DRAFT

Social learning in lung transplant decision

Preliminary and incomplete. Please do not circulate.[1]

Laura Doval[2] Federico Echenique[3] Wanying Huang[4]
Yi Xin[5]

August 16, 2024

[1]Prepared for submission to NBER Market Design 2024. Do not cite or circulate without the authors' explicit permission. The data reported here have been supplied by the Hennepin Healthcare Research Institute (HHRI) as the contractor for the Scientific Registry of Transplant Recipients (SRTR). The interpretation and reporting of these data are the responsibility of the author(s) and in no way should be seen as an official policy of or interpretation by the SRTR or the U.S. Government.

[2]Columbia Business School and CEPR. E-mail: laura.doval@columbia.edu
[3]University of California, Berkeley. E-mail: fede@econ.berkeley.edu
[4]Monash University. E-mail: kate.wanyinghuang@gmail.com
[5]California Institute of Technology. E-mail: yixin@caltech.edu

Abstract

We study the allocation of deceased-donor lungs to patients in need of a transplant. Patients make sequential decisions in an order dictated by a priority policy. Using data from a prominent Organ Procurement Organization in the United States, we provide reduced-form evidence of social learning: because patients accept or reject in sequence, they may "herd" on rejecting an organ that would otherwise be accepted. We develop and estimate a structural model, using it to quantify different policy proposals and informational regimes. While blinding patients to their position in a sequence leads to less organ waste, it also results in an allocation of organs that provide patients with smaller average utility (the utilitarian objective is, however, higher). A policy of ordering patients by their probability of acceptance exacerbates social learning, leading to higher waste, but also raises utility per accepted organ and speeds up allocation.

KEYWORDS: *market design; organ allocation; social learning.*

JEL CODES: D47, D61.

1 Introduction

We study the allocation of deceased-donor organs to patients in need of a transplant, focusing on lung donations. An organ is offered, in sequence, to an ordered list of transplant centers who make a decision on behalf of a patient. Whether to accept an organ offer is one of the most important decisions made by a transplant center, and one that is marred with uncertainty. Within a small window of time, the transplant center needs to assess the match quality between their patient and the organ offer. The transplant center's decision is consequential: patients who reject offers for which they have high priority are more likely to leave the system unmatched, prolonging their compromised quality of life. Meanwhile, accepting a poor match quality offer can lead to a lower probability of the patient surviving after transplantation (Gilroy et al., 2019).

To facilitate decision-making, transplant centers are provided with information about donor characteristics, often including the patient's position in the order of offers for the organ. In particular, when a transplant patient is offered a lung, their transplant center is informed of the patient's *sequence number in the match run*–the number indicating their priority, with a lower number corresponding to a higher priority–along with the reason why previous offers, if any, were rejected (Harhay et al., 2019). More recently, the United Network for Organ Sharing (UNOS) has ensured that information about the patient's sequence number is readily and systematically available for all organs.[1]

At the same time, rejections early in the match run often lead to further rejections later in the match run (Gilroy et al., 2019), motivating policy proposals to hide match run positions from transplant centers (Cohen et al., 2018a; Benvenuto and Aversa, 2019). Indeed, when evaluating whether to accept an offer, a transplant center aware of their patient's sequence number can aggregate both the information provided about the donor and the details of earlier rejections. Since organ quality is the most frequently cited reason for rejecting an organ offer, receiving an offer at a

[1]See https://unos.org/news/match-summary-view-now-available-in-donornet-desktop/.

higher sequence number may indicate a harder-to-place organ (Goldberg et al., 2016). Thus, while providing information about early rejections could facilitate decision-making, it may also be detrimental to organ placement if this information crowds out a transplant center's own evaluation of the organ.

In this paper, we use detailed data from a large organ procurement organization to assess how the availability of information on sequence number affects organ discard rates and patient welfare in the allocation of deceased-donor lungs. We make two contributions. First, we provide reduced-form evidence of *social learning*: even after controlling for all donor- and patient-risk factors that determine a patient's position in the match run, the probability of accepting a pair of deceased-donor lungs decreases as the sequence number increases. In other words, patients with lower priority on a match run are more likely to reject the organ than those with higher priority. This finding suggests that transplant centers interpret rejections by centers with a low sequence number as an indication of a lower-quality organ. Second, we estimate a structural model of sequential assignment and use the model estimates to evaluate the welfare implications of different priority and information policies in the presence of social learning.

To assess the presence of social learning, we exploit a feature of the allocation of deceased-donor organs in the US, known as *provisional yes*. Under this system, organs are allocated in a two-stage process: in the first stage, compatible patients are asked whether they are interested in the organ; in the second stage, the organ is offered among those patients who responded "yes," and it is assigned to the lowest sequence number patient who accepts it. By comparing provisional-yes decisions–which are unconditional on receiving the organ–with final acceptance decisions–which are conditional on receiving the offer, after it potentially being rejected–we can measure the effect of sequence numbers on the probability of accepting an organ. Note, however, that the match run position is itself a measure of match-quality between the patient and the organ, making it natural for sequence number to be negatively correlated with acceptance probability. To measure the informational content of the match run position–above and beyond the factors that go into the sequence

number–we analyze the relationship between match run position and provisional-yes and final-acceptance decisions, while controlling for all patient and donor factors that determine a patient's priority for the organ.

Section 3 presents reduced-form evidence that information about match run position negatively affects the probability of organ placement. Controlling for all donor and patient risk factors, we find that, all else being equal, the probability of finally accepting an organ decreases as the sequence number increases, while the probability of provisionally accepting an organ remains unaffected by sequence number. These results suggest that an organ offered at a higher sequence number is perceived as bad news of the organ's quality. At the same time, our data features a significant fraction of discarded organs: 47% of lungs are not allocated to any patient. Because a match run is only initiated for organs deemed viable, this high discard rate indicates substantial waste. Inefficiency and waste is always undesirable, but they are especially troubling when it involves the allocation of scarce organs to patients in need of a transplant.

