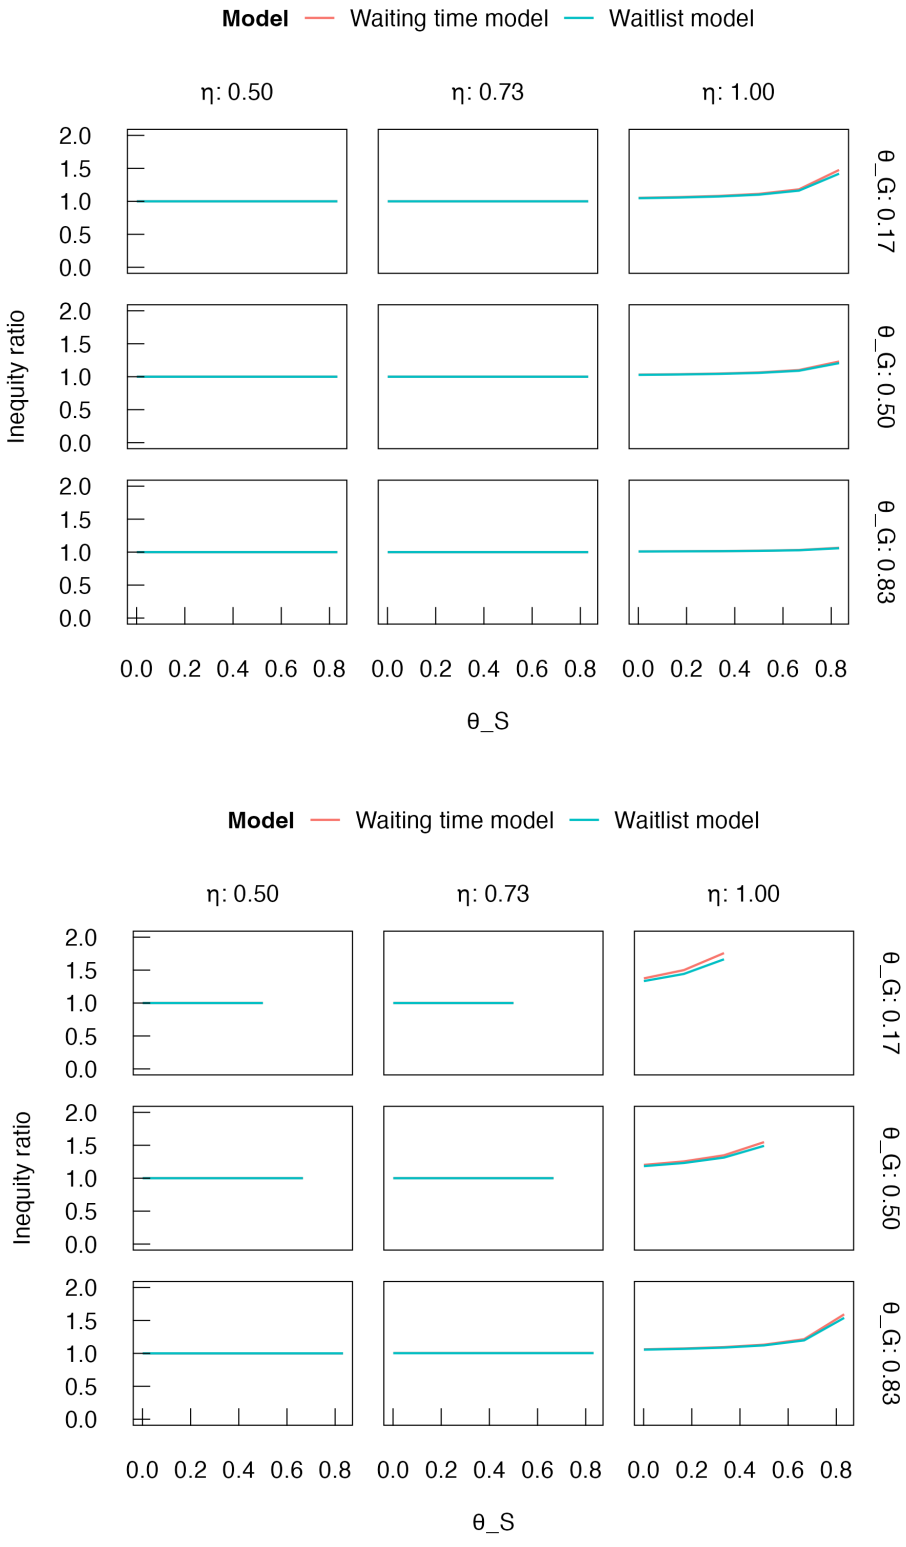

Figure 3-7: Inequality ratio (varying with θ_S) across multiple combinations of η (columns) and θ_G (rows), for $\lambda_{\text{tot}} = 100$, with $\rho = 0.1$ (top) and $\rho = 0.5$ (bottom).

Inequality decreases with proportion of good organs

The inequality ratio was decreasing in θ_G in both the waitlist and waiting time models. As the proportion of arriving organs which were good increased, inequality between selective and non-selective patients decreased, and the effect was more pronounced for larger values of η. Figure 3-8 shows this effect across multiple parameter combinations with θ_G on the horizontal axis.

The result is reasonably intuitive. As θ_G increases, there are more good organs to be distributed within the system. The selective queue accepts more organs, especially with larger values of η, which applies pressure to reduce their aggregate queue length. This ensures that even more of the newly arriving good organs are offered to the non-selective queue.

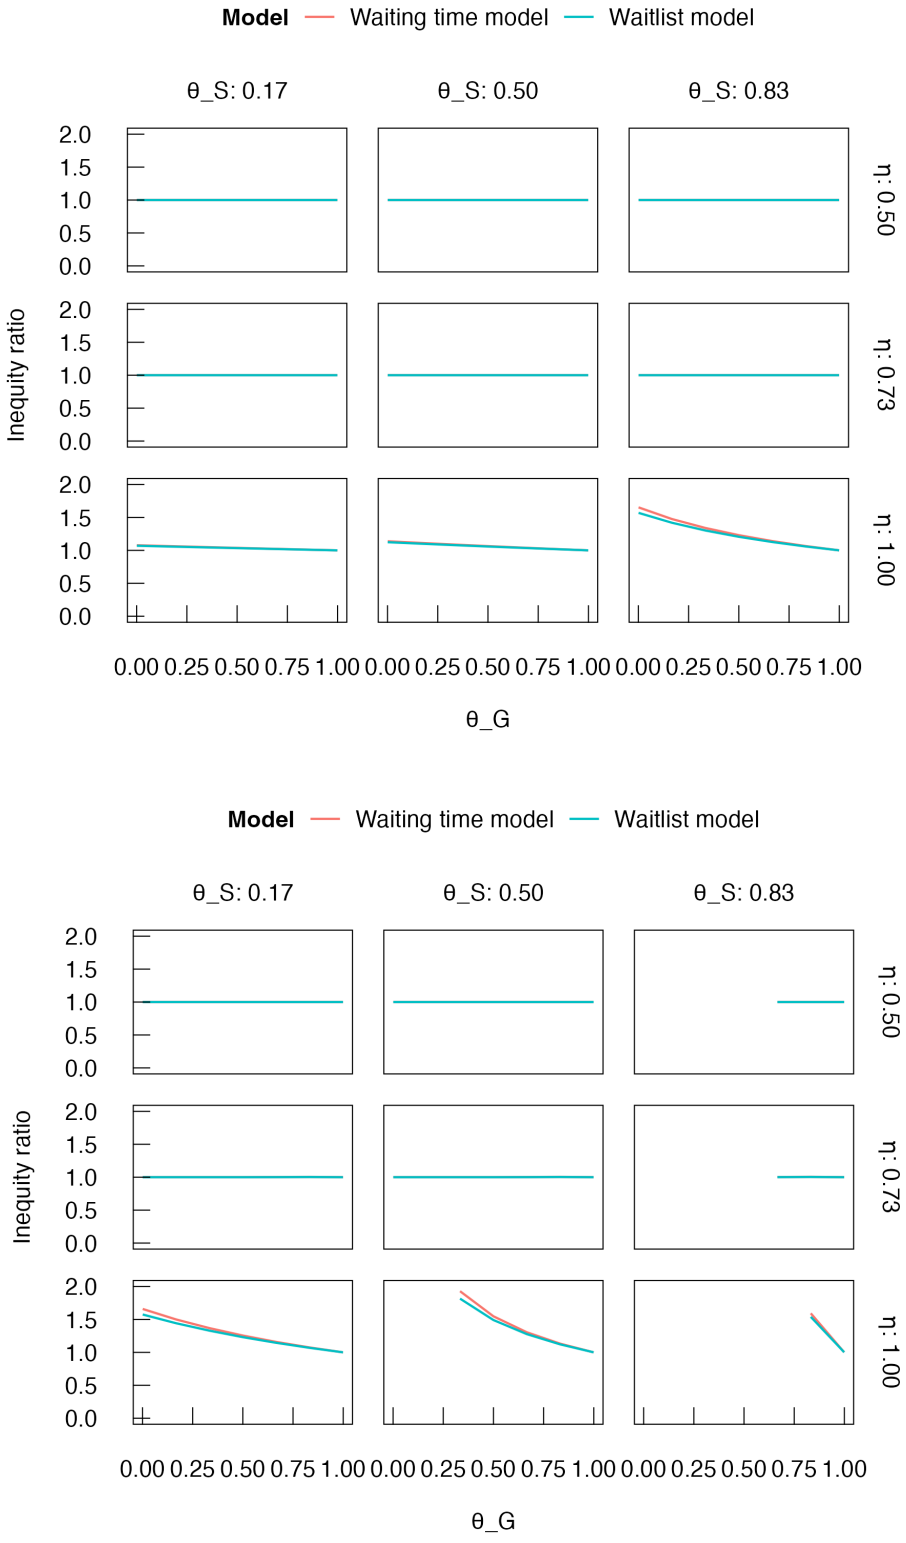

Figure 3-8: Inequality ratio (varying with θ_G) across multiple combinations of θ_S (columns) and η (rows), for $\lambda_{\text{tot}} = 100$, with $\rho = 0.1$ (top) and $\rho = 0.5$ (bottom).

3.7 Discussion and Conclusions

Allocation systems that are controlled by a central authority typically have equity as one of their primary goals. Unfortunately, inequity can creep into a system through subtle mechanisms that are not well-understood by the authority, and it is therefore important to propose models that identify these mechanisms. This chapter studied recipient choice as a novel source of inequity.

We defined a game-theoretic queueing model, with both selective (strategic) and non-selective (non-strategic) recipients, and resources of good and bad quality. Numerical experiments showed that two common (and intuitively equitable) approaches to allocating the resources were susceptible to inequity introduced through strategic play by the selective recipients.

The experiments showed little difference in the inequity introduced by the waitlist and waiting time models. In both cases, inequity increased as the value of a good resource increased relative to the value of a bad resource. This result is intuitive: since selective recipients were able to directly influence their chance of receiving a good resource, it makes sense that utility is concentrated among the selective recipients when this resource is highly-valued. Similarly, in both cases, inequity increased as the fraction of recipients who were selective increased.

The results suggest some specific situations in which the authority who supervises an allocation system with recipient choice must take particular care. In particular, significant inequity may be introduced in a setting where the resources being allocated have widely-varying qualities or valuations, and when a significant fraction of the population is selective. Identifying the former mostly requires domain knowledge, but the latter can be identified using data that are typically collected as part of an allocation system. For example, most deceased-donor organ allocation systems track counts of accepted and declined organs, and it is straightforward to identify the relative selectiveness of different patients and transplant centers using these data.

There are certainly limitations to the study in this chapter. First, several conjectures are left in place of results; these deal with the uniqueness of the solution to

each system of balance equations defining the models, the existence and uniqueness of equilibria in the models, and the monotonicity of the inequality measure as a function of separate variables. Analytical results to prove these conjectures would be an excellent direction for future research.

Second, as has been mentioned, the underlying model in this chapter is a game played on a network of queues with Poisson processes governing all inflows and outflows. Lengths of the queues are defined by the solution to a system of balance equations. This type of modelling framework was not identified in the literature review, and could possibly be generalized to describe other settings in future research.

Finally, we presented a stylized model as a first step in studying the phenomenon of inequity due to patient choice. A data-driven study aimed at determining whether the effect is significant in practice could be an extremely useful next step. This type of study could make use of simulation tools that exist for some systems (such as organ allocation), or be conducted with a model fitted to observed data. We strongly believe that inequity due to recipient choice is an effect that warrants further study.

Chapter 4

Modeling Mass Self-Testing for SARS-CoV-2

4.1 Introduction

The global emergence of the SARS-CoV-2 virus caught the whole world off guard. During 2020 and 2021, widespread community transmission of the pathogen led to massive demand for healthcare resources that stretched many national health systems beyond their capacity. To slow the transmission of the virus and reduce the stress placed on health systems to manageable levels, authorities implemented various non-pharmaceutical interventions (NPIs). These ranged from public advisories emphasizing the importance of minimizing close contacts, to face mask mandates, travel restrictions, and lockdowns [42, 43].

As they developed strategies for intervention, policymakers had to make difficult trade-offs. On the one hand, the most strict interventions (including lockdowns) were very effective at relieving pressure on health systems and averting deaths due to COVID-19 [44, 45]. On the other hand, they came with wide-ranging and potentially huge negative consequences that were both explicit [46] and implicit [47]. These included direct financial costs, disrupted education, and adverse effects on mental health, among many others. Less strict interventions (such as public advisory campaigns) were less effective at curbing transmission of the virus, but came with milder

negative consequences.

In general, predicting the effect of an intervention is a very challenging problem. Some interventions which had been adopted in previous respiratory pandemics similar to COVID-19 had already been studied in the literature, which gave policymakers a good understanding of their positive and negative consequences. However, other interventions were novel, and though they were implemented based on sound reasoning, their effectiveness had not been studied prior to the COVID-19 pandemic. It is easy to see why analyzing these interventions is extremely important for informing responses to similar crises that may occur in the future.

The novel NPI studied in this chapter is mass screening, and in particular, mass self-testing. Mass screening programs involve large, susceptible populations taking tests preemptively and routinely to detect infected individuals before they have the chance to infect many others. Mass screening programs are particularly attractive interventions because they come with limited negative consequences, which are primarily financial costs that can be estimated prior to implementation. During the COVID-19 pandemic, advancements in diagnostics, manufacturing, and supply chain capabilities enabled several countries, including Greece, to procure and distribute self-testing kits to the public at massive scale. However, testing the entire population remained infeasible, and therefore the programs had to choose how to allocate their finite supply of tests among the population.