Motivated by our reduced-form findings, in Section 4 we propose a structural model of sequential assignment. The model adapts the canonical social learning model of Bikhchandani et al. (1992) to the specific institutional details of deceased-donor organ allocations. In our model, each patient compatible with a given donor first decides whether to provisionally accept an offer and, if they do, then decides whether to finally accept the organ in question. Organ quality is uncertain, and each patient learns imperfectly about quality from a private signal. The decision to provisionally accept an offer is influenced only by patient idiosyncratic factors, while the final acceptance decision depends both on the patient's private signal and the information inferred from the rejections by patients with smaller sequence numbers. In Section 5, we present the model estimates, which reveal a high prior probability of low-quality organs and a high precision in the patient's private signals. Our estimated model closely replicates both the data and the reduced-form evidence.

Using our estimated model, we analyze the welfare impact of different priority policies

and informational regimes, while taking into account the effects of social learning. Our first set of counterfactual experiments examines two contrasting priority ranking policies: under the *greedy priority* policy, patients are ranked in decreasing order of their probability of accepting the organ in question, while under the *reverse greedy priority* policy, patients are ranked in increasing order of their probability of acceptance. The greedy policy offers organs to patients who value it the most, and would seem at first to lead to faster acceptances. Social learning, however, complicates this dynamic: the decision to reject made by a patient who is *a priori* more likely to accept, carries a stronger negative inference about the organ than the same decision made by a patient who is less likely to accept. This amplifies the negative perception of organ quality under the greedy policy, as a series of rejections by high-priority patients reinforces doubts about the organ.

Our counterfactual results (see Table 6 in Section 6) show that the greedy priority policy turns out to allocate fewer organs compared to the current policy, but these allocations occur to patients with lower sequence numbers. Moreover, under the greedy priority policy, patients who ultimately receive the organ are better off overall since they value the organ highly, and low-quality organs are rejected more frequently than under other policies. In sum, while the greedy policy leads to higher waste–through social learning–it ensures that organs are allocated to patients who value them more.

In contrast, the reverse greedy policy allocates more organs than the current policy, resulting in less waste. The allocated organs are accepted by patients with higher sequence numbers, and the utility of patients who accept the organs is lower under the reverse greedy policy. This lower utility is partly because early rejections under reverse greedy provide less information about organ quality, leading to a higher acceptance rate of low-quality organs.

Our second set of counterfactual experiments compares different informational regimes under the current priority policy, which serves as our baseline. First, we implement a *no social learning* regime by blinding patients to their position in the match run,

and allocating the organ to the patient with the lowest sequence number among those who accept the offer. Compared to the baseline, this policy allows for less information aggregation, but results in the highest acceptance rate by both increasing the acceptance of high-quality organs and reducing the rejection of low-quality ones. Consequently, it also leads to lower utility among patients who receive a transplant. Second, we consider an *information sharing* regime that allows patients to access the signals of all other patients in the same match run. Compared to the baseline, this regime facilitates greater information aggregation while eliminating the accumulation of negative news through social learning. As a result, this policy achieves a higher acceptance rate and yields high utility for patients who accept an offer, similar to the outcomes achieved by the greedy priority policy (see further details in Table 7).

Altogether our results highlight the trade-offs of allowing patients to have access to information about their position in the match run. As our evidence shows, early rejections contain information about match quality above and beyond that which is used to compute a patient's priority for an organ. This is consistent with proposals to have organ procurement organizations share more and better information with transplant centers: absent this, information about who and why rejected an organ can supplement the information that a transplant center has access to, which partly justifies UNOS' recent decision to make information about match run position and early rejections readily available across all organs. In other words, information about match run positions allows the system to aggregate the information contained in early rejections. At the same time, early rejections garble the information leading to those rejections, which in the presence of social learning can lead to excessive rejections. Thus, the more information is available to a transplant center, the less will be the role for inferences based on the coarse information obtained from other centers' decisions. Our counterfactual analysis suggests that the greedy priority performs the best in terms of achieving higher patient utility and a faster allocation process, with the caveat of a slightly lower acceptance rate due to social learning.

Related Literature We are closest to the contribution of Zhang (2010), who studies the role of observational learning in the allocation of deceased donor kidneys. Our result that, all else equal, sequence number negatively affects the acceptance probability replicates a similar observation in Zhang (2010) in the context of deceased donor kidneys. Through counterfactual simulations, the author shows that eliminating social learning leads to less organ wastage. More recently, De Mel et al. (2020) study the ability of transplant centers to assess organ quality and how ability heterogeneity impacts the matching process between organs and transplant centers. In contrast to Zhang (2010), De Mel et al. (2020) find that eliminating the possibility of observational learning across transplant centers does not significantly improve the organ discard rate.[2]

Several market design papers (empirically) study the efficiency and welfare properties of deceased donor organ allocation, suggesting improvements. Focusing on the allocation of deceased-donor kidneys, Agarwal et al. (2018, 2021) study the effect of dynamic incentives on organ wastage. Shi and Yin (2022) show that eliminating organ wastage is at odds with prioritizing patients on the basis of their waiting time. Muñoz-Rodriguez (2021) studies the effects of asymmetric information in the allocation of hearts and livers, in which doctors are better informed about patients' urgency and readiness status than the organ procurement agency.

Lastly, the connection between match run position at which an organ is accepted and transplantation outcomes is extensively studied in the medical literature (see, e.g., Cohen et al., 2018a for kidneys, Goldberg et al., 2016 for liver, and Harhay et al., 2019 for lungs). Most studies conclude that while transplantation outcomes–mortality and graft survival–and match run position at which an organ is accepted are negatively correlated, these correlations are not statistically significant once one controls for the transplant center that accepts the organs. The conclusion is then that transplant centers that accept organs at later positions, which are typically higher

[2]Social learning has been empirically studied in other contexts (see, e.g., Cipriani and Guarino (2014) for a financial application, and Newham and Midjord (2018) for a voting application within the FDA).

volume centers, have more experience in assessing organ quality and are willing to take more organs because they have a larger patient population.

This strand of medical literature, however, acknowledges several limitations: (i) the analysis is done solely on organs that are eventually transplanted, without accounting for selection bias, (ii) other outcomes such as hospital length of stay, duration of mechanical ventilation, postoperative functional status, quality of life, and transplant cost are not considered, and (iii) the conclusions rely on the assumption that the data about donor-risk factors available to the econometrician is equally available to the transplant center at the time of offer evaluation, whereas many recent policy proposals focus on making donor data more readily available to evaluating transplant centers (Gilroy et al., 2019; Hackmann et al., 2022). For these reasons, the medical community has requested a randomized control trial in which some transplant centers are blinded to their position in the match run to meaningfully evaluate the impact of match run position on transplant outcomes (Cohen et al., 2018b).