The Greek self-testing program ran from April to December, 2021. The program initially targeted primary and high-school students and their teaching staff. In May 2021, the program was expanded to private sector employees and civil servants who were working in person and not vaccinated or previously infected. For the most part, each individual in the targeted groups was required by law to take two self tests per week (regardless of whether they were symptomatic) and report the results through a centralized online platform. The entire population was also encouraged to take tests, and free kits were occasionally distributed to everyone, particularly before or after holidays that involved large gatherings of people.

This chapter aims to quantify the effect of the novel Greek self-testing program on

curbing the COVID-19 pandemic. In particular, we aimed to estimate the impact of the program on the reproduction number of the virus, hospitalizations, and deaths, as well as to obtain insights on best practices and lessons learned from the operational decisions made when implementing the program. It is hoped that this analysis can serve as a reference point for policymakers who aim to develop effective intervention strategies during future pandemics.

4.1.1 Main Results and Contributions

The primary methodological contribution of this chapter is a novel epidemiological model that describes the spread of an infectious disease through a population in the presence of self-testing. Crucially, the model makes use of highly granular data that were collected by the Greek NPHO through the online platform used for reporting results of self-tests. It is a compartmental model, and Section 4.2 provides a detailed description of its structure, input data, parameters, and the outcomes it estimates.

Next, the model was fitted to the historical data and used to perform inference, with exact methods described in Section 4.3. Broadly speaking, we aimed to answer the following two questions:

1. What was the overall impact of the program in curbing the COVID-19 pandemic in Greece?

2. What can be learnt from the operational decisions made when implementing the self-testing program in Greece?

Overall Impact of the Program

The overall impact of the self-testing program was assessed according to its effect on three important metrics: reproduction number of the virus, hospitalizations, and deaths. Our analysis found that the program reduced the reproduction number by an average of 4% over the self-testing period, and by as much as 24% during periods when virus transmission was high. Our most conservative estimates also suggest that

discarding the program would have resulted in hospitalizations increasing by approximately 25% and deaths by 20%, corresponding to nearly 22 000 hospitalizations and 2 300 deaths.

Impact of Operational Decisions

Our analysis focused on three key operational decisions made during the implementation of the program. The first related to the scale of the program: what would have happened if the Greek government had increased or decreased the number of tests administered? Our results suggest that if 20% fewer tests had been administered, hospitalizations and deaths would have increased by 10% and 8% respectively. On the other hand: if 20% more tests had been administered, hospitalizations and deaths would have decreased by 8% and 6% respectively. Importantly, this shows a pattern of diminishing returns.

The second question we sought to answer was related to the subpopulations targeted by the distribution of the tests: given a fixed number of tests, how should the Greek government have distributed them among the various age groups in the population? Our results suggest that, after fixing the total number of tests to what was observed, distributing approximately 40% to the 0 to 18 age group, 60% to the 19 to 64 age group, and 0% to the 65-plus age group would have resulted in the largest reduction in hospitalizations and deaths. However, in both metrics, the estimated reduction of 1-2% was small.

The final question related to the quality of the self-tests themselves: how important was the clinical accuracy of the self-tests used? Our results suggest that higher-quality self-tests would have significantly reduced hospitalizations and deaths. For example, if the sensitivity of the self-tests was increased from 60% to 80%, we estimate that hospitalizations and deaths would have reduced by 12% and 10% respectively.

4.2 Model Description

This section provides a detailed description of our model. It begins with a qualitative description of the compartments and transitions in the model (Section 4.2.1), and then elaborates on the dynamics; in other words, how transitions are modeled using a mix of parameters and available historical data (Section 4.2.2). This is followed by a discussion of the methods used to fit the model (Section 4.2.3), while a detailed description of the procedure which minimized the loss function is left to Appendix C.2.

4.2.1 Model Compartments and Transitions

The model is reminiscent of a standard SIR model, but with important modifications that enable relevant dynamics, such as the self-testing and vaccination programs, to be described. All compartments are replicated across three age groups: 0 to 18 years, 19 to 64 years, and 65-plus years. These groups are chosen for two reasons. First, they approximately correspond to the age groups which were subject to different self-testing regimes imposed by the Greek government (details are included in Appendix C.1). Second, they correspond to the age groups which are generally assumed to have significant differences in the health outcomes of infected individuals. Depending on the age group, compartments are indexed with subscript a taking values in $\mathcal{G} = \{0-18,\ 19-64,\ 65+\}$.

Similarly, each compartment is replicated across two vaccination statuses. These are indexed by $v \in \{0, 1\}$ to indicate vaccination. Figure 4-1 illustrates the model compartments and possible transitions between them for a single age group, $a \in \mathcal{G}$.

First, we focus on describing the unvaccinated compartments (indexed by $a0$). Most individuals are susceptible and unvaccinated (S_{a0}) to begin with. From here, they may become vaccinated (S_{a1}) or infected. If infected, they move to one of three compartments: asymptomatic and mild (IAM_{a0}), symptomatic and mild (ISM_{a0}), or symptomatic and severe (ISS_{a0}). Those with mild disease are assumed to recover without requiring treatment, whereas those with severe disease eventually become sick enough to end up in hospital. Populations in all three of these compartments can

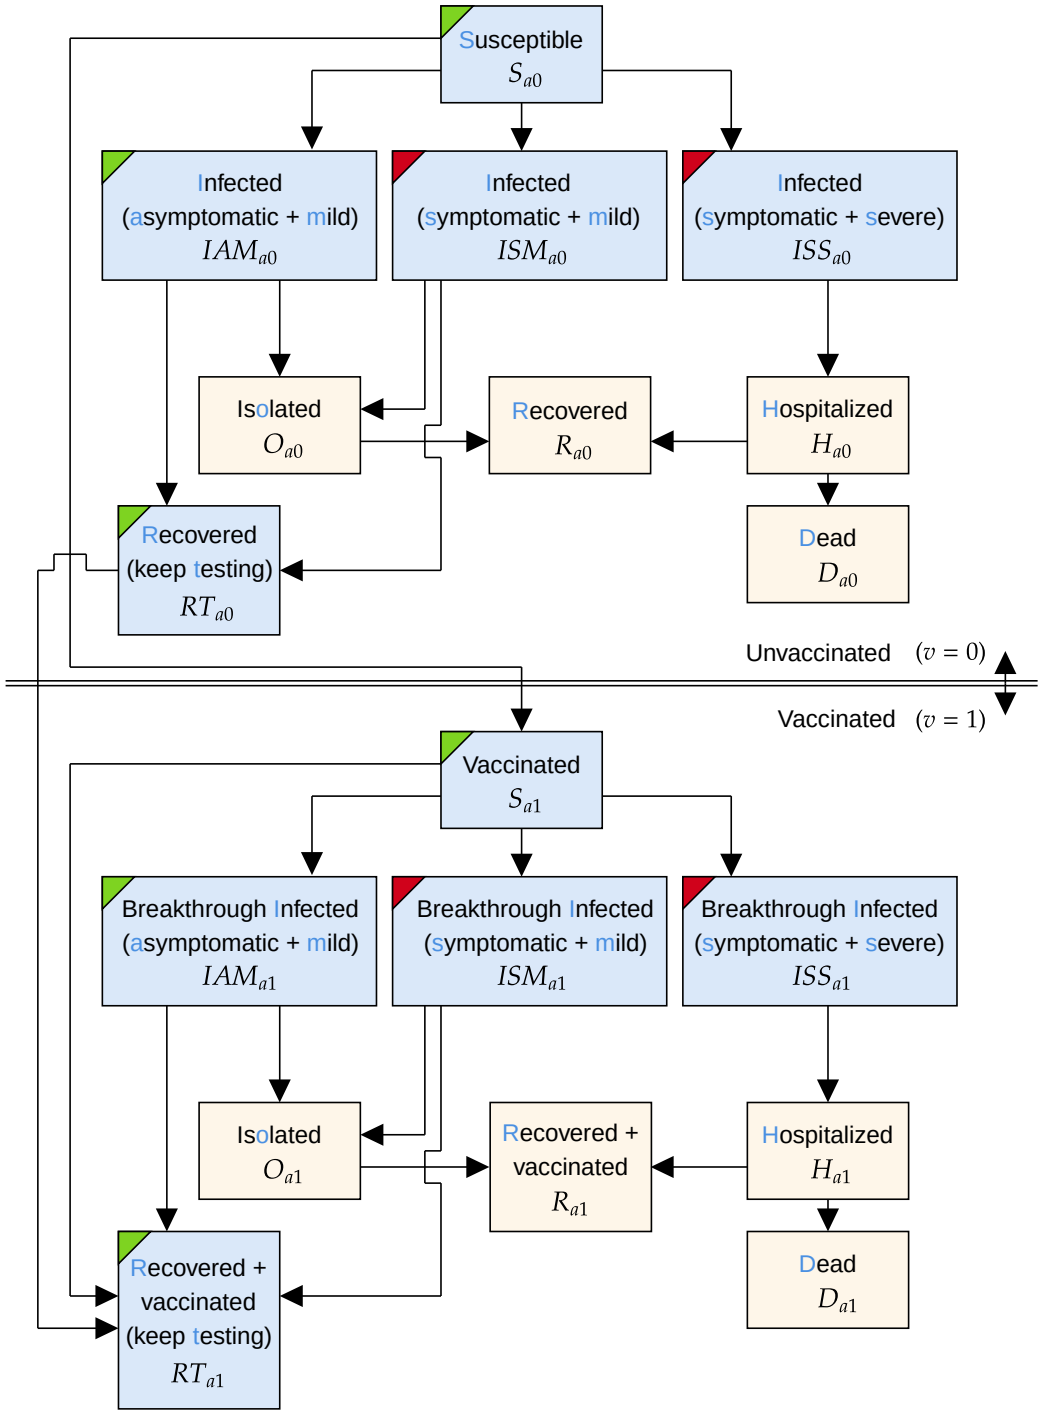

Figure 4-1: The structure of the model within a particular age group, $a \in \mathcal{G}$. For vaccination group $v \in \{0, 1\}$, states are indexed by av. Blue compartments are subject to testing, while green and red corners indicate asymptomatic and symptomatic compartments, respectively.

spread the disease to other susceptible populations in any age group or vaccination status.

Infected individuals may be identified by taking a test, at which point those with mild disease move to an isolated compartment (O_{a0}), whereas those with severe disease move to a hospitalized (H_{a0}) compartment. Infected individuals with mild disease, however, may never be identified through testing – in which case they eventually recover into a compartment that continues to be subject to testing (RT_{a0}). Infected individuals with severe disease are always assumed to eventually be identified.

Individuals in isolation with mild disease are assumed to eventually recover, while hospitalized individuals may recover or die. Recovered individuals transition to R_{a0}, and those who die to D_{a0}.

For the vaccinated population (indexed by $a1$), transitions between compartments follow the same structure as those for the unvaccinated, but with the following exception: a fraction of the vaccinated population is assumed to gain full immunity. Therefore, vaccinated individuals in S_{a1} may transition directly to the recovered compartment that continues to be subject to testing (RT_{a1}) after being exposed to the virus.

4.2.2 Model Dynamics and Data

Guided by data obtained through the Greek NPHO, the model operates in discrete time steps corresponding to days. Time steps are indexed with $t \in \{0, 1, \ldots, T\}$, where T is the model horizon. The actual model was fitted to data spanning January 21, 2021, to December 15, 2021.

The dynamics of the model can be described in terms of the number of individuals that transition between each pair of compartments in a single time step. Let $\mathcal{X}$ be the set of all compartments in the model and $\mathcal{X}_{av}$ be the compartments for a given age group a and vaccination status v. For compartments $X, Y \in \mathcal{X}$, the number of individuals that transition from X to Y at time step t is denoted $\Delta_{X \to Y}(t)$. If compartment X has a population at time step t equal to $X(t)$, the population at the

next time step is given by:

$$X(t+1) = X(t) + \sum_{Y \in \mathcal{X}} \Delta_{Y \to X}(t) - \sum_{Y \in \mathcal{X}} \Delta_{X \to Y}(t)$$

In the model, transitions corresponding to historical data are generally fixed to this data (for example, the number of new vaccinations). Other transitions are modeled using parameters learned as part of the fitting process. To disambiguate, we use the following notational convention: quantities explicitly available in data are denoted with a bar, whereas quantities learned during the fitting process are denoted without a bar. For example: $\bar{x}$ denotes a quantity from the data, whereas x is a quantity that must be learned.

We discuss the transitions in four categories: (1) vaccination, (2) infection, (3) recovery, hospitalization, and death, and (4) testing and isolation.

Vaccination

Vaccinated individuals are defined as those who have received two doses of the vaccine. Let $\bar{v}_a(t)$ denote the recorded number of individuals in age group a who received their second dose on day t. These data are provided by Greek NPHO records [48].

The available data do not specify whether individuals who were vaccinated at time t belonged in the susceptible (S_{a0}) or recovered compartment that is subject to testing (RT_{a0}). We assume a proportional split between the two, so that:

$$\Delta_{S_{a0} \to S_{a1}}(t) = \left(\frac{S_{a0}}{S_{a0} + RT_{a0}} \right) \bar{v}_a(t)$$

$$\Delta_{RT_{a0} \to RT_{a1}}(t) = \left(\frac{RT_{a0}}{S_{a0} + RT_{a0}} \right) \bar{v}_a(t)$$

Infection

Infected individuals expose susceptible individuals to the virus when coming into close contact with them, which may result in the susceptible individuals becoming infected. It is useful to examine two quantities that model these dynamics: first,

the total number of newly infected individuals in each time step, and second, the number of these who transition into the different infected compartments in the model (depending on disease severity and vaccination status).

We first focus on transmission at time t from age group $b \in \mathcal{G}$ to age group $a \in \mathcal{G}$. Consider an infected individual from age group b and, to begin with, assume that the entire age group a population is unvaccinated and susceptible. Then, let $\beta_{ab}(t)$ be the number of individuals from age group a that the infected individual from age group b comes into close contact with and infects. Note that these parameters, which are referred to as *mixing parameters*, capture several relevant factors including the infectivity of the pathogen, the contagiousness of the infected individuals, and the contact patterns between different age groups.

However, each infected individual from age group b need not infect as many as $\beta_{ab}(t)$ individuals from age group a due to two reasons. First, some of those in age group a may not be susceptible, and second, individuals in age group a may have developed full immunity through vaccination.