Organization The rest of the paper is organized as follows. Section 2 describes the institutional setting and our data. Section 3 provides reduced form evidence of social learning in the allocation of deceased donor lungs. Section 4 presents the structural model, and Section 5 discusses the results of the model estimation. Finally, using the model estimates as a baseline, Section 6 analyzes the outcomes of counterfactual experiments on different information and allocation policies across different measures of efficiency and welfare.

2 Institutional Background and Data

2.1 Institutional Background

An organ procurement organization (OPO) is a non-profit agency responsible for recovering organs from deceased donors for transplantation in the United States. There are 56 such agencies, each legally permitted to operate within an assigned donation service area. The OPO from which we obtained data is one of the largest

organ procurement organizations in the United States, recovering organs and tissues for transplantation and serving 215 hospitals and 9 transplant centers in a region with a population of 20 million. In assigning organs, this OPO implements the assignment rules outlined by the Organ Procurement and Transplantation Network (henceforth, OPTN), which we will discuss next.

OPTN Priority Assignment In the OPTN system for organ allocations, patients are ordered according to a priority system. A patient's priority is reflected in their *sequence number*: patients are offered an organ in the order of their sequence number, with those holding lower numbers receiving offers before those with larger numbers. The process of sequentially offering an organ to a list of patients is called a *match run*.

According to the OPTN policy guide (OPTN, 2019), a lung patient's sequence number is determined by four factors: (i) their lung allocation score (LAS), ranging from 0 to 100, calculated based on a combination of waitlist urgency and post-transplant survival probabilities; (ii) their blood-type compatibility with the donor; (iii) their total wait time; and (iv) their physical proximity to the donor. Specifically, each donor-patient pair is initially classified into a priority group based on physical proximity and blood type compatibility. Within each classified priority group, donor-patient pairs are prioritized in descending order of the patient's LAS. For patients with the same LAS, ties are broken by prioritizing those who have waited longer.

In summary, when assigning sequence numbers to patients, priority is given first to those with closer proximity and better blood-type compatibility with the donor. Following this initial assignment of priority groups, the final sequence number is determined by tie-breaking, which awards higher priority to patients with a larger LAS, and then, to those who have waited longer.

The Allocation System When a viable lung becomes available for transplantation, the OPTN system offers the organ to each patient in the order dictated by their sequence number in a match run. Before receiving the offer, each patient has the

opportunity to opt out of the match run through a practice known as "provisional yes". Provisional yes allows patients on the waitlist to *provisionally* accept or reject a deceased donor lung offer within an hour. Patients who provisionally reject an organ opt out of the process entirely. According to OPTN, the purpose of provisional yes is to inform the acting organ procurement agency that the patient (and their medical team) has received the offer and is interested in moving ahead, either by accepting the organ or receiving more information (OPTN, 2019). Introduced in 2018, the provisional yes practice addresses concerns about patients who are not fit for transplantation but still appear on the waiting list, causing delays in the allocation process. By enabling these unfit patients to opt out, provisional yes aims to streamline and expedite organ allocation by maintaining an active list of potential recipients.

The allocation process then moves to a second stage, where only those who have provisionally accepted the offer can decide whether to *finally* accept or reject it. When making decisions, patients are informed of their sequence number, along with the reason why prior offers, if any, were rejected (Harhay et al., 2019). If all final decisions result in rejection, the organ is then discarded. Importantly, neither provisional nor final rejections adversely affect patients' priority in future assignments.

Notice that when receiving a second-stage offer, a patient with lower priority, i.e., a higher sequence number, should apply the following contingent reasoning: they should make their final decision conditional on the event that all higher-priority patients who provisionally accepted the organ have subsequently rejected it, since otherwise, the lower-priority patient would not receive it even if they chose to accept. Consequently, if these higher-priority patients' rejection decisions are informative about organ quality, this reasoning creates channels for social learning.

2.2 Data

Our data includes all lung donations managed by the concerned OPO between January 2019 and December 2020. For each donor organ, we observe a rich set of donor

and patient characteristics, as well as detailed records of their acceptance and rejection decisions.[3] The only key missing variable is a patient's total wait time at the time of the offer. We construct this variable using data from the Scientific Registry of Transplant Recipients, which includes the initial wait-listing date for each patient.[4] This allows us to calculate the patient's total wait time for each offer by taking the difference between the date of the offer and their initial wait-listing date.[5]

To keep our empirical model tractable, we focus on patients in need of a bilateral lung transplant, which constitutes about 91% of the total number of patients. We exclude offers involving patients under 12 years of age, as they are prioritized differently by the OPTN allocation system. In addition, we exclude patients who require both heart and lung transplants, as the challenges faced by those in need of a bilateral lung transplant differ substantially from those faced by patients needing a heart-lung transplant or a single lung transplant.[6]

Our final dataset comprises 32,786 offers, involving 548 deceased donors and 1,348 patients actively wait-listed for a bilateral lung transplant during our observational period. It is worth noting that while all patients in our final dataset require a bilateral lung transplant, about 5% of all accepted donors (out of 290) had their lungs allocated to two different patients. For these donors, we treat the two allocations as a single acceptance and use the lower sequence number among the two allocated patients as the corresponding accepted sequence number.

The high rate of organ waste is a striking feature of our data. With approximately

[3]This includes the timing of decisions and the recorded reason for rejections.

[4]This study used data from the Scientific Registry of Transplant Recipients (SRTR), which includes data on all donors, wait-listed candidates, and transplant recipients in the US, submitted by members of the Organ Procurement and Transplantation Network (OPTN). The Health Resources and Services Administration (HRSA) of the U.S. Department of Health and Human Services provides oversight to the activities of both the OPTN and SRTR contractors.

[5]We exclude periods in which the patient was recorded as inactive, as wait time does not accumulate over these periods.

[6]Patients in need of a bilateral lung transplant typically suffer from more severe lung disease or impaired lung function. Compared to a single lung transplant, the procedure for a bilateral lung transplant generally involves higher risks and longer recovery times, but offers better long-term survival prospects.

1,300 patients in need of a bilateral lung transplant and only 548 donors available, demand is more than double the supply. Despite this, the final allocation rate for these organs is only about 53%. Specifically, out of the 548 deceased donors, about 258 had their lungs discarded. This substantial level of wastage underscores the need to evaluate the efficiency and welfare outcomes of the current priority policy and compare it with alternative policies, as discussed in Section 6.

Patient and Donor summary statistics Table 1 summarizes patients' and donors' characteristics in our sample. Panel (A) of Table 1 reports sample statistics for patient characteristics, including health indicators such as body mass index (BMI) and diabetic status, together with medical history, such as whether the patient had any prior lung transplant. On average, patients wait just over three months before receiving their first lung transplant offer. The average LAS is approximately 43 points, with more than 80% of patients having LAS between 30 and 50. Recall that a higher LAS indicates a greater likelihood of benefiting from a lung transplant, and the OPTN allocation system prioritizes patients with higher LAS.

During the sample period, patients received an average of 24 offers in total. Note, however, the median number of offers is only about 6, indicating that many patients received substantially fewer offers. The positive skewness in the number of offers received is likely due to a small portion of patients with lower LAS remaining on the waitlist for extended periods, thereby receiving a large number of offers.

Panel (B) of Table 1 reports the sample statistics for donor characteristics. Donors are, on average, 18 years younger than the patients. About 20% of donors have either a history of heavy alcohol consumption, or a status for increased risk, suggesting considerable variability in the quality of donor lungs. Despite this variability, the overall health of donor lungs is good, as evidenced by the variable "P/F ratio," which averages around 420, falling within the normal and healthy range of 400 to 500 measured at sea level.[7] Nevertheless, although each donor is offered to an average

[7]The P/F ratio is defined as the ratio between the partial pressure of arterial oxygen (PaO2, measured in mmHg) and the fraction of inspired oxygen (FiO2, expressed as a decimal). A P/F

Table 1: Patient and Donor Characteristics

	Mean	S.D.	Min	Max
Panel A: Patient Characteristics ($N = 1,348$)				
Age	55.9	13.5	12	79
Female	51.7%	50.0%	0	100
Caucasian	67.5%	46.9%	0	100
Body Mass Index (BMI)	25.5	4.5	15	37
Diabetic patient	16.8%	37.4%	0	100
Previous lung transplant	3.6%	18.7%	0	1
Lung Allocation Score (LAS)	43.3	14.3	5.9	95.4
Wait time (months)	3.4	7.5	0	87.1
Number of offers per patient	24.3	42.7	1	347
Panel B: Donor characteristics ($N = 548$)				
Age	38.2	14.1	7	75
Weight (kg)	79.0	20.8	23.5	170.7
Height (cm)	169.5	10.8	119	198.5
Heavy alcohol history	17.7%	38.2%	0	1
IV drug use history	6.0%	23.8%	0	1
Increasing risk status	21.2%	40.9%	0	1
P/F ratio	419.3	102.3	80	1400
Number of offers per donor	59.8	48.4	1	222

*N*ote: For all patient characteristics that can vary throughout our observational period, such as age, BMI, LAS, and wait time, we use the value recorded at the first offer that each patient receives.

of 60 patients, fewer than 50% of them are ultimately accepted.

To see if there are any differences in lung quality between accepted and rejected donors, Table 2 reports summary statistics of donor characteristics by their final allocation outcome. As expected, accepted donors generally exhibit higher quality compared to rejected donors. For example, accepted donors are, on average, three years younger and have slightly lower probabilities of a history of heavy alcohol consumption or intravenous drug use compared to rejected donors. More importantly, the direct measure of lung quality, the P/F ratio, is approximately 10% higher in

ratio below 300 suggests some level of impaired lung function.

Table 2: Donor Characteristics by Allocation Outcome

	Accepted Donors		Rejected Donors	
	Mean	S.D.	Mean	S.D.
Age (year)	36.9	13.4	39.7	14.7
Weight (kg)	80.1	20.3	77.8	21.3
Height (cm)	169.5	10.6	169.4	11.1
History of heavy alcohol consumption	17.2%	37.8%	18.2%	38.7%
History of IV drug use	5.5%	22.9%	6.6%	24.9%
Increasing risk status	21.4%	41.1%	20.9%	40.8%
P/F ratio	437.2	83.0	399.2	117.4
Number of offers per donor	68.5	48.2	50.1	46.9
Number of observations	$N = 290$		$N = 258$	

accepted donors than in rejected ones, indicating that the former has better lung functionality than the latter. We also note that lungs from accepted donors are offered to approximately 20 more patients on average than rejected donors.

Acceptance Rate and Offers As described in Section 2.1, patients first decide whether to provisionally accept the offer before making their final decision. So the observed low final acceptance rate–only about 53%–could stem from either a high provisional rejection rate or a scenario where a large proportion of donors are provisionally accepted but ultimately rejected. Our data suggests that the latter is the case: more than 95% of all offers are provisionally accepted, yet only 0.9% of them result in final acceptance. It is important to note that this low final acceptance rate is an underestimation, as our data on final acceptance decisions are truncated at the first acceptance decision in a match run. In other words, we do not observe the hypothetical final decisions of patients with lower priority who could have made decisions after the first acceptance. Therefore, we exclude all offers following the first acceptance from our calculation of the final acceptance rate. We are left with 15,773 offers for which both provisional and final decisions are recorded.

Table 3 reports the key characteristics of offers at different decision stages. The table

Table 3: Offers Characteristics by Decision Stage

	All		Provisional Yes		Final Yes	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
Lung Allocation Score (LAS)	39.74	10.13	39.68	10.03	54.47	20.27
Primary blood type match (%)	76.40	42.47	76.43	42.44	96.90	17.37
Wait time (month)	8.46	11.25	8.45	11.20	5.84	9.98
Distance (per NM)	247.95	212.58	239.80	186.84	119.00	130.79
Decision Time (minute)	-	-	10.28	91.68	207.48	442.71
Number of observations	$N = 32,786$		$N = 31,329$		$N = 290$	

Note: The variable "primary blood type match" is a categorical variable with three levels: 0 for incompatible, 1 for moderate compatibility and 2 for the highest compatibility. All observations in the dataset have a primary blood type match rating of at least 1. In addition, the unit for wait time is in months, and the unit for distance is measured in nautical miles.

sheds light on the acceptance decisions from the patients' perspective. As most offers are provisionally accepted, the characteristics among all offers and those provisionally accepted are very similar. On average, an offer is provisionally accepted by a patient with a LAS of about 40, who has been waiting for approximately 8 months, has about a 76% chance of being a primary blood type match with the donor, and is located within 240 nautical miles of the donor.[8] The only noticeable difference between all offers and those provisionally accepted is the slightly shorter distance between the donor and the patient.