Let us specifically focus on infections of unvaccinated individuals in age group a. As mentioned above, each infected individual in age group b may infect fewer than $\beta_{ab}(t)$ unvaccinated individuals in age group a, because only a fraction of those they come into contact with are susceptible and unvaccinated. We refer to the population within age group a that an infected individual comes into contact with as the *community population* for that age group. This population is drawn from all compartments except the isolated, hospitalized, and dead compartments. It is denoted $\mathcal{C}_a$:

$$\mathcal{C}_a = \{S_{av},\ IAM_{av},\ ISM_{av},\ ISS_{av},\ RT_{av},\ R_{bv} : v \in \{0,1\}\}$$

and the size of the associated community population is denoted $C_a(t)$:

$$C_a(t) = \sum_{X \in \mathcal{C}_a} X(t)$$

With this notation, the fraction $S_{a0}(t)/C_a(t)$ of all contacts made by an individual

from age group b with age group a are with susceptible and unvaccinated individuals. Therefore, an infected individual from age group b infects $\beta_{ab}(t) \times (S_{a0}(t)/C_a(t))$ unvaccinated individuals from age group a in time step t.

To calculate the total number of unvaccinated individuals from age group a that become infected in a given time step, infections due to each age group must be considered. Let $I_b(t)$ be the number of infected individuals at time t from age group b, given by:

$$I_b(t) = \sum_{v \in \{0,1\}} IAM_{bv}(t) + ISM_{bv}(t) + ISS_{bv}(t)$$

Then, letting $NI_{a0}(t)$ be the total number of newly infected unvaccinated individuals in age group a at time t, the quantity $NI_{a0}(t)$ can be expressed as:

$$NI_{a0}(t) = \sum_{b \in \mathcal{G}} \beta_{ab}(t) \left(\frac{S_{a0}(t)}{C_a(t)} \right) I_b(t)$$

Infections of vaccinated individuals from age group a are similar, with one difference: some vaccinated individuals who are exposed to the virus are not infected due to vaccine-induced immunity. We introduce a parameter termed the *probability of vaccine immunity* to model this effect, denoted $\bar{p}^{\mathrm{v-imm}}$.

More precisely, a vaccinated individual who is exposed to the virus and would have been infected had they not been vaccinated, remains uninfected with probability $\bar{p}^{\mathrm{v-imm}}$. This probability is assumed to remain the same across age groups. Also, note that it relates directly to the effectiveness of the vaccine.

Vaccine effectiveness for the Omicron variant is estimated to be approximately 80% [49]. This value is adjusted to $\bar{p}^{\mathrm{v-imm}} = 85\%$ to account for the greater vaccine effectiveness against non-Omicron variants that were prevalent in Greece throughout 2021 (accounting for vaccine types and variant proportions in the population [50]). Using this new parameter, the number of new infections in the vaccinated from age group a at time step t, denoted $NI_{a1}(t)$, can be written as:

$$NI_{a1}(t) = (1 - \bar{p}^{\mathrm{v-imm}}) \sum_{b \in \mathcal{G}} \beta_{ab}(t) \left(\frac{S_{a1}(t)}{C_a(t)} \right) I_b(t)$$

Next, we describe the transitions that determine the symptoms and severity of disease in an infected individual by introducing two parameters. The first is the *probability of asymptomatic disease*, denoted $\bar{p}_a^{\text{asymp}}$, which represents the probability that an infected individual in age group a does not develop symptoms. We have

$$\bar{p}_{0-18}^{\text{asymp}} = 0.6 \qquad \bar{p}_{19-64}^{\text{asymp}} = 0.35 \qquad \bar{p}_{65+}^{\text{asymp}} = 0.35$$

based on estimates reported for unvaccinated infected individuals [51]. Because there are no data to suggest that the proportion of asymptomatic infection changes due to vaccination, the same probabilities are also used for breakthrough infections.

The second is the *probability of severe disease*, denoted p_{av}^{severe}, which represents the probability that an infected individual who is symptomatic will develop severe disease (as opposed to mild disease). Recall that those with severe disease eventually require hospitalization, whereas those with mild disease recover without treatment. Estimates [52] for the probabilities of infected and unvaccinated individuals who are hospitalized (1% for 0 to 18 years, 7% for 19 to 64 years, and 30% for 65-plus years) do not quite correspond to p_{av}^{severe} in our model, which is a probability conditional on symptomatic disease. Therefore, this probability is left as a free parameter to be learned in the fitting process.

Combining all expressions and parameters, transitions of the susceptible population into the infected compartments can be expressed as follows:

$$\begin{aligned}
\Delta_{S_{av} \to IAM_{av}}(t) &= \bar{p}_a^{\text{asymp}} &&\times && NI_{av}(t) \\
\Delta_{S_{av} \to ISS_{av}}(t) &= (1 - \bar{p}_a^{\text{asymp}}) p_{av}^{\text{severe}} &&\times && NI_{av}(t) \\
\Delta_{S_{av} \to ISM_{av}}(t) &= (1 - \bar{p}_a^{\text{asymp}})(1 - p_{av}^{\text{severe}}) &&\times && NI_{av}(t)
\end{aligned}$$

Finally, we must define the transition from S_{a1} to RT_{a1}. This represents the susceptible and vaccinated individuals in age group a who develop full immunity by taking the vaccine but continue to take tests because they are not themselves aware of their immunity. This transition takes place at the point where such an individual

is exposed to the virus, and therefore can be expressed as:

$$\Delta_{S_{a1}\to RT_{a1}}(t) = \bar{p}^{\text{v-imm}} \sum_{b\in\mathcal{G}} \beta_{ab}(t) \left(\frac{S_{a1}(t)}{C_a(t)}\right) I_b(t)$$

Recovery, Hospitalization, and Death

Infected individuals with mild disease spend an average of $\bar{\tau}^{\text{rec}} = 10$ days [53] in the infected compartments (IAM_{av} or ISM_{av}) before they recover and infectiousness subsides. Therefore, in the absence of a positive test, the following rates for each age group a and vaccination status v can be written:

$$\Delta_{IAM_{av}\to RT_{av}}(t) = \frac{1}{\bar{\tau}^{\text{rec}}} IAM_{av}(t)$$

$$\Delta_{ISM_{av}\to RT_{av}}(t) = \frac{1}{\bar{\tau}^{\text{rec}}} ISM_{av}(t)$$

Similarly, it is assumed that isolated individuals also spend an average of $\bar{\tau}^{\text{rec}}$ days in their compartment before recovery:

$$\Delta_{O_{av}\to R_{av}}(t) = \frac{1}{\bar{\tau}^{\text{rec}}} O_{av}(t)$$

Infected individuals with severe disease are hospitalized an average of $\bar{\tau}^{\text{hosp}} = 10$ days [54] after infection:

$$\Delta_{ISS_{av}\to H_{av}}(t) = \frac{1}{\bar{\tau}^{\text{hosp}}} ISS_{av}(t)$$

Hospitalised individuals are discharged either due to death or recovery. The average length of stay in the hospital, denoted by $\bar{\tau}^{\text{los}}_{av}(t)$, varies with age, vaccination status, and month and is provided by the Greek NPHO [55]. Among the discharged individuals, the fraction p^{die}_a of those who die could in principle be estimated from the raw data. Unfortunately, this estimate was observed to vary substantially over time, and therefore it was left as a parameter to be learned through the fitting process.

Combining the above gives the following transitions:

$$\Delta_{H_{av}\to R_{av}}(t) = (1 - p_a^{\text{die}})\frac{1}{\bar{\tau}_{av}^{\text{los}}(t)}H_{av}(t)$$

$$\Delta_{H_{av}\to D_{av}}(t) = p_a^{\text{die}}\frac{1}{\bar{\tau}_{av}^{\text{los}}(t)}H_{av}(t)$$

Testing and Isolation

The model allows for two types of testing: regular tests and self tests. Both are conducted with the aim of isolating infected individuals and preventing them from spreading the disease, or providing them with the appropriate hospital treatment in the case of severe illness. However, the tests themselves are quite different in nature.

Regular tests are voluntary tests sought out by individuals (usually due to symptoms) and administered by a medical provider. There is little the government can do to pressure an individual into taking a regular test – instead, this process is driven almost entirely by the behavior of the individual.

Regular tests are accurate, with clinical sensitivity and specificity given by $\bar{\sigma}^{\text{reg}} = 0.8$ and $\bar{\mu}^{\text{reg}} = 1$ respectively. Available data [56] provide the total regular tests administered across the country, and these are combined with Greek census data [57] to estimate the number of regular tests performed daily in age group a and vaccination status v, denoted by $\overline{T}_{av}^{\text{reg}}(t)$. Tests are assumed to be allocated proportionally to the sizes of different age and vaccination groups.

At each time step t, the $\overline{T}_{av}^{\text{reg}}(t)$ tests must be split among different compartments. Letting $T_X^{\text{reg}}(t)$ be the tests taken by compartment $X \in \mathcal{X}$, we model the split as follows. First, because newly hospitalized individuals were routinely tested upon admission:

$$T_{ISS_{av}}^{\text{reg}}(t) = \Delta_{X\to H_{av}}(t)$$

The remaining $\overline{T}_{av}^{\text{reg}}(t) - T_{ISS_{av}}^{\text{reg}}(t)$ tests are assumed to be split in the following way. Because regular tests are voluntary, individuals in different compartments had different *propensities* to take them based on their symptoms. θ_a and $1 - \theta_a$ are introduced

to denote the testing propensities of asymptomatic and symptomatic individuals respectively. The total propensity of age group a and vaccination group v, denoted Θ_{av}, can be computed:

$$\Theta_{av} = \theta_a(S_{av}(t) + RT_{av}(t) + IAM_{av}(t)) + (1-\theta_a)ISM_{av}(t)$$

Then, the $\overline{T}^{\text{reg}}_{av}(t) - T^{\text{reg}}_{ISS_{av}}(t)$ tests are split in proportion with the testing propensity and size of each compartment:

$$T^{\text{reg}}_X(t) = \begin{cases} \left(\frac{\theta_a X(t)}{\Theta_{av}}\right)\left(\overline{T}^{\text{reg}}_{av}(t) - T^{\text{reg}}_{ISS_{av}}(t)\right) & \text{if } X \in \{S_{av}, RT_{av}, IAM_{av}\} \\ \left(\frac{(1-\theta_a)X(t)}{\Theta_{av}}\right)\left(\overline{T}^{\text{reg}}_{av}(t) - T^{\text{reg}}_{ISS_{av}}(t)\right) & \text{if } X = ISM_{av} \\ 0 & \text{otherwise} \end{cases}$$

As a brief note, this expression is scale-free in θ_a in the sense that $T^{\text{reg}}_X(t)$ does not change when scaling θ_a and $1-\theta_a$ by the same positive constant. This is why the pair of testing propensities can be normalized to sum to one.

Given $T^{\text{reg}}_X(t)$, the total number of positive regular tests taken by compartment X can therefore be calculated in a straightforward way:

$$P^{\text{reg}}_X(t) = \begin{cases} \bar{\sigma}^{\text{reg}} T^{\text{reg}}_X(t) & \text{if } X \in \{IAM_{av}, ISM_{av}\} \\ \Delta_{X \to H_{av}}(t) & \text{if } X = ISS_{av} \\ 0 & \text{otherwise} \end{cases}$$

In contrast to regular tests, *self-tests* are mandatory. Unvaccinated individuals in certain age groups were required by the government to take these tests at regular intervals regardless of whether they were symptomatic or not, and by increasing or decreasing supply the government was able to increase or decrease testing coverage among these populations.

Self-tests are generally less accurate than regular tests, with clinical sensitivity and specificity given by $\bar{\sigma}^{\text{self}} = 0.6$ and $\bar{\mu}^{\text{self}} = 1$ respectively. Available data [56]

provide the total number of self-tests administered to age group a and vaccination status v each week. These were assumed to be uniformly distributed throughout that week, and we let $\overline{T}^{\text{self}}_{av}(t)$ denote the number of self-tests performed at time t. Because self-tests were only for unvaccinated individuals, note that all $\overline{T}^{\text{self}}_{a1}(t) = 0$.

Because self-tests are mandatory, we assume a uniform distribution among the eligible compartments. Note that no self-tests are taken by individuals in ISS_{av}, because it is assumed that those with severe disease instead seek a regular test. The total number of self-tests administered to compartment X can be calculated as:

$$T^{\text{self}}_X(t) = \begin{cases} \left(\frac{X(t)}{S_{av}(t)+IAM_{av}(t)+ISM_{av}(t)+RT_{av}(t)}\right)\overline{T}^{\text{self}}_{av}(t) & \text{if } X \in \{S_{av}, IAM_{av}, ISM_{av}, RT_{av}\} \\ 0 & \text{otherwise} \end{cases}$$

and the total number of positive self-tests in each compartment is:

$$P^{\text{self}}_X(t) = \begin{cases} \bar{\sigma}^{\text{self}} T^{\text{self}}_X(t) & \text{if } X \in \{IAM_{av}, ISM_{av}, ISS_{av}\} \\ 0 & \text{otherwise} \end{cases}$$

We conclude by noting that all individuals who are identified as positive and are not suffering from a severe infection move to the isolated compartment:

$$\Delta_{X \to O_{av}}(t) = P^{\text{reg}}_X(t) + P^{\text{self}}_X(t) \qquad \text{for } X \in \{IAM_{av}, ISM_{av}\}$$

4.2.3 Model Fitting

The model was fitted against several outcomes for which daily historical data were available: hospitalizations, deaths, total cases, and cases reported through the self-testing program. Parameters were selected to minimize the sum of squared log errors between the model estimates for these outcomes and the data. This subsection provides details on these outcomes, leaving the minimization procedure to Appendix C.2.

Hospitalizations and Deaths

The number of daily hospitalizations (by age group and vaccination status) and deaths (by age group) are available from data recorded by the NPHO of Greece [58]. The corresponding model predictions are given by:

$$\Delta_{ISS_{av} \to H_{av}}(t) \qquad \text{and} \qquad \sum_{v \in \{0,1\}} \Delta_{H_{av} \to D_{av}}(t)$$

Total Cases

The total number of new cases by age group are available from data recorded by the Greek NPHO [59]. The corresponding model predictions are given by:

$$\sum_{v \in \{0,1\}} \sum_{X \in \mathcal{X}_{av}} P_X^{\text{self}}(t) + P_X^{\text{reg}}(t)$$

Cases Reported Through the Self-Testing Program

The total number of new cases (by age group) that were identified through the self-testing program are available from the Greek NPHO [59]. While these cases were predominantly recorded following a positive self test, the data contain artifacts to suggest that some were erroneously recorded after a positive regular test (without an associated positive self test).