In terms of decision time, provisional decisions occur much faster than final decisions. On average, the former takes only about 10 minutes, whereas the latter requires an additional 3 hours. We also observe that offers that are ultimately accepted tend to be associated with patients with higher LAS, are more likely to have a primary blood type match with the donor, are geographically closer, and have shorter wait times.

[8]Donor-patient pairs that are of primary blood type match include all four identical blood type combinations (O-O, A-A, B-B, AB-AB), as well as specific cross-type matches: O donor with B patient, A donor with AB patient, and B donor with AB patient. Pairs that are of secondary blood type match consist of O donor with a patient with blood type A or AB. All offers in our sample are either primary or secondary blood type matches.

These relationships, except for wait time, are as expected, since they all contribute to a lower sequence number, thereby leading to a higher priority for the patient in any match, where the probability of acceptance is usually higher.

Our descriptive results suggest that patients treat the two stages of decision-making differently. Initial decisions are made quickly and potentially influenced by factors such as the immediate availability of clinicians or the patient's physical conditions for transplantation. In contrast, final decisions require more deliberation on the quality of the match and thus take more time. Together, these observations motivate us to formulate these two decision stages separately in our model.

3 Reduced-form Evidence of Social Learning

In this section, we present reduced-form evidence of social learning in lung transplant acceptance decisions: even after controlling for all variables that determine a patient's sequence number, we show that a patient's acceptance probability decreases with their sequence number. That is, patients with a larger sequence number, who are offered an organ later in a match run, are less likely to accept compared to those with a smaller sequence number. In particular, we find that (i) the four key variables–patient LAS, primary ABO, distance, and wait time–affect the sequence number in a manner consistent with the description of the OPTN priority assignment; (ii) sequence number has no effect on provisional acceptance decisions once we condition on these key variables and other controls; and (iii) sequence number exhibits a significant negative effect on final acceptance decisions, even after conditioning on all controls. These findings suggest that patients treat provisional and final acceptance as separate decisions. All else being equal, receiving an offer at a higher sequence number signals bad news about the offer, resulting in lower acceptance probabilities.

Sequence Number OPTN policy describes how sequence number is computed, as described in Section 2. Here we present an alternative statistical model of sequence number, looking at how the sequence number of a patient relates to her own charac-

teristics, the donor characteristics, as well as the compatibility between the matched donor-patient pair. The importance of this model lies in its role as a control in the main reduced-form evidence supporting social learning. Specifically, we estimate the following model:

$$s_{i,j} = \mathbf{x}_j'\boldsymbol{\beta} + \mathbf{x}_i'\boldsymbol{\gamma} + \mathbf{c}_{i,j}'\boldsymbol{\alpha} + \varepsilon_i, \tag{3.1}$$

where $s_{i,j}$ is the sequence number of patient j when matched with donor i; $\mathbf{x}_j$ is a vector of patient characteristics that include their LAS, total wait time, BMI, as well as their health condition indicators such as diabetic status and history of previous lung transplant; $\mathbf{x}_i$ is a vector of donor characteristics including age, weight, height, P/F ratio and dummies for whether the donor has a history of IV drug or heavy alcohol use. We also include a vector $\mathbf{c}_{i,j}$ of donor-patient compatibility indicators such as whether donor i and patient j are of a primary blood type match, their differences in age, weight, height, and the distance between the donor and the patient.[9]

The first column of Table 4 reports the results of regression (3.1). The coefficients of interest are those that correspond to the four factors used in the OPTN system for assigning sequence numbers, namely patient LAS, primary ABO, wait time, and distance. All four coefficients are statistically significant, and exhibit the anticipated signs: patients with higher LAS, a better blood type match, longer wait times, and closer geographical proximity to the donor are more likely to have lower sequence numbers; thus receiving higher priority in the queue. In terms of magnitude, the coefficients for primary ABO and distance are much larger than those for LAS and wait time, suggesting their first-order effects in determining priority levels, which is also consistent with the OPTN policy.

Acceptance Decisions Next, we examine the effect of sequence number on patient acceptance decisions, controlling for all previous variables, including those used in

[9]More specifically, the geographical distance between the donor and the patient is characterized by "Zones," where Zone A is for distances between 0-250 nautical miles (NM); Zone B is between 250-500 NM; Zone C is between 500-1000 NM; Zone D is between 1000-1500 NM; and Zone E is between 1500-2500 NM.

Table 4: Reduced-Form Analysis: Sequence Number, Provisional Acceptance and Final Acceptance

	(1) Sequence number	(2) Provisional yes	(3) Final yes
Sequence number		-0.001	-0.054***
		(0.001)	(0.010)
Patient LAS	-0.596***	-0.018***	0.041***
	(0.016)	(0.003)	(0.004)
Primary ABO	-39.938***	0.020	0.759*
	(0.382)	(0.077)	(0.390)
Wait time (month)	-0.064***	0.005*	-0.022**
	(0.013)	(0.003)	(0.010)
Distance (reference group–Zone A: 0-250 NM)			
Zone B: 250-500 NM	40.807***	0.922***	-0.888***
	(0.313)	(0.078)	(0.156)
Zone C: 500-1000 NM	31.356***	-1.739***	-1.810*
	(1.770)	(0.094)	(1.025)
Zone D: 1000-1500 NM	23.893***	-1.180***	-
	(2.892)	(0.234)	-
Zone E: 1500-2500 NM	18.203***	-2.485***	-
	(2.419)	(0.176)	-
Patient characteristics	X	X	X
Donor characteristics	X	X	X
Patient-donor differences	X	X	X
Provisional yes	-	-	X
Observations	32,697	32,697	14,587
R-squared	0.490	0.125	0.246
R-squared w/o $s_{i,j}$	-	0.125	0.207

*N*ote: For all regressions, we control for a wide range of variables about the patient, donor and donor-patient characteristics which we do not report due to space constraints. Note that since we restrict the sample to offers that are provisionally accepted in Column (3), the variable "sequence number" is reordered based on the initial sequence number of these patients. Robust standard errors are reported in parentheses. ***Significant at 1% level, **Significant at 5% level, *Significant at 10% level. Columns (2)-(3) show results from a logit regression, so pseudo R-squared is reported.