To account for this possibility, let $\gamma_a(t)$ be the fraction of regular positive tests that were reported through the self-testing program by age group a at time t. The model estimates of cases reported through the self-testing program are given by:

$$\sum_{v \in \{0,1\}} \sum_{X \in \mathcal{X}_{av}} P_X^{\text{self}}(t) + \gamma_a(t) P_X^{\text{reg}}(t)$$

Note that $\gamma_a(t)$ is assumed to vary with age and over time to account for different types of reporting errors as the system processed varying amounts of self-tests for different segments of the population.

4.3 Methods for Inference

Recall that the study aimed to quantify the overall impact of the self-testing program in Greece and understand the effect of several operational decisions related to the scale of the program, the distribution of self-tests among age groups, and the clinical sensitivity of the self-tests.

This section describes how estimates were obtained by performing sensitivity analyses on the fitted model to simulate outcomes for different input values of either self-tests or their clinical sensitivity. To obtain confidence intervals in addition to point estimates, simulations were performed within a bootstrapping framework whose description is left to Appendix C.2.

4.3.1 Sensitivity Analysis

The sensitivity analyses on the clinical sensitivity of self-tests was performed in a straightforward way. Given the fitted model, the fitted parameters were simply used to simulate outputs of the model for various values of the clinical sensitivity, $\bar{\sigma}^{\text{self}}$.

On the other hand, sensitivity analyses which required modifying the number of self-tests were more complicated. In all cases, self-tests were modified by scaling the input values uniformly across time and different age groups. To be more precise, at time t, the number of self-tests taken by age group a and vaccination status v was set to $(1+\eta)\overline{T}^{\text{self}}_{av}(t)$, where η is the scaling parameter. Note that $\eta = 0$ recovers the observed scenario, while $\eta < 0$ and $\eta > 0$ consider scenarios in which the scale of the program is decreased or increased, respectively.

Changing the scale of the self-testing program would likely have modified voluntary testing behavior compared with what was observed. For example, if self-tests were reduced, some individuals who self-tested in practice and could no longer self-test under this reduced-scale scenario might seek to take a regular test instead (particularly if they were symptomatic). Ignoring this potential for behavioral change could have overstated the effect of self-tests, especially when their scale is reduced ($\eta < 0$).

To address this concern, the following modification was made to each simulation.

First, under the observed scenario and for individuals within a given community compartment $X \in \mathcal{C}_a$, the fraction who sought a regular test among those who did not take a self test was computed:

$$q_X(t) = \min\left(1, \frac{T_X^{\text{reg}}(t)}{X(t) - T_X^{\text{self}}(t)}\right) \tag{4.1}$$

The fraction $q_X(t)$ is the fraction of individuals in X who took a regular test, given that they did not take a self-test. It was assumed that this fraction remained the same regardless of the scale of the self-testing program. Now let $q_X(t)$ represent the fractions that were produced after fitting the model to the data. Then, in the simulations, the number of regular tests taken were modified to be:

$$q_X(t)\left(X(t) - T_X^{\text{self}}(t;\eta)\right) \tag{4.2}$$

where $T_X^{\text{self}}(t;\eta)$ is the number of self-tests in compartment X, after the total number of self-tests are scaled by $(1+\eta)$.

To estimate the impact of the fractions of self-tests that were allocated among different age groups, we varied the fractions of self-tests allocated to each group while constraining them to sum to one (so that to total number of tests is held constant). For each set of fractions considered, the number of tests taken within each age group was scaled uniformly across all time periods.

4.3.2 Conservative Estimates for Overall Impact

The obvious approach to obtaining a simple estimate for the overall impact of the program on reproduction number, deaths, and hospitalizations is to simply set $\eta = -1$. This directly removes self-tests from the model, and we refer to this approach as the *direct method*. It was the only method used to estimate the effect of the program on the effective reproduction number, R_t, which was calculated for each time step according to the method proposed by Arroyo-Marioli et al. [60].

However, during inference, we expect estimates made by the model to become

less accurate as η strays further from zero. This reduced confidence is reflected in wider bootstrapped confidence intervals for estimates of hospitalizations and deaths, and conservative estimates are provided by the lower ends of the intervals. However, it is useful to obtain even more conservative estimates which we can be extremely confident in lower bounding the true values. For this purpose, we made use of the *indirect method* for estimating the overall impact of the program on hospitalizations and deaths, which worked as follows.

First, the sensitivity analyses were used to estimate local derivatives at the observed scenario for deaths and hospitalizations, which we denote with ∂D and ∂H. In particular, the analysis was run for $\eta = -0.01$ and $\eta = 0.01$ to consider $\pm 1\%$ perturbations of the scale of the program, and then finite differences were used to calculate ∂D and ∂H. Given that the total number of deaths and hospitalizations are concave in the number of self-tests conducted (in other words, additional self-tests provide diminishing returns), conservative estimates of the total number of deaths and hospitalizations are provided by:

$$(\text{total number of self tests}) \times \partial D$$

$$(\text{total number of self tests}) \times \partial H$$

4.4 Results

4.4.1 Overall Impact of the Program

The overall impact of the self-testing program was assessed according to its effect on three important metrics: the reproduction number of the virus, hospitalizations, and deaths. As described in Section 4.3, the effect on reproduction number was quantified using standard sensitivity analyses, whereas the effect on hospitalizations and deaths was quantified using standard sensitivity analyses (the direct method) in addition to a more conservative approach (the indirect method).

Table 4.1 presents 80% confidence intervals for the average reduction and the maximum weekly reduction in the effective reproduction number, R_t, over the period of study. The analysis suggests that the program reduced the transmissibility of the virus by 4.7% on average, with a largest weekly reduction of approximately 24%.

Table 4.1 also presents 80% confidence intervals on the number of deaths and hospitalizations averted by the self-testing program, as obtained by the direct and indirect methods of estimation. For reference, the total number of deaths observed in the historical data over the self-testing period was 10 336, and the total number of hospitalizations was 76 299.

Metric	**Estimate**
Percentage reduction in R_t	
Mean	4.72 (3.93 – 5.37)
Maximum	24 (17.3 – 25.6)
Reduction in deaths	
Using direct method	4 888 (3 698 – 6 808)
Using indirect method	3 434 (2 350 – 4 387)
Reduction in hospitalizations	
Using direct method	40 655 (33 625 – 50 699)
Using indirect method	28 379 (21 763 – 34 425)

Table 4.1: Point estimates and 80% confidence intervals for percentage reduction in R_t and absolute reduction in deaths and hospitalizations due to the self-testing program in Greece).

The most conservative estimates on the effect of the self-testing program suggest a mortality reduction of at least 20%, which corresponds to approximately 2 000 deaths. Furthermore, the program yielded a reduction in hospital admissions of at least 25%, corresponding to approximately 20 000 hospitalizations.

4.4.2 Impact of Operational Decisions

We also quantified the effect of three main operational decisions: the scale of the program, the fractions of self-tests allocated among each age group, and the sensitivity of the self-tests themselves. As described in Section 4.3, each effect was estimated using a variation on a standard sensitivity analysis.

Scale of the Program

For policymakers, the trade-offs associated with scale are straightforward: more tests come at higher cost, while increasing the chance of early detection and isolation of infections. Across the period of the study, approximately 60 million self-tests were taken by the population and on average 2.1% of the population were tested daily.

Figure 4-2 shows point estimates and 80% confidence intervals of the percentage changes in hospitalizations and deaths during 2021 had the program in Greece been scaled up or down. For example, if 20% fewer tests had been administered, total hospitalizations would have increased by approximately 10% and total deaths by approximately 8%. If 20% more tests had been administered, total hospitalizations would have decreased by approximately 8% and total deaths by approximately 6%.

Allocation of Self-Tests Among Age Groups

In the program implemented in Greece, 56.2% of the total self-tests were allocated to the 0 to 18 age group, 43.3% were allocated to the 19 to 64 age group, and 0.5% were allocated to the 65-plus age group. Percentage changes in deaths and hospitalizations for all possible alternative allocations of self-tests are provided in Figure 4-3.

The largest reductions are as follows. A strategy that distributed 30% of self-tests

Figure 4-2: 80% confidence intervals for the percentage changes in deaths (left) and hospitalizations (right) as a function of the percentage change in the self-tests administered.

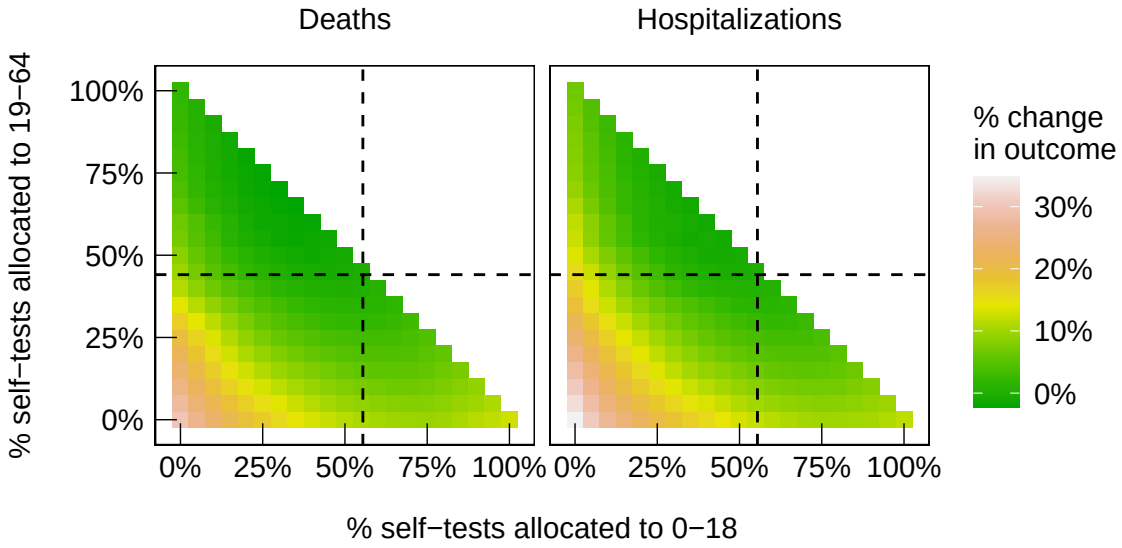

Figure 4-3: Percentage changes in deaths (left) and hospitalizations (right) as the fractions of tests allocated between different age groups changes. Remaining tests are allocated to the 65-plus group, and dashed lines are fractions observed in Greece.

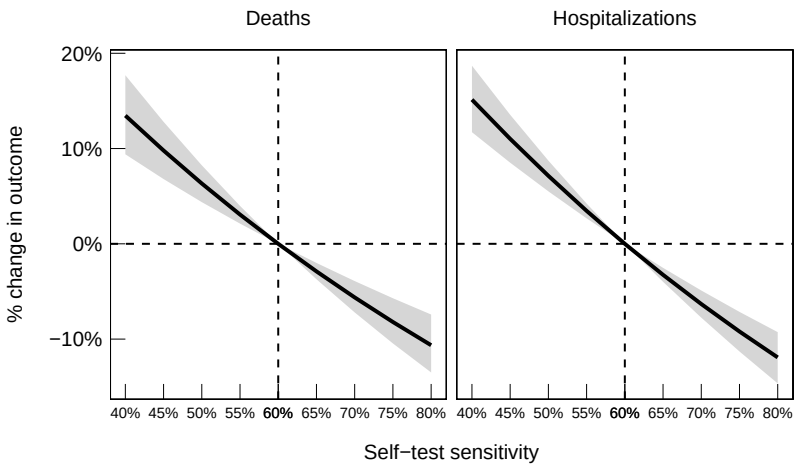

Figure 4-4: 80% confidence intervals for percentage changes in deaths (left) and hospitalizations (right) as self-test sensitivity changes. Dashed line indicates sensitivity reported by the manufacturers of the testing kits used in Greece.

to the 0 to 18 age group and 70% to the 19 to 64 age group resulted in a 2.23% reduction in deaths (80% CI: -3.85 to 6.84%). A strategy that distributed 40% of self-tests to the 0 to 18 age group and 60% to the 19 to 64 age group would have resulted in a 1.16% reduction in hospitalizations (80% CI: -2.34 to 4.04 %).

Accuracy of Tests

Because they must be inexpensive and provide a fast turnaround for results, self-tests are less accurate than the PCR and antigen tests used for regular testing. In order to determine whether the program could have benefitted from more accurate tests, and also to assess the effect of possibly reduced sensitivity due to errors made by the public in administering the tests, the model was used to perform sensitivity analyses for this parameter on hospitalizations and deaths.

Figure 4-4 provides point estimates and 80% confidence intervals of the percentage changes in hospitalizations and deaths during 2021 had the program in Greece used self-tests with different sensitivities. The results indicate that higher quality tests would have contributed to averting more deaths and hospitalizations. For example, if self-test sensitivity was increased from 60% to 80%, we estimate that hospitalizations would have decreased by 12% and deaths by 10%.

4.5 Discussion and Conclusions

NPIs based on social distancing, such as bans on public events, school closures, and full-blown lockdowns, have been shown to reduce transmissibility of a pathogen [45] but are associated with large financial and societal costs. In contrast, mass self-testing programs offer a low-cost solution that has minimal impact on social and economic activity, though the effectiveness of these programs in response to a widespread infectious disease has not been explored previously. Our results suggest that the implementation of the self-testing program in Greece during 2021 was at least as effective in reducing the transmissibility of SARS-CoV-2 as the previously mentioned social distancing measures and averted a significant number of hospitalizations and deaths.

Our results also show the scale of the self-testing program in Greece (as measured by the number of tests) yielded diminishing returns on deaths and hospitalizations avoided. On the one hand, more deaths and hospitalizations would have been averted had the program been scaled up and more tests been administered, but on the other, disproportionally more deaths and hospitalizations would have occurred had the program been scaled down and fewer tests administered. The analysis also shows that alternative allocations of self-tests amongst the age groups could have been more effective in Greece. In particular, increasing the fraction of tests distributed to the 19 to 64 age group could have averted slightly more deaths and hospitalizations, which is consistent with findings that this age group show higher transmissibility compared to children [61]. Finally, the analysis shows that using tests with higher sensitivity would have averted more deaths and hospitalizations.

Given that the analysis was based on observational data and not a (natural or designed) experiment, several assumptions had to be made. First, it was assumed that the dynamics of the disease followed a structure defined by a compartmental model, though structures of this kind are widely used in the mathematical epidemiology literature. Furthermore, when performing the sensitivity analyses, other changes related to population behavior, policy modifications, or additional stress on the healthcare system caused by the perturbation were not accounted for (besides the potential test

substitution effect discussed in Section 4.3).

Moreover, the model assumed that false positive self-tests in susceptible individuals do not lead to them being isolated or hospitalized, since in principle these cases were followed up with a PCR or antigen test and resolved. Another assumption was that individuals die only after being hospitalized, implying there are no deaths at home. The data fully support this latter assumption — there were almost no COVID-19-related deaths outside of a hospital in Greece. Finally, although individuals age, these transitions between compartments of different age groups were not modelled for simplicity.

Despite its shortcomings, the model is rich enough to produce insights for policymakers and public health practitioners on the deployment and optimization of self-testing programs. In late 2020, Slovakia tested their entire population over two consecutive weekends using antigen tests administered by medical professionals [62]. This intervention led to a decrease in the reproduction number (and therefore the number of reported cases) immediately after the testing program [63], but the effect quickly disappeared because high-intensity testing administered by the medical system was not sustainable [64]. In contrast, we found that the Greek program was effective because it focused on testing a large proportion of the population with regularity, even though the tests themselves were of lower quality [65]. This sets an important priority for policymakers responding to future pandemics.

Our findings also suggest that, assuming full population coverage is not feasible due to limited resources or prohibitive cost, policymakers should target the most active subsets of the population with self-tests [66]. Furthermore, tests with higher sensitivity should always be preferred due to their ability to avert more deaths and hospitalizations, but an effective implementation can still be achieved with relatively low sensitivity tests provided the coverage and frequency of testing within the population is high.

Potential for future pandemics seems to be increasing as the global population grows and ages, and as advances in transportation increase global mobility. It is important to document interventions that were used during the COVID-19 pandemic

to develop a set of best practices for future pandemic preparedness. Mass screening through self-testing is one such best practice that could provide an invaluable tool to slow the spread of a highly infectious pathogen with minimal social and financial cost.

Appendix A

Appendix for Chapter 2

A.1 Characterization of Feasible Allocations

Here, we provide a proof of Proposition 2.1 for the case where $J = 2$. Similar reasoning extends the proof easily to $J > 2$.

Proof. First suppose $y = (a, b)$ is a feasible allocation. This means:

$$\mathbf{w}^\intercal f_{a1} > \mathbf{w}^\intercal \bar{\mathbf{f}}_1,\ \forall \bar{\mathbf{f}}_1 \in \mathcal{P}_1 \setminus \{\mathbf{f}_{a1}\} \qquad \text{and} \qquad \mathbf{w}^\intercal f_{b2} > \mathbf{w}^\intercal \bar{\mathbf{f}}_2,\ \forall \bar{\mathbf{f}}_2 \in \mathcal{P}_1 \setminus \{\mathbf{f}_{b2}\}$$

It follows that $\mathbf{w}^\intercal (\mathbf{f}_{a1} + \mathbf{f}_{b2}) > \mathbf{w}^\intercal \left(\bar{\mathbf{f}}_1 + \bar{\mathbf{f}}_2\right), \forall \bar{\mathbf{f}}_1 \in \mathcal{P}_1 \setminus \{\mathbf{f}_{a1}\}, \bar{\mathbf{f}}_2 \in \mathcal{P}_2 \setminus \{\mathbf{f}_{b2}\}$. Therefore, $\mathbf{f}_{a1} + \mathbf{f}_{b2}$ is an extreme point of $\mathcal{P}_1 + \mathcal{P}_2$.

Now suppose that $\mathbf{f}_{a1} + \mathbf{f}_{b2}$ is an extreme point of $\mathcal{P}_1 + \mathcal{P}_2$. By definition there is some $\mathbf{w}$ for which:

$$(\mathbf{f}_{a1} + \mathbf{f}_{b2})^\intercal \mathbf{w} > (\mathbf{g}_1 + \mathbf{g}_2)^\intercal \mathbf{w}$$

for any $\mathbf{g}_1 \in \mathcal{P}_1$ and $\mathbf{g}_2 \in \mathcal{P}_2$ such that $\mathbf{g}_1 + \mathbf{g}_2 \neq \mathbf{f}_{a1} + \mathbf{f}_{b2}$.

If there is some $i \neq a$ for which $\mathbf{f}_{i1}^\intercal \mathbf{w} \geq \mathbf{f}_{a1}^\intercal \mathbf{w}$ (meaning that recipient type i scores at least as highly as a on resource type 1) we would have:

$$(\mathbf{f}_{i1} + \mathbf{f}_{b2})^\intercal \mathbf{w} \geq (\mathbf{f}_{a1} + \mathbf{f}_{b2})^\intercal \mathbf{w}$$

which violates the previous observation given that $\mathbf{f}_{i1} + \mathbf{f}_{b2} \neq \mathbf{f}_{a1} + \mathbf{f}_{b2}$ and both $\mathbf{f}_{i1} \in \mathcal{P}_1$

and $\mathbf{f}_{b2} \in \mathcal{P}_2$.

So $\mathbf{f}_{i1}^\intercal \mathbf{w} < \mathbf{f}_{a1}^\intercal \mathbf{w}$ for any $i \neq a$ and recipient type a is the unique top-scoring recipient type for resource type 1. Similar reasoning applies to b on resource type 2. $\square$

A.2 NP-Hardness of Priority Scoring Optimization

This section provides a proof that the optimization in Problem 2.2 is NP-hard. Notation is preserved from Section 2.2.

We reduce from an instance of the MAX-FLS (feasible linear subsystem) problem in Amaldi and Kann [26]. This problem deals with finding the largest subset of relations that can be satisfied simultaneously in a system of linear inequalities. An instance is defined by a set of J halfspaces given by $H_j = \{\mathbf{w} \in \mathbb{R}^K : \mathbf{a}_j^\intercal \mathbf{w} \geq b_j\}$.

Amaldi and Kann [26] show that this problem remains NP-hard when the inequalities are homogeneous ($b_j = 0$) and the trivial solution $\mathbf{w} = \mathbf{0}$ is excluded. This is the version of the problem we reduce from.

To see the reduction, set $I = 2$, and then $\mathbf{f}_{1j} = \mathbf{a}_j$, $\mathbf{f}_{2j} = \mathbf{0}$, $r_{1j} = 1$ and $r_{2j} = 0$ for all j. If we are able to find some $\mathbf{w} \neq \mathbf{0}$ so that $\mathbf{a}_j^\intercal \mathbf{w} \geq 0$ then we also have $(\mathbf{f}_{1j} - \mathbf{f}_{2j})^\intercal \mathbf{w} \geq \mathbf{0}$ and therefore claim the reward $r_{1j} = 1$ on the jth resource type. The objective values in both problems is clearly identical and the reduction is complete.

A.3 Projection Optimality Bound

Our approach to deriving the bound is to fix $\hat{\mathbf{r}}$, and let an adversary pick a reward vector $\mathbf{r}$ which forms an angle no greater than γ in order to maximize the suboptimality of the projection solution. The adversary is only allowed to perturb the reward vector perpendicular to $\overline{\mathcal{S}}$ in order to ensure that $\hat{\mathbf{r}}$ remains the projected scores.

Note that any vector of reward coefficients can be represented in terms of the perturbation away from $\hat{\mathbf{r}}$ as $\mathbf{r} = \hat{\mathbf{r}} + \boldsymbol{\delta}$, where $\boldsymbol{\delta} \perp \overline{\mathcal{S}}$ is a perturbation orthogonal to $\overline{\mathcal{S}}$. The relationship between γ and $||\boldsymbol{\delta}||$ is:

$$\tan(\gamma) = \frac{||\boldsymbol{\delta}||}{||\hat{\mathbf{r}}||}$$

We wish to define the set $\boldsymbol{\Delta}$ to be all perpendicular perturbations with length less than or equal to some value $\theta := ||\hat{\mathbf{r}}|| \tan(\gamma) > 0$. Let a collection of orthonormal vectors spanning the space orthogonal to $\overline{\mathcal{S}}$ be arranged in columns to get $\mathbf{G} \in \mathbb{R}^{IJ \times (IJ-K)}$. Then $\boldsymbol{\Delta}$ can be defined as a degenerate ball with radius θ in terms of $\mathbf{G}$:

$$\boldsymbol{\Delta} = \{\mathbf{G}\mathbf{u} : ||\mathbf{u}||_2 \leq \theta\}$$

We start by establishing the following lemma:

Proposition A.1. *Let $J = 1$, and the projected reward coefficients be given by $\hat{\mathbf{r}} \in \mathbb{R}^I$ with $\hat{r}_1 \geq \hat{r}_2 \geq \ldots \geq \hat{r}_I$. Let γ be the angle between $\mathbf{r}$ and $\hat{\mathcal{S}}$ as in (2.4).*

Let z^* be the optimal objective value of Problem 2.2, and let z be the objective value of the projection solution. Let $\hat{\mathbf{x}}$ be any set of allocation fractions. The following bound holds for any angle $0 \leq \gamma < \pi/2$:

$$\frac{z^* - z}{||\hat{\mathbf{r}}||} \leq \max\left(\frac{\hat{r}_1}{||\hat{\mathbf{r}}||} + ||\mathbf{e}_1 - \hat{\mathbf{x}}||_2 \tan(\gamma), \ldots, \frac{\hat{r}_I}{||\hat{\mathbf{r}}||} + ||\mathbf{e}_I - \hat{\mathbf{x}}||_2 \tan(\gamma)\right) - \frac{\hat{\mathbf{r}}^\intercal \hat{\mathbf{x}}}{||\hat{\mathbf{r}}||}$$

Proof. We look for the largest suboptimality that can be induced by the adversary across all feasible perturbations, referring to this quantity as α. Since z^* and z depend on the perturbation we write them as $z^*(\delta)$ and $z(\boldsymbol{\delta})$. When the adversary

perturbs, they are restricted by the fact that $\mathbf{r} \geq \mathbf{0}$, which means that by dropping this restriction we get an upper bound on α as $\alpha \leq \max_{\boldsymbol{\delta} \in \boldsymbol{\Delta}} (z^*(\boldsymbol{\delta}) - z(\boldsymbol{\delta}))$.

An upper bound on $z^*(\boldsymbol{\delta})$ is $\max(r_1, \ldots, r_I) = \max(\hat{r}_1 + \delta_1, \ldots, \hat{r}_I + \delta_I)$. Also note that, since the allocation fractions of the projection solution do not change as the adversary perturbs the reward coefficients, we have $z(\boldsymbol{\delta}) = (\hat{\mathbf{r}} + \boldsymbol{\delta})^\intercal \hat{\mathbf{x}}$. We can write:

$$\begin{aligned}
\alpha &\leq \max_{\boldsymbol{\delta} \in \boldsymbol{\Delta}} (z^*(\boldsymbol{\delta}) - (\hat{\mathbf{r}} + \boldsymbol{\delta})^\intercal \hat{\mathbf{x}}) \\
&\leq \max_{\boldsymbol{\delta} \in \boldsymbol{\Delta}} (\max(\hat{r}_1 + \delta_1, \ldots, \hat{r}_I + \delta_I) - (\hat{\mathbf{r}} + \boldsymbol{\delta})^\intercal \hat{\mathbf{x}}) \\
&= \max_{\boldsymbol{\delta} \in \boldsymbol{\Delta}} (\max(\hat{r}_1 + \delta_1 - \boldsymbol{\delta}^\intercal \hat{\mathbf{x}}, \ldots, \hat{r}_I + \delta_I - \boldsymbol{\delta}^\intercal \hat{\mathbf{x}}) - \hat{\mathbf{r}}^\intercal \hat{\mathbf{x}})
\end{aligned}$$

The RHS optimization is nonconvex (as the maximization of a piecewise linear convex function) but we can write it as:

$$\alpha \leq \max \left(\hat{r}_1 + \max_{\boldsymbol{\delta} \in \boldsymbol{\Delta}} (\delta_1 - \boldsymbol{\delta}^\intercal \hat{\mathbf{x}}), \ldots, \hat{r}_I + \max_{\boldsymbol{\delta} \in \boldsymbol{\Delta}} (\delta_I - \boldsymbol{\delta}^\intercal \hat{\mathbf{x}}) \right) - \hat{\mathbf{r}}^\intercal \hat{\mathbf{x}}$$

We let $v_i^* = \max_{\boldsymbol{\delta} \in \Delta} (\delta_i - \boldsymbol{\delta}^\intercal \hat{\mathbf{x}})$. A closed form solution for each of these values can be obtained by optimizing over the ellipsoid $\boldsymbol{\Delta}$ (parameterized by θ for this part of the argument):

$$v_i^* = \theta \sqrt{(\mathbf{e}_i - \hat{\mathbf{x}})^\intercal \mathbf{G} \mathbf{G}^\intercal (\mathbf{e}_i - \hat{\mathbf{x}})}$$

Since the columns of the nonsquare matrix $\mathbf{G}$ are orthonormal, we can extend them into an orthonormal square matrix $\overline{\mathbf{G}}$ to derive a bound on each v_i^*:

$$v_i^* \leq \theta \sqrt{(\mathbf{e}_i - \hat{\mathbf{x}})^\intercal \overline{\mathbf{G}} \overline{\mathbf{G}}^\intercal (\mathbf{e}_i - \hat{\mathbf{x}})} \leq \theta ||\mathbf{e}_i - \hat{\mathbf{x}}||_2$$

and use this to return to the bound on α:

$$\alpha \leq \max (\hat{r}_1 + \theta ||\mathbf{e}_1 - \hat{\mathbf{x}}||_2, \ldots, \hat{r}_I + \theta ||\mathbf{e}_I - \hat{\mathbf{x}}||_2) - \hat{\mathbf{r}}^\intercal \hat{\mathbf{x}}$$

Substituting in the relationship $\tan(\gamma) = \theta / ||\hat{\mathbf{r}}||$ leaves us with the result we

intended to show:

$$\frac{\alpha}{||\hat{\mathbf{r}}||} \leq \max\left(\frac{\hat{r}_1}{||\hat{\mathbf{r}}||} + ||\mathbf{e}_1 - \hat{\mathbf{x}}||_2 \tan(\gamma), \ldots, \frac{\hat{r}_I}{||\hat{\mathbf{r}}||} + ||\mathbf{e}_I - \hat{\mathbf{x}}||_2 \tan(\gamma)\right) - \frac{\hat{\mathbf{r}}^\intercal \hat{\mathbf{x}}}{||\hat{\mathbf{r}}||}$$

□

Next we refine Proposition A.1 to prove Proposition 2.3 for the specific case of the priority allocation.