OPTN priority assignment to determine sequence number. To do so, we run a Logit model as follows:

$$y_i = \mathbb{1}\{\beta_s s_{i,j} + \mathbf{x}_j'\boldsymbol{\beta} + \mathbf{x}_i'\boldsymbol{\gamma} + \mathbf{c}_{i,j}'\boldsymbol{\alpha} + \varepsilon_i \geq 0\}, \tag{3.2}$$

where, as before, $s_{i,j}$ is the sequence number of patient j when matched with donor i, and the vectors $\mathbf{x}_j$, $\mathbf{x}_i$ and $\mathbf{c}_{i,j}$ represent the observed characteristics of patients, donors, and donor-patient compatibility. In the specifications below, y_i is a dummy variable indicating whether the donor is provisionally accepted or finally accepted.

Provisional Yes The second column of Table 4 presents the result of regression (3.2) when y_i represents the provisional acceptance decision. As reported in the first row, after controlling for a rich set of observables, the negative coefficient β_s on sequence number is not statistically significant, which suggests that sequence number itself does not affect patients' provisional decisions above and beyond the information contained in the controls. The coefficient on patient LAS is negative and statistically significant, meaning that patients with a higher LAS are less likely to accept provisionally. We also observe a weak positive coefficient for wait time, suggesting that patients who have waited longer are slightly more likely to accept provisionally. Finally, compared to patients in Zone A (within 0-250 NM of the donor), only patients in Zone B (within 250-500 NM of the donor) are more likely to provisionally accept the offer, while patients in more distant zones are less likely to do so. This is consistent with the notion that the provisional decision reflects a basic unavailability for a potential transplant operation.

Final Decisions The third column of Table 4 presents the result of regression (3.2) when y_i represents the final acceptance decision. As mentioned before, our data on patient final acceptance is truncated at the first acceptance.[10] The coefficient of interest here is the one associated with the sequence number. Unlike provisional

[10]Recall that we exclude all offers following the first acceptance, resulting in 15,773 offers where both provisional and final decisions are recorded.

acceptance decisions, this coefficient is negative ($\hat{\beta}_s = -0.054$) and statistically significant at the 1 percent level, even after controlling for all factors used to determine sequence numbers, as well as all other observables.

In other words, all else equal, having a higher sequence number alone decreases the final acceptance probability of a patient who provisionally accepted the offer. This significant negative correlation between sequence number and final acceptance probability, which is absent in provisional acceptance decisions, suggests that patients with higher sequence numbers make negative inferences about the donor quality once they condition on the event that they will be matched to the donor. Since the donor can only be accepted by the next patient in line if it is rejected by all patients with higher priority, only information about rejections is transmitted as patients move further down the priority queue. The negative coefficient on sequence number implies that previous rejections are interpreted as "bad news" about the donor quality, thereby reducing the likelihood that the donor will finally be accepted.

We note that all coefficients for patient LAS, wait time, and distance are statistically significant in final acceptance decisions, but exhibit the opposite sign compared to the coefficients observed in provisional acceptance decisions. For example, the coefficient for LAS has changed from negative in column (2) to positive in column (3) of Table 4. This change suggests that, contrary to provisional decisions, patients with a higher LAS—indicating a more urgent status and better survival probability—are more inclined to accept the offer in final decisions, despite their lower likelihood of provisional acceptance. Similarly, a negative coefficient for Zone B (compared to Zone A) implies that patients residing further away from the donor are less likely to accept the offer in the final decision, despite their higher likelihood of provisional acceptance.

4 Structural Model

We propose a model that combines a discrete-choice specification with sequential decision-making and the effects of social learning. In modeling social learning, we

build on the classic theoretical models of sequential decision-making (Banerjee, 1992; Bikhchandani et al., 1992), adapted for the purposes of structural estimation. Our model incorporates four key components that influence patient acceptance decisions: (i) unobserved lung quality that is common to all patients; (ii) observed donor and patient individual characteristics; (iii) observed donor-patient match quality characteristics; (iv) unobserved heterogeneity. The first component, an unobserved common-value component, is the object of social learning. Patients update negatively about the value of the common quality component from past rejections in a match run. The second and third components account for observed heterogeneity in patients, donors, and donor-patient pairs.

Donor quality Each donor $i \in N$ is described by a quality $\theta_i \in \{-1, 1\}$, identically and independently distributed across donors. We denote by $p \in (0, 1)$ the prior probability that the donor quality is $\theta = 1$. Throughout, we refer to $\theta_i = 1$ as donor i being of high quality and $\theta_i = -1$ as donor i being of low quality. The variable θ_i captures any additional unobserved common aspects of the donor's lung quality, in addition to the observed donor characteristics. Each θ_i is unobserved to both patients and the econometrician, and we assume that it is fixed throughout the allocation process. This assumption is supported by our empirical observation that nearly all offers–about 98% in our dataset–receive a final decision before the organ is physically extracted from the donor. Thus, in our data, concerns about organ quality deterioration during the allocation process are minimal.

Allocation When donor i becomes available, a match run is formed involving the patients compatible with donor i, which we denote N_i. These patients are first asked for their provisional decision, and those who say provisional yes are then asked for their final decision. Provisional decisions are made on the basis of idiosyncratic factors such as the transplant center having a team available for transplantation and/or the patient being ready for transplantation. Instead, final acceptance decisions are made based on the information available to the patient regarding organ quality.

Formally, we assume each patient in N_i is endowed with two signals, denoted by

$(\omega^1_{i,j}, \omega^2_{i,j})$, each taking values in $\{-1, 1\}$. The first signal indicates whether patient j is ready for transplant ($\omega^1_{i,j} = 1$) or not ($\omega^1_{i,j} = -1$); its distributed independently from the donor quality, and independently and identically across patients. We denote by μ the probability that $\omega^1_{i,j} = 1$. We let P_i denote the patients who are ready for transplant, i.e., patients with $\omega_{i,j} = 1$ are those who say provisional yes. Conditional on being ready for transplant, patient j's second signal, $\omega^2_{i,j}$ is informative about the organ quality. We denote by $\alpha \in (1/2, 1)$ the signal's precision, that is, the probability that $\omega^2_{i,j} = \theta_i$ conditional on the organ quality being θ_i. We assume that conditional on donor quality, the second signal is independently and identically distributed across patients.

Final acceptance When offered the organ, patient $j \in P_i$ decides whether to accept ($d_{i,j} = A$) or reject ($d_{i,j} = R$) the organ.[11] Patient j's utility from rejecting the organ is given by $\epsilon^R_{i,j}$. Instead, if patient i accepts a donor of quality θ_i, they receive payoff:

$$u^A(\tilde{\mathbf{x}}_{i,j}, \theta_i, \epsilon^A_{i,j}) = \tilde{\mathbf{x}}'_{i,j}\tilde{\boldsymbol{\beta}} + \gamma\theta_i + \epsilon^A_{i,j}, \tag{4.1}$$

where $\tilde{\mathbf{x}}'_{i,j} := (1, \mathbf{x}'_i, \mathbf{x}'_j, \mathbf{c}'_{i,j})$ consists of three observable components, including a constant that normalized to one: $\mathbf{x}_i$ is a vector of donor i's characteristics; $\mathbf{x}_j$ is a vector of patient j's characteristics; and $\mathbf{c}_{i,j}$ is a vector of characteristics capturing the compatibility between donor i and patient j (just as in the reduced-form model in Section 3). The variables in $\tilde{\mathbf{x}}_{i,j}$ are observable to both patient j and the econometrician. The parameter $\tilde{\boldsymbol{\beta}}$ is a vector consisting of the utility weights associated with $\tilde{\mathbf{x}}_{i,j}$ and γ is the utility weight associated with θ_i. Finally, $\epsilon^A_{i,j}$ denotes the idiosyncratic utility shock that patient j receives when accepting donor i's lung. We assume that $\epsilon^R_{i,j}$ and $\epsilon^A_{i,j}$ are distributed i.i.d. across patients and donors, each following a Type I extreme value distribution. Consequently, the difference between them, denoted by $\delta_{i,j} := \epsilon^A_{i,j} - \epsilon^R_{i,j}$, follows a logistic distribution centered at zero with a scale parameter of one.

At the moment of making their decision, patient j knows the idiosyncratic shocks–

[11] The decisions of patients in $N_i \setminus P_i$ are automatically set to R.

which remain unobserved to other patients and the econometrician–their private signal, $\omega^2_{i,j}$, and their sequence number. Patient j accepts the donor if conditional on all the available information, the expected utility from accepting is larger that $\epsilon^R_{i,j}$, that is,

$$\mathbb{E}[u^A(\tilde{\mathbf{x}}_{i,j}, \theta_i, \epsilon^A_{i,j}) - u^R(\epsilon^R_{i,j})|\omega^2_{i,j}, s_{i,j}] \geq 0.$$

From the perspective of patient j, the only unknown variable in the above expression is θ_i. Thus, the above inequality can be written as

$$\delta_{i,j} \geq -v(\tilde{\mathbf{x}}_{i,j}, \omega^2_{i,j}, s_{i,j}), \tag{4.2}$$

where

$$v(\tilde{\mathbf{x}}_{i,j}, \omega^2_{i,j}, s_{i,j}) = \tilde{\mathbf{x}}'_{i,j}\tilde{\boldsymbol{\beta}} + \gamma\mathbb{E}[\theta_i|\omega^2_{i,j}, s_{i,j}] \tag{4.3}$$

is the final acceptance threshold of patient j for donor i. This threshold is determined by a combination of donor and patient characteristics (including characteristics describing their mutual compatibility), their private information, and the information conveyed through the allocation process through their sequence numbers. The information transmitted through the patient's sequence number is crucial: a patient evaluates θ_i conditional on the event that all patients with a lower sequence number have turned down the donor.

The Inference Problem Under provisional yes, when it is patient j's turn to decide, they infer that all higher-priority patients who provisionally accepted the offer must have eventually rejected it. According to (4.2), these final rejections result from a combination of private information and patient characteristics. For model tractability, we assume that each patient $j \in P_i$ knows the characteristics of other patients in P_i with lower sequence numbers.[12] Hence, given these characteristics, patient j can infer the extent to which past rejections are due to higher-priority patients' private signals. Recall that the expected quality of θ_i conditioned on $\omega^2_{i,j}$ and $s_{i,j}$ is $\mathbb{E}[\theta_i|\omega^2_{i,j}, s_{i,j}]$. As $s_{i,j}$ increases, the number of patients rejecting donor i

[12]This assumption also helps ease the computational complexity in estimating our structural model.

also weakly increases, so one would expect $\mathbb{E}[\theta_i|\omega_{i,j}^2, s_{i,j}]$ to be weakly decreasing in $s_{i,j}$. Intuitively, as the sequence number increases, the number of past rejections also increases, leading to the accumulation of negative information. This, in turn, reduces the probability of acceptance, resulting in a herd behavior of rejecting the donor.

We view the extent of negative information transmitted through the sequence number as reflecting the level of social learning in the current OPTN system. In Section 6, we report counterfactual experiments to evaluate the impact of social learning by varying the amount of information contained in $s_{i,j}$. Specifically, we will examine scenarios where $s_{i,j}$ ranges from containing no information to containing the maximum amount of private information. Furthermore, we will explore how different allocation algorithms affect the informational content of $s_{i,j}$ and interact with social learning.

5 Estimation

Motivated by our empirical evidence (see Section 3), we estimate our model using maximum likelihood in two steps. First, we estimate the probability of provisional acceptance μ directly from the data. Second, we estimate the remaining parameters by a direct search method, the Nelder-Mead simplex algorithm. The parameters that need to be identified in the second step include: the precision of the second private signals, α; the prior on the unobserved quality of the lung being high, p; the coefficient for the expected lung quality, γ; and the vector of coefficients for donor, patient and donor-patient match quality characteristics, $\tilde{\boldsymbol{\beta}}$. Together, conditional on provisional acceptance–which occurs with probability μ–the vector of parameters $\Phi := (\alpha, p, \gamma, \tilde{\boldsymbol{\beta}})$ collectively governs the final decision process.