Proof. In this selection procedure, the allocation is determined by the highest-ranking projected reward coefficient (this is the recipient type with the highest score), so $\hat{\mathbf{x}} = \mathbf{e}_1$. Substituting into the bound obtained in Proposition A.1 gives:

$$\begin{aligned}
\frac{z^* - z}{||\hat{\mathbf{r}}||} &\leq \max\left(\frac{\hat{r}_1}{||\hat{\mathbf{r}}||} + 0, \ldots, \frac{\hat{r}_I}{||\hat{\mathbf{r}}||} + ||\mathbf{e}_I - \mathbf{e}_1||_2 \tan(\gamma)\right) - \frac{\hat{r}_1}{||\hat{\mathbf{r}}||} \\
&= \max\left(0, \frac{\hat{r}_2 - \hat{r}_1}{||\hat{\mathbf{r}}||} + \sqrt{2}\tan(\gamma), \ldots, \frac{\hat{r}_I - \hat{r}_1}{||\hat{\mathbf{r}}||} + \sqrt{2}\tan(\gamma)\right) \\
&= \max\left(0, \frac{\hat{r}_2 - \hat{r}_1}{||\hat{\mathbf{r}}||} + \sqrt{2}\tan(\gamma)\right)
\end{aligned}$$

Finally, by using the relationship $||\hat{\mathbf{r}}|| = ||\mathbf{r}|| \cos(\gamma)$ we get:

$$\frac{z^* - z}{||\mathbf{r}||} \leq \max\left(0, \frac{\hat{r}_2 - \hat{r}_1}{||\hat{\mathbf{r}}||} + \sqrt{2}\sin(\gamma)\right)$$

□

A.4 Lookahead Optimality Bound

The proof in this section establishes the result for $J = 3$, but the same approach can easily be extended to any $J > 3$.

Proof. Let $(i_1, j_1), (i_2, j_2), (i_3, j_3)$ be the sequence of matches selected by the lookahead heuristic, and let $(i'_2, j_2), (i'_3, j_3)$ be the matches appearing in the upper bound used by the heuristic when it selects its first match.

Next, let $(i^*_1, j_1), (i^*_2, j_2), (i^*_3, j_3)$ be the matches in the optimal solution. The algorithm followed by the heuristic ensures that the following two inequalities hold:

1. $r_{i_2 j_2} + r_{i_3 j_3} \geq \max(r_{i'_2 j_2}, r_{i'_3 j_3})$.

2. $r_{i_1 j_1} + r_{i'_2 j_2} + r_{i'_3 j_3} \geq r_{i^*_1 j_1} + r_{i^*_2 j_2} + r_{i^*_3 j_3}$.

If the first inequality did not hold, then the heuristic would have selected (i'_2, j_2) or (i'_3, j_3) before (i_2, j_2). If the second did not hold, then it would have selected (i^*_1, j_1) before (i_1, j_1).

Piecing these observations together gives the ratio:

$$\frac{z}{z^*} = \frac{r_{i_1 j_1} + r_{i_2 j_2} + r_{i_3 j_3}}{r_{i^*_1 j_1} + r_{i^*_2 j_2} + r_{i^*_3 j_3}} \tag{A.1}$$

$$\geq \frac{r_{i_1 j_1} + \max(r_{i'_2 j_2}, r_{i'_3 j_3})}{r_{i_1 j_1} + r_{i'_2 j_2} + r_{i'_3 j_3}} \tag{A.2}$$

This ratio is at its minimum when $r_{i_1 j_1} = 0$, and once this restriction is enforced it becomes clear that the minimum value of the ratio is $1/2$. □

Note that, in the final ratio constructed in the proof, replacing $\max(r_{i'_2 j_2}, r_{i'_3 j_3})$ with $\max(r_{i'_2 j_2}, \ldots, r_{i'_J j_J})$ in the numerator and $r_{i'_2 j_2} + r_{i'_3 j_3}$ with $r_{i'_2 j_2} + \cdots + r_{i'_J j_J}$ in the denominator leads to a minimum ratio of $1/(J-1)$. The remaining logic all holds for $J > 3$, and so does the result.

A.5 Margin Formulation with $\eta > 0$

In the general case when $\eta > 0$, (2.5.3) can be formulated as a linear MIP rather than a linear program. Problem A.1 provides the full formulation:

Problem A.1 (General Margin Formulation)**.**

$$\max_{\mathbf{w},\, \mathbf{s},\, \mathbf{x},\, \boldsymbol{\gamma},\, z} \quad \sum_{i=1}^{I} \sum_{j=1}^{J} r_{ij} x_{ij} + \eta z \tag{A.3a}$$

$$\text{subject to} \qquad \mathbf{e}^\intercal \mathbf{w} = 1 \tag{A.3b}$$

$$s_{ij} = \mathbf{w}^\intercal \mathbf{f}_{ij} \qquad \forall i \in [I], j \in [J] \tag{A.3c}$$

$$s_{ij} - s_{i'j} \geq M(x_{ij} - 1) \qquad \forall i \in [I],\ j \in [J],\ i' \neq i \tag{A.3d}$$

$$\sum_{i=1}^{I} x_{ij} \leq 1 \qquad \forall j \in [J] \tag{A.3e}$$

$$z \leq \sum_{i=1}^{I} x_{ij} s_{ij} - \gamma_j \qquad \forall j \in [J] \tag{A.3f}$$

$$\gamma_j \geq s_{ij} - M x_{ij} \qquad \forall j \in [J] \tag{A.3g}$$

$$\mathbf{x} \in \{0, 1\}^{I \times J} \tag{A.3h}$$

The $\sum_{i=1}^{I} x_{ij} s_{ij}$ term in (A.3f) picks out the top score for resource type j, and the constraints in (A.3g) are only switched on for recipient types which are not top-ranked. The product terms each involve a binary and continuous variable and can be linearized using standard techniques.

A.6 Fixed Point Uniqueness in Lottery Allocation

We numerically verified that Banach-Picard iteration using the operator $T(\mathbf{x})$ (that was defined in Section 2.6) converges to a unique fixed-point. The computation was conducted as follows: for each $I, J \in \{2, 3, 5, 10, 20, 50\}$ (representing the number of recipient and resource types, respectively), 1000 instances of data were randomly generated to satisfy the scarcity assumption. For each instance, 1000 random vectors in $\mathcal{X}$ were selected and used as starting points for the Banach-Picard iteration scheme.

In all instances and starting points, iterations converged to a unique fixed point. Across all data instances, the median number of iterations never exceeded 20, and the median time to converge never exceeded 0.1 seconds.

Appendix B

Appendix for Chapter 3

B.1 Waitlist Model with $m \to \infty$

For the waitlist model, we can write the rates of the three Poisson processes that individual patients are subject to as:

$$G(y) = \frac{\mu_G}{L_G} \qquad B(y) = \frac{\mu_B}{L_B} \qquad D = m\gamma$$

Even though these quantities are the same across all queues, the expressions for the outcome probabilities remain cumbersome. By taking $m \to \infty$ with $y = \bar{y}m$ being held as a fixed fraction $\bar{y} \in [0, 1]$ of m (and similar for z), we end up with simpler expressions for the probabilities and an intuitive interpretation of them. In the notation, we now replace $\bar{y}, \bar{z}$ with y, z, since they still effectively represent thresholds.

For selective patients, the probability of receiving a good organ in this regime is:

$$p_S^{\text{good}}(y, z) = \left(1 - \exp\left(\frac{-G(y)z}{\gamma}\right)\right) + \exp\left(\frac{-G(y)z}{\gamma}\right)\left(\frac{G(y)}{G(y) + B(y)}\right)\left(1 - \exp\left(\frac{-(G(y) + B(y))(1 - z)}{\gamma}\right)\right)$$

which can be interpreted in terms of only two health stages. Under this interpretation, a selective patient arrives into the system, and if they have not received a good organ

before z/γ time has elapsed, they move into the second health stage. Here, they are subject to the processes of both good and bad organ arrivals, and if they do not receive either before $(1-z)/\gamma$ time has elapsed, they die.

The expression for the probability of a selective patient receiving a bad organ is consistent with this interpretation:

$$p_S^{\text{bad}}(y,z) = \exp\left(\frac{-G(y)z}{\gamma}\right)\left(\frac{B(y)}{G(y)+B(y)}\right)\left(1-\exp\left(\frac{-(G(y)+B(y))(1-z)}{\gamma}\right)\right)$$

and so is the expression for the probability of death:

$$p_S^{\text{die}}(y,z) = \exp\left(\frac{-G(y)z}{\gamma}\right)\exp\left(\frac{-(G(y)+B(y))(1-z)}{\gamma}\right)$$

Things are very similar for the non-selective group, except there is only one health stage where the individual is subject to both good and bad organ processes:

$$\begin{aligned}
p_N^{\text{good}}(y) &= \left(\frac{G(y)}{G(y)+B(y)}\right)\left(1-\exp\left(\frac{-(G(y)+B(y))}{\gamma}\right)\right)\\
p_N^{\text{bad}}(y) &= \left(\frac{B(y)}{G(y)+B(y)}\right)\left(1-\exp\left(\frac{-(G(y)+B(y))}{\gamma}\right)\right)\\
p_N^{\text{die}}(y) &= \exp\left(\frac{-(G(y)+B(y))}{\gamma}\right)
\end{aligned}$$

As well as new expressions for the individual probabilities, the $m \to \infty$ regime also gives us an expression that relates L_G and y so that, although we do not have $L_G(y)$ in closed form, we have the relationship implicitly defined in a single equation. In other words, given some y that the entire selective population sticks to, solving the following equation provides us with the unique L_G from which we can determine the lengths of all individual queues using the balance equations presented previously:

$$\begin{aligned}
\mu_G f(y,L_G)+\mu_B = \lambda_S \exp\left(\frac{-G(y)y}{\gamma}\right)\left(1-\exp\left(\frac{-(1-y)}{\gamma}\left(\frac{\mu_G f(y,L_G)+\mu_B}{Lf(y,L_G)}\right)\right)\right)+\\
\lambda_N\left(1-\exp\left(\frac{-1}{\gamma}\left(\frac{\mu_G f(y,L_G)+\mu_B}{L_G f(y,L_G)}\right)\right)\right)
\end{aligned}$$

where:

$$f(y, L_G) = 1 - \frac{\lambda_S}{\mu_G}\left(1 - \exp\left(\frac{-G(y)y}{\gamma}\right)\right)$$

$f(y, L_G)$ can be interpreted as the fraction of good organs which are distributed amongst the queues which also accept bad organs. To see this, note that the expression $\lambda_S\left(1 - \exp\left(\frac{-G(y)y}{\gamma}\right)\right)$ is the rate of selective arrivals receiving a good organ in the first queue of the two-stage model. $\mu_G f(y, L_G) + \mu_B$ is therefore the rate at which organs are offered to the queues which accept bad organs, and the equality effectively balances this rate with the rate at which these organs are received by these queues.

B.2 Supplementary Proofs

Proposition B.1. *Let $g(x) \geq 0$ and $h(x) \geq 0$ be defined on $[0.5, 1]$, and satisfy the following properties:*

(i) $g(0.5) = h(0.5)$.

(ii) $g(x)$ *is convex.*

(iii) $h(x)$ *is concave.*

Then $g(x)/h(x)$ is increasing over its domain.

Proof. Take any $0.5 \leq x_1 \leq x_2 \leq 1$ in the domain. Since $g(x)$ is convex, we have:

$$\frac{g(x_2) - g(0.5)}{x_2 - 0.5} \geq \frac{g(x_1) - g(0.5)}{x_1 - 0.5}$$

which is equivalent to:

$$g(x_2) \geq \left(\frac{g(x_1) - g(0.5)}{x_1 - 0.5} \right) (x_2 - 0.5) + g(0.5)$$

Similarly, we can show that:

$$h(x_2) \leq \left(\frac{h(x_1) - h(0.5)}{x_1 - 0.5} \right) (x_2 - 0.5) + h(0.5)$$

and then, piecing these two relationships together, we end up with:

$$\begin{aligned} \frac{g(x_2)}{h(x_2)} &\geq \frac{\left(\frac{g(x_1) - g(0.5)}{x_1 - 0.5} \right) (x_2 - 0.5) + g(0.5)}{\left(\frac{h(x_1) - h(0.5)}{x_1 - 0.5} \right) (x_2 - 0.5) + h(0.5)} \\ &= \frac{g(x_1) - g(0.5)(x_2 - x_1)}{h(x_1) - h(0.5)(x_2 - x_1)} \\ &\geq \frac{g(x_1)}{h(x_1)} \end{aligned}$$

which shows that $g(x)/h(x)$ is increasing. □

Proposition B.2. *Let $g(x) \geq 0$ and $h(x) \geq 0$ be defined on $[0,1]$, and satisfy the following properties:*

(i) $g(x) \geq h(x)$.

(ii) $g'(x) \geq h'(x)$.

(iii) $h'(x) \leq 0$.

Then $g(x)/h(x)$ is increasing over its domain.

Proof. Let $f(x) \coloneqq g(x)/h(x)$. Then we have:

$$f'(x) = \frac{g'(x)h(x) - g(x)h'(x)}{h(x)^2}$$

At any $x \in [0,1]$, we can examine two possible cases. If $g'(x) \geq 0$, then since $-g(x)h'(x) \geq 0$, the numerator in $f'(x)$ is the sum of two nonnegative terms. Hence, $f'(x) \geq 0$.

On the other hand, if $g'(x) < 0$, then we have:

$$f'(x) \geq \frac{h'(x)(h(x) - g(x))}{h(x)^2} \geq 0$$

which confirms that $f(x)$ is increasing. □

Appendix C

Appendix for Chapter 4

C.1 Additional Program Implementation Details

In March 2021, when a new wave of the COVID-19 pandemic placed significant pressure on the Greek National Health System, authorities chose to launch mandatory weekly testing, starting with students and staff in schools. Self-testing kits were distributed through pharmacies across the country on a weekly basis, using social security numbers. Control was carried out either with the display of a solemn declaration stating that the individual was tested negative, or through an online form available at `self-testing.gov.gr`. For positive self-test results, an amendment of legislation required a second test (either PCR or antigen) by a healthcare professional to confirm the result of the first test and update the National Registry for COVID-19 Patients. The authorities also proceeded occasionally with the distribution of free self-testing kits to the entire population, particularly before or after public holidays.

The distribution of self-testing kits and the implementation of mandatory checks began on April 7, 2021, on the reopening date of high schools, vocational high schools and junior high schools as well as junior high schools with senior classes, with the provision of two self-testing kits per week for all students (as well as the staff at these schools). On April 19, 2021, the government passed a law imposing mandatory self-testing for private sector employees and civil servants working physically at their workplace. Self-testing kits were distributed to these employees throughout the sum-

mer period of 2021, while distribution to teaching staff and students paused during the summer holidays.

In September 2021, authorities imposed once-per-week mandatory self-testing for private sector employees and civil servants who were not fully vaccinated or who have not previously had COVID-19, at personal cost. With the beginning of the new school year, weekly self-testing was mandatory only for students aged 4-18 who were not vaccinated or had not previously had COVID-19. It should be noted that when students were considered as close contacts of a COVID-19 case, the Education Division would provide additional self-tests to ensure daily testing. Finally, private sector employees and civil servants who were not vaccinated or had not had COVID-19 were obliged to carry out two weekly tests from November 5, 2021, onwards, at personal cost, considering the role of seasonality in the spread of SARS-CoV-2.

During the entire period, a total of 97 000 000 self-testing kits were purchased. The kits were selected based on the rates of sensitivity (above 85% for samples testing positive with a qPCR up to the 33rd reaction cycle) and specificity (over 99%), which was supported through a peer reviewed process available at the FIND test directory (`https://www.finddx.org`). Importantly, these tests were safe, inexpensive, acceptable by the target population and simple to take at home or anywhere, providing rapid results.

C.2 Loss Function Minimization

The model was fitted by finding the parameter values that minimized the sum of squared log errors between the data and the model predictions for the outcomes discussed in Section 4.2.3.

To introduce some notation, consider an age group a and time period t. Let $\overline{\Delta H}_a(t)$ and $\overline{\Delta D}_a(t)$ be the corresponding number of new hospitalizations and deaths, respectively; let $\overline{P}_a^{\text{tot}}(t)$ be the corresponding total number of cases recorded; and let $\overline{P}_a^{\text{conf}}(t)$ be the corresponding number of cases reported through the self-testing program. The data and corresponding predictions are summarized in Table C.1.

A non-convex optimization algorithm was used to minimize the sum of squared log errors between the model predictions and the target data. In order to avoid overfitting, a regularization term was included in the loss function.

Data	**Model Estimate**
Hospitalizations	$\Delta_{ISS_{av}\to H_{av}}(t)$
Deaths	$\sum_{v\in\{0,1\}} \Delta_{H_{av}\to D_{av}}(t)$
Total Cases	$\sum_{v\in\{0,1\}} \sum_{X\in\mathcal{X}_{av}} P_X^{\text{self}}(t) + P_X^{\text{reg}}(t)$
Cases Through Self-Testing Program	$\sum_{v\in\{0,1\}} \sum_{X\in\mathcal{X}_{av}} P_X^{\text{self}}(t) + \gamma(t) P_X^{\text{reg}}(t)$

Table C.1: List of outcomes the compartmental model was fitted against, with corresponding expressions for the model estimates.

C.2.1 Parameters and Initialization

Certain model parameters were constrained in order to facilitate the fitting process. These are discussed next and summarized in Table C.2.

The probabilities of severe disease p_{av}^{severe}, though learned from the fitting process defined below, were constrained to lie within the following intervals:

$$0 \le p_{\text{0-18},v}^{\text{severe}} \le 0.02 \qquad 0 \le p_{\text{19-64},v}^{\text{severe}} \le 0.14 \qquad 0 \le p_{\text{65+},v}^{\text{severe}} \le 0.6$$

Parameter	Count (Wave 1)	Count (Wave 2)	Notes
$IAM_{av}(0)$	3	6	initial prevalence $< 0.3\%$
$ISM_{av}(0)$	3	6	initial prevalence $< 0.3\%$
$ISS_{av}(0)$	3	6	initial prevalence $< 0.3\%$
$\beta_{ab}(t)$	36	42	constant over 4 week intervals
p_{av}^{severe}	6	6	intervals from [52]
p_a^{die}	3	3	intervals from [58]
$\gamma_a(t)$	18	21	constant over 6 week intervals
Total	72	90	

Table C.2: Number of parameters learned by the model in each of the two waves.

which have widths twice the size of estimated values [52]. Similarly, all p_a^{die} were learnt and constrained to lie in intervals $[0, 2q_a]$, where q_a represents the total deaths divided by the total hospitalizations in age group a in a given wave, obtained from Greek NPHO data [58].

The mixing parameters, $\beta_{ab}(t)$, and the reporting parameters, $\gamma_a(t)$, were assumed to be constant in 4 and 6 week intervals respectively. These limits were empirically established with the goal of enabling tractability and avoiding overfitting.

The data include all dates between January 21, 2021 and December 15, 2021. During this period, Greece experienced two epidemic waves, one lasting from January 21, 2021 until June 20, 2021, and the other from June 21, 2021 until December 15, 2021. The self-testing program was introduced in the middle of the first waves and continued through the second wave. The model is fitted to both waves in separate runs. For the first wave, the sizes of all compartments were initialized at 0 with the following exceptions. The number of initially hospitalized patients were directly estimated from the raw data using Little's Law. The size of the susceptible compartments and infected compartments were allowed to be nonzero (and learned through the fitting process) subject to the following constraints: (i) the prevalence was assumed to be less than 0.3% (an upper bound supported by Bastani et al. [67]), (ii) the sum of all compartments summed to the population data per age group provided in the Greek census [57].

For the second wave, the sizes of all compartments were initialized at their levels learned from the model fitting to the first wave, with the exception of the infected compartments that were learned through fitting on the second wave. Similar to the first wave, a constraint was imposed that the sum of all compartments in each age group was equal to values provided in the Greek census [57]. Note that fitting the sizes of the infected compartments for the second wave was necessary, because the fit for these quantities from the first wave may not be credible due to end-of-horizon effects.

C.2.2 Loss Function

The loss function used to fit the model was the sum of squared log errors, across time steps, age groups, and waves for all outcomes presented in Table C.1.

There were more than 70 parameters to fit for each wave. In addition, the loss function was complex with no analytical gradients, making the underlying optimization problem a difficult task with the potential for overfitting. In order to obtain robust solutions, a regularization scheme was used on the time-varying parameters ($\beta_{ab}(t)$ and $\gamma_a(t)$) with a two-stage block minimization procedure to ease the optimization burden.

Regularization was a standard penalty on the absolute differences in successive values of the time-varying parameters. In particular, two regularization parameters were introduced, λ_1 and λ_2, and the following component was added to the loss function:

$$\lambda_1 \sum_{a\in\mathcal{A}} \sum_{b\in\mathcal{A}} |\beta_{ab}(t+1) - \beta_{ab}(t)| + \lambda_2 \sum_{a\in\mathcal{A}} |\gamma_a(t+1) - \gamma_a(t)|$$

The block-minimization approach worked in two steps:

1. Numerically minimize the regularized loss function with the cross-age-group mixing parameters fixed to zero ($\beta_{ab}(t) = 0$ for $a \neq b$), to obtain an initial estimate on the unknown parameters.

2. Numerically minimize the loss function by varying all mixing parameters, $\beta_{ab}(t)$,

keeping the remaining parameters fixed to their values from Step 1.

Both fitting steps were completed using the Levenberg-Marquardt algorithm [68] with random restarts. The regularization parameters were fixed in the optimization and chosen by searching through a grid (Figure C-1). The parameters were chosen to be as large as possible while retaining a good fit to the data (low mean absolute percentage error). All computational experiments were run on the SuperCloud infrastructure [69].

C.2.3 Bootstrapping for Confidence Intervals

To derive confidence intervals for our sensitivity analyses, an approach was followed that is similar to traditional bootstrapping. It relied on the data being split according to the 74 regional units of Greece, and used this split to construct *virtual* populations that were obtained by randomly sampling regional units (with replacement) until the total population of the collection was at least 2 million.

It was observed that the time series within individual regional units with small populations suffered from high variance, and therefore the 2 million resident threshold was selected so that the time series of the virtual populations were relatively stable. Figure C-2 shows the profiles of the 25 bootstrap samples.

In our analysis, we set $N = 25$ for tractability, and report 80% confidence intervals. Confidence intervals for the quantities of interest were derived by completing the fitting procedure and analysis on each of the bootstrapped datasets.

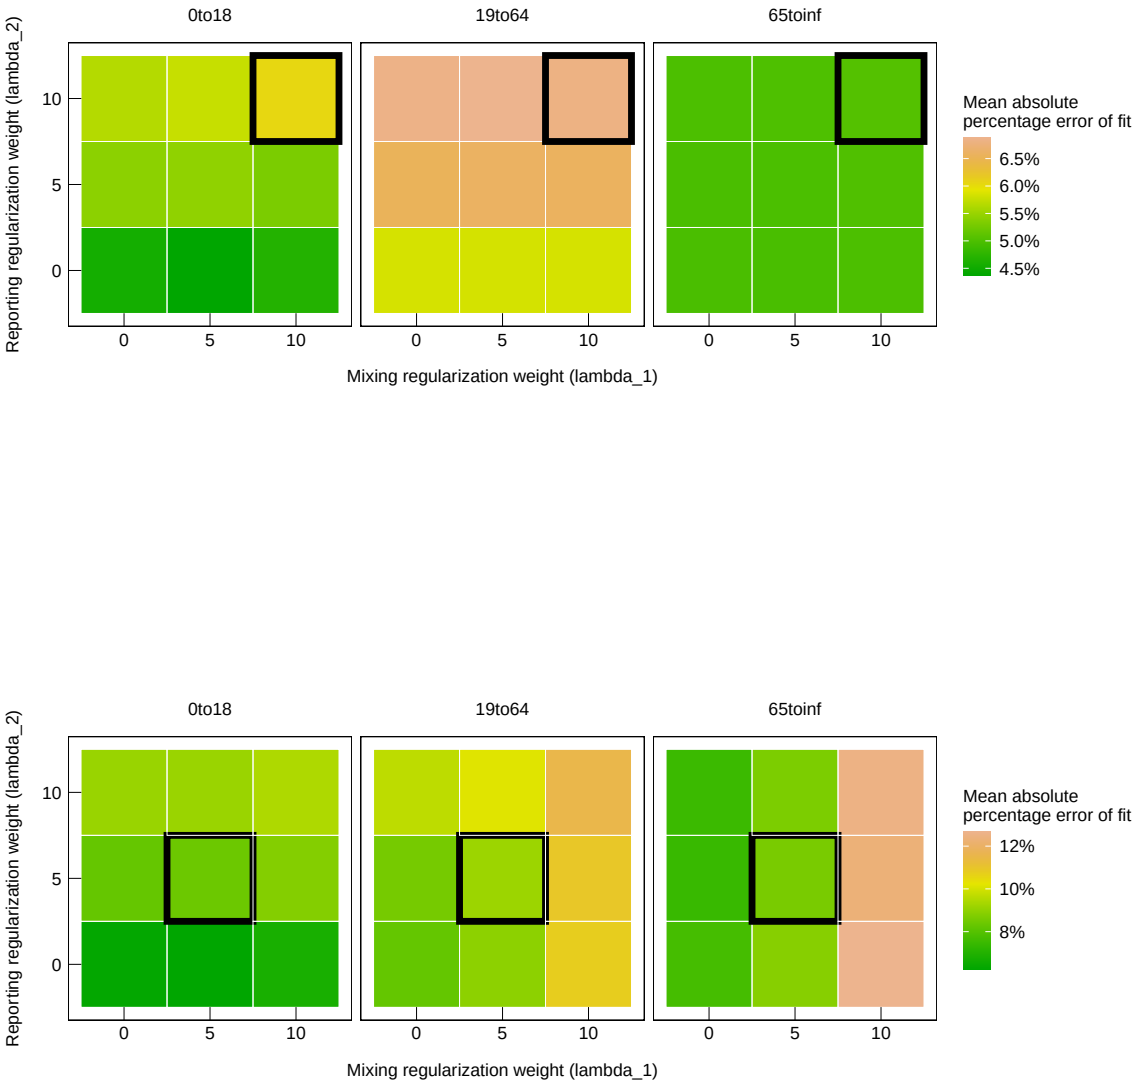

Figure C-1: Median values of the mean absolute percentage achieved by the fitted model across all bootstrap datasets as the regularization weights vary, for the first wave (top) and second wave (bottom).

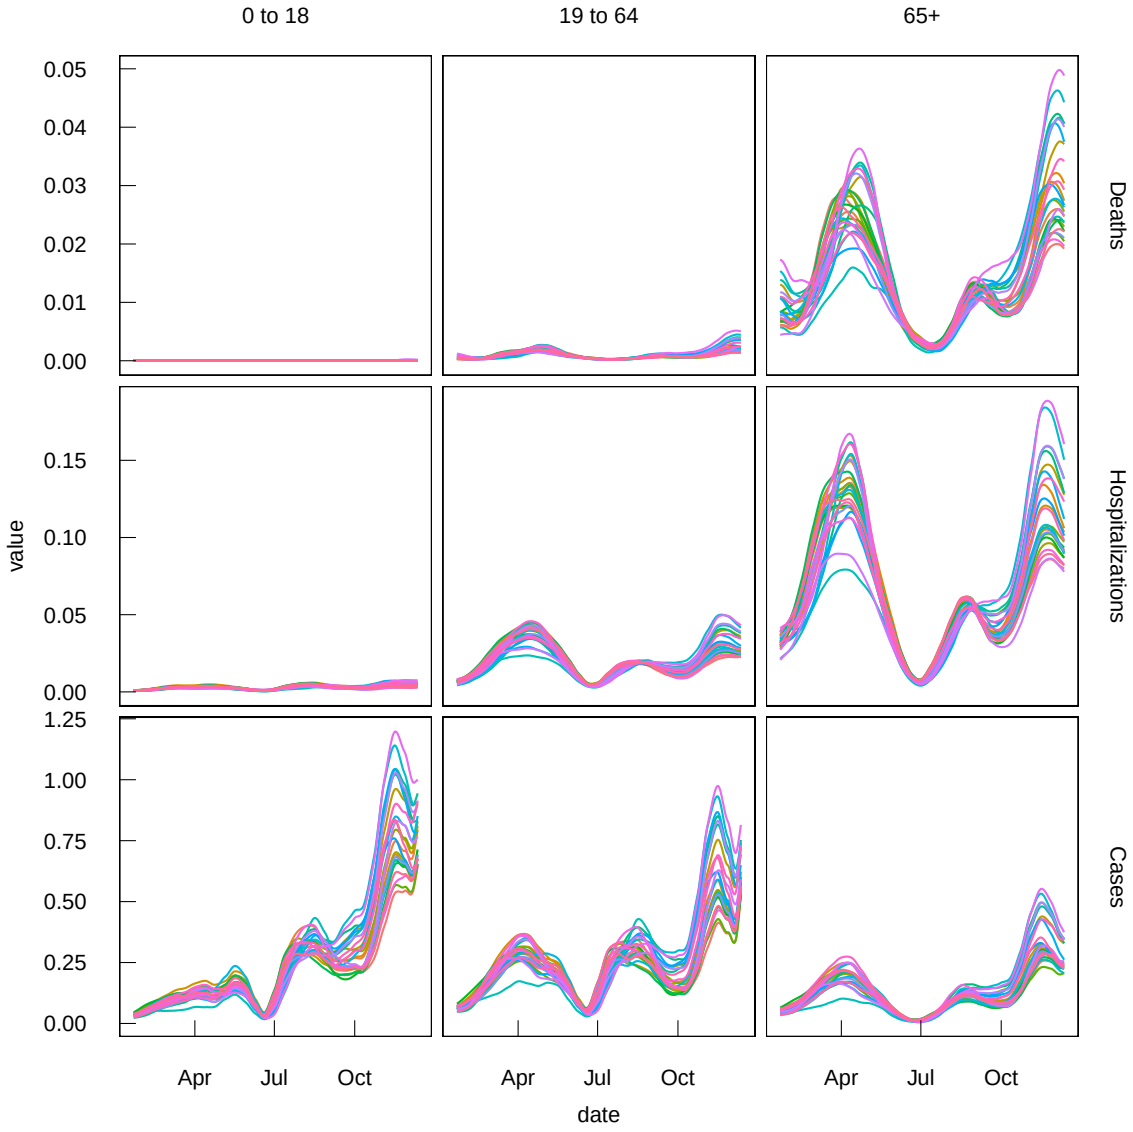

Figure C-2: Deaths, hospitalizations, and cases per person split by age group. Each colored series is a time series from a single bootstrap sample.

Bibliography

[1] Charles J. Dougherty. "Body futures: the case against marketing human organs". In: *Health Progress (Saint Louis, MO)* 68.5 (1987), pp. 51–55.

[2] Alexander Supady et al. "Allocating scarce intensive care resources during the COVID-19 pandemic: practical challenges to theoretical frameworks". In: *The Lancet Respiratory Medicine* (2021).

[3] Erika Satomi et al. "Fair allocation of scarce medical resources during COVID-19 pandemic: ethical considerations". In: *Einstein (Sao Paulo)* 18 (2020).

[4] Tal Zarsky. "The trouble with algorithmic decisions: an analytic road map to examine efficiency and fairness in automated and opaque decision making". In: *Science, Technology, & Human Values* 41.1 (2016), pp. 118–132.

[5] H. Peyton Young. *Equity: In Theory and Practice*. Princeton University Press, 1994.

[6] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. "Fairness, efficiency, and flexibility in organ allocation for kidney transplantation". In: *Operations Research* 61.1 (2013), pp. 73–87.

[7] Yichuan Ding, S. Thomas McCormick, and Mahesh Nagarajan. "A fluid model for one-sided bipartite matching queues with match-dependent rewards". In: *Operations Research* 69.4 (2021), pp. 1256–1281.

[8] Israel David and Uri Yechiali. "One-attribute sequential assignment match processes in discrete time". In: *Operations Research* 43.5 (1995), pp. 879–884.

[9] Rhonda Righter. "A resource allocation problem in a random environment". In: *Operations Research* 37.2 (1989), pp. 329–338.

[10] Peng Shi. "Optimal priority-based allocation mechanisms". In: *Management Science* 68.1 (2022), pp. 171–188.

[11] Itai Ashlagi and Peng Shi. "Optimal allocation without money: an engineering approach". In: *Management Science* 62.4 (2016), pp. 1078–1097.

[12] Aaron L. Bodoh-Creed. "Optimizing for distributional goals in school choice problems". In: *Management Science* 66.8 (2020), pp. 3657–3676.

[13] Philipp Afeche, Rene Caldentey, and Varun Gupta. "On the optimal design of a bipartite matching queueing system". In: *Operations Research* 70.1 (2022), pp. 363–401.

[14] Michael E. Sisselman and Ward Whitt. "Value-based routing and preference-based routing in customer contact centers". In: *Production and Operations Management* 16.3 (2007), pp. 277–291.

[15] Vijay Mehrotra et al. "Routing to manage resolution and waiting time in call centers with heterogeneous servers". In: *Manufacturing & Service Operations Management* 14.1 (2012), pp. 66–81.

[16] John C. Sparrow. *History of personnel demobilization in the United States Army*. Vol. 20. 210. Office of the Chief of Military History, Department of the Army, 1951.

[17] Vicki Been et al. "Allocation of the limited subsidies for affordable housing". In: *New York Times* (2018).

[18] Neil Thakral. *The public-housing allocation problem*. Tech. rep. Technical report, Harvard University, 2016.

[19] Hank Greely. "The equality of allocation by lot". In: *Harvard Civil Rights-Civil Liberties Law Review* 12 (1977), p. 113.

[20] Rhiannon Tudor Edwards. "Points for pain: waiting list priority scoring systems". In: *BMJ* 318.7181 (1999), pp. 412–414.

[21] Stefanos A. Zenios, Lawrence M. Wein, and Glenn M. Chertow. "Evidence-based organ allocation". In: *The American journal of medicine* 107.1 (1999), pp. 52–61.

[22] Ajay K. Israni et al. "New national allocation policy for deceased donor kidneys in the United States and possible effect on patient outcomes". In: *Journal of the American Society of Nephrology* 25.8 (2014), pp. 1842–1848.

[23] UNOS. *Transplant trends – UNOS*. `https://unos.org/data/transplant-trends/`. Last accessed on 2021-08-23. 2021.

[24] NLIHC. *The Gap | NLIHC*. `https://reports.nlihc.org/gap`. Last accessed on 2021-08-23. 2021.

[25] Vincent Delos and Denis Teissandier. "Minkowski sum of polytopes defined by their vertices". In: *Journal of Applied Mathematics and Physics* 45 (2015), pp. 62–67.

[26] Edoardo Amaldi and Viggo Kann. "The complexity and approximability of finding maximum feasible subsystems of linear relations". In: *Theoretical Computer Science* 147.1-2 (1995), pp. 181–210.

[27] Stephen P Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[28] Ezekiel J. Emanuel et al. "Fair allocation of scarce medical resources in the time of COVID-19". In: *New England Journal of Medicine* 382.21 (2020), pp. 2049–2055.

[29] Govind Persad, Alan Wertheimer, and Ezekiel J. Emanuel. "Principles for allocation of scarce medical interventions". In: *The Lancet* 373.9661 (2009), pp. 423–431.

[30] Ben Saunders. "A defence of weighted lotteries in life saving cases". In: *Ethical Theory and Moral Practice* 12.3 (2009), pp. 279–290.

[31] Douglas B White and Derek C Angus. "A proposed lottery system to allocate scarce COVID-19 medications: promoting fairness and generating knowledge". In: *JAMA* 324.4 (2020), pp. 329–330.

[32] Alexander A. Iyer, Saskia Hendriks, and Annette Rid. "Advantages of using lotteries to select participants for high-demand COVID-19 treatment trials". In: *Ethics & Human Research* 42.4 (2020), pp. 35–40.

[33] Lynn A. Jansen and Steven Wall. "Weighted lotteries and the allocation of scarce medications for COVID-19". In: *Hastings Center Report* 51.1 (2021), pp. 39–46.

[34] Olle Ten Cate. "Rationales for a lottery among the qualified to select medical trainees: decades of Dutch experience". In: *Journal of Graduate Medical Education* 13.5 (2021), pp. 612–615.

[35] OPTN. *Ethical principles in the allocation of human organs*. `https://optn.transplant.hrsa.gov/professionals/by-topic/ethical-considerations/ethical-principles-in-the-allocation-of-human-organs/`. Last accessed on 2023-05-02. 2015.

[36] Andrew Wey et al. "Influence of kidney offer acceptance behavior on metrics of allocation efficiency". In: *Clinical Transplantation* 31.9 (2017), e13057.

[37] Ashley Y. Choi et al. "Transplant center variability in organ offer acceptance and mortality among US patients on the heart transplant waitlist". In: *JAMA Cardiology* 5.6 (2020), pp. 660–668.

[38] Xuanming Su and Stefanos Zenios. "Patient choice in kidney allocation: the role of the queueing discipline". In: *Manufacturing & Service Operations Management* 6.4 (2004), pp. 280–301.

[39] Xuanming Su and Stefanos A. Zenios. "Recipient choice can address the efficiency-equity trade-off in kidney transplantation: a mechanism design model". In: *Management Science* 52.11 (2006), pp. 1647–1660.

[40] Refael Hassin and Moshe Haviv. *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Vol. 59. Springer Science & Business Media, 2003.

[41] Folkert O. Belzer and James H. Southard. "Principles of solid-organ preservation by cold storage". In: *Transplantation* 45.4 (1988), pp. 673–676.

[42] Nikolaos Askitas, Konstantinos Tatsiramos, and Bertrand Verheyden. "Estimating worldwide effects of non-pharmaceutical interventions on COVID-19 incidence and population mobility patterns using a multiple-event study". In: *Scientific Reports* 11 (Jan. 2021), p. 1972.

[43] Berber T Snoeijer et al. "Measuring the effect of non-pharmaceutical interventions (NPIs) on mobility during the COVID-19 pandemic using global mobility data". In: *NPJ Digital Medicine* 4.1 (2021), p. 81.

[44] Sebastian Mader and Tobias Rüttenauer. "The effects of non-pharmaceutical interventions on COVID-19 mortality: A generalized synthetic control approach across 169 countries". In: *Frontiers in Public Health* (2022), p. 740.

[45] Seth Flaxman et al. "Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe". In: *Nature* 584.7820 (2020), pp. 257–261.

[46] Antoine Mandel and Vipin Veetil. "The economic cost of COVID lockdowns: an out-of-equilibrium analysis". In: *Economics of Disasters and Climate Change* 4 (2020), pp. 431–451.

[47] Anton Pak, Oyelola A Adegboye, and Emma S McBryde. "Are we better-off? The benefits and costs of Australian COVID-19 lockdown". In: *Frontiers in Public Health* (2021), p. 1973.

[48] NPHO of Greece. *COVID-19 Vaccinations in Greece [Unpublished raw data]*. 2021.

[49] Tommy Nyberg et al. "Comparative analysis of the risks of hospitalisation and death associated with SARS-CoV-2 omicron (B. 1.1. 529) and delta (B. 1.617. 2) variants in England: a cohort study". In: *The Lancet* 399.10332 (2022), pp. 1303–1312.

[50] Toon Braeye et al. "Vaccine effectiveness against onward transmission of SARS-CoV-2-infection by variant of concern and time since vaccination, Belgian contact tracing, 2021". In: *Vaccine* 40.22 (2022), pp. 3027–3037.

[51] Qiuyue Ma et al. "Global percentage of asymptomatic SARS-CoV-2 infections among the tested population and individuals with confirmed COVID-19 diagnosis: a systematic review and meta-analysis". In: *JAMA Network Open* 4.12 (2021), e2137257–e2137257.

[52] Tjede Funk et al. "Age-specific associations between underlying health conditions and hospitalisation, death and in-hospital death among confirmed COVID-19 cases: a multi-country study based on surveillance data, June to December 2020". In: *Eurosurveillance* 27.35 (2022), p. 2100883.

[53] Andrew William Byrne et al. "Inferred duration of infectious period of SARS-CoV-2: rapid scoping review and analysis of available evidence for asymptomatic and symptomatic COVID-19 cases". In: *BMJ Open* 10.8 (2020), e039856.

[54] Christel Faes et al. "Time between symptom onset, hospitalisation and recovery or death: statistical analysis of Belgian COVID-19 patients". In: *International Journal of Environmental Research and Public Health* 17.20 (2020), p. 7560.

[55] NPHO of Greece. *Length of Hospital Stays for COVID-19 Patients in Greece [Unpublished raw data]*. 2021.

[56] NPHO of Greece. *Self-Testing Figures and Cases Confirmed Positive [Unpublished raw data]*. 2021.

[57] Hellenic Statistical Authority. *2011 Population-Housing Census*. https://www.statistics.gr/el/statistics/-/publication/SAM03/2011. Last accessed on 2023-05-02. 2011.

[58] NPHO of Greece. *Hospitalizations and Deaths for COVID-19 Patients in Greece [Unpublished raw data]*. 2021.

[59] NPHO of Greece. *Daily COVID-19 Cases*. https://eody.gov.gr/category/covid-19/. Last accessed on 2023-05-02. 2021.

[60] Francisco Arroyo-Marioli et al. "Tracking R of COVID-19: a new real-time estimation using the Kalman filter". In: *PloS one* 16.1 (2021), e0244474.

[61] Spyros Sapounas et al. "Cold-season epidemic dynamics of COVID-19 in two major metropolitan areas in Greece: hypotheses and implications for public health interventions". In: *Frontiers in Medicine* 9 (2022).

[62] Ed Holt. "Slovakia to test all adults for SARS-CoV-2". In: *The Lancet* 396.10260 (2020), pp. 1386–1387.

[63] Jaroslav Frnda and Marek Durica. "On pilot massive COVID-19 testing by antigen tests in Europe. Case study: Slovakia". In: *Infectious Disease Reports* 13.1 (2021), pp. 45–57.

[64] Martin Pavelka et al. "The impact of population-wide rapid antigen testing on SARS-CoV-2 prevalence in Slovakia". In: *Science* 372.6542 (2021), pp. 635–641.

[65] Daniel B. Larremore et al. "Test sensitivity is secondary to frequency and turnaround time for COVID-19 screening". In: *Science Advances* 7.1 (2021), eabd5393.

[66] Mélodie Monod et al. "Age groups that sustain resurging COVID-19 epidemics in the United States". In: *Science* 371.6536 (2021), eabe8372.

[67] Hamsa Bastani et al. "Efficient and targeted COVID-19 border testing via reinforcement learning". In: *Nature* 599.7883 (2021), pp. 108–113.

[68] Matt Newville et al. "lmfit/lmfit-py: 1.0.3". In: *Zenodo* (2021).

[69] Albert Reuther et al. "Interactive supercomputing on 40,000 cores for machine learning and data analysis". In: *2018 IEEE High Performance Extreme Computing Conference*. IEEE. 2018, pp. 1–6.