For each donor i and each patient $j \in N_i$ matched with donor i, we observe the following set of variables in the data:

$$\Lambda_{i,j} = (\mathbb{1}_{j \in P_i}, \mathbb{1}_{d_{i,j}=A}, s_{i,j}, \mathbf{x}_i, \mathbf{x}_j, \mathbf{c}_{i,j})$$

where $\mathbb{1}_{j\in P_i}$ is an indicator for whether patient j provisionally accepts donor i; $\mathbb{1}_{d_{i,j}=A}$ is an indicator for whether patient j finally accepts donor i; $s_{i,j}$ is the sequence number assigned to patient j for donor i; $\mathbf{x}_i$ is a vector of donor i's characteristics; $\mathbf{x}_j$ is a vector of patient j's characteristics; and $\mathbf{c}_{i,j}$ is a vector of donor-patient compatibility characteristics. Let $\Lambda_i = (\Lambda_{i,j})_{j\in N_i}$ denote the data observed for donor i, and $\Lambda = (\Lambda_i)_{i\in N}$ denote the data observed for all donors.

5.1 The Likelihood Function

For each donor $i \in N$, the likelihood of observing $|P_i|$ provisional acceptances among $|N_i|$ matched patients is $\mu^{|N_i|}(1-\mu)^{|N_i\setminus P_i|}$. Thus, the maximum likelihood estimator of μ in the first step is given by $\frac{\sum_{i\in N}|P_i|}{\sum_{i\in N}|N_i|}$.

Now, in the second step, for a given vector of parameters Φ, we focus on the probability that donor i is accepted by some patient $j \in P_i$. According to (4.2), the probability of patient j accepting donor i is

$$\mathbb{P}[d_{i,j} = A|\tilde{\mathbf{x}}_{i,j}, \omega^2_{i,j}, s_{i,j}; \Phi] = 1 - F(-v(\tilde{\mathbf{x}}_{i,j}, \omega^2_{i,j}, s_{i,j}; \Phi)),$$

where F is the cumulative distribution function of the standard logistic distribution. The complementary probability, $F(-v(\tilde{\mathbf{x}}_{i,j}, \omega^2_{i,j}, s_{i,j}; \Phi))$ is thus the probability of patient $j \in P_i$ rejecting donor i. Denote the patient in P_i with the lowest sequence number who finally accepts donor i by

$$j^*_i := \arg\min_{s_{i,j}}\{j \in P_i : d_{i,j} = A\}.$$

Since individual shocks are i.i.d., the likelihood that donor i is finally accepted by patient $j^*_i \in P_i$ given Φ is given by the following equation:

$$\mathscr{L}^A_i(\Phi; \Lambda_i) = \big(1 - F(-v(\tilde{\mathbf{x}}_{i,j^*_i}, \omega^2_{i,j^*_i}, s^i_{j^*_i}; \Phi))\big) \cdot \prod_{j\in P_i: s_{i,j}<s^i_{j^*_i}} F(-v(\tilde{\mathbf{x}}_{i,j}, \omega^2_{i,j}, s_{i,j}; \Phi)),$$

Note that the first term in the above product is the probability that donor i is

accepted by patient j_i^* *conditional* on being rejected by all patients with provisional acceptance and a lower sequence number than $s_{i,j}$. The second term is the probability of this conditional event, i.e., that donor i is eventually rejected by these provisionally accepted patients with lower sequence numbers. In contrast, the probability of donor i being rejected by all patients in P_i is

$$\mathscr{L}_i^R(\Phi;\Lambda_i) := \prod_{j\in P_i} F(-v(\tilde{\mathbf{x}}_{i,j},\omega_{i,j}^2,s_{i,j};\Phi)).$$

Let $S \subseteq N$ denote the set of donors that are successfully allocated, i.e., it has been accepted by at least one patient in P_i. The log-likelihood function of the observed donor allocations is

$$\mathscr{LL}(\Phi;\Lambda) := \sum_{i\in S} \log \mathscr{L}_i^A(\Phi;\Lambda_i) + \sum_{i\in S^c} \log \mathscr{L}_i^R(\Phi;\Lambda_i). \tag{5.1}$$

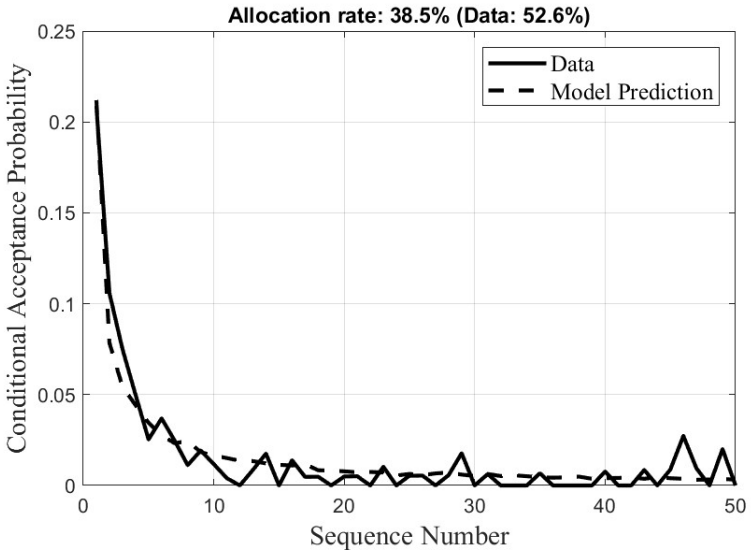
The first term in the above equation represents the log-likelihood of all donors in S being accepted by patient j_i^*, while the second term represents the log-likelihood of all donors in S^c never being accepted by any matched patients.

5.2 Estimation Results

Table 5 reports our coefficient estimates and the corresponding standard deviations, for the parameters in Equation (5.1). Recall that α is the precision of the patients' private information regarding the (unobserved) quality of the donor's lungs, while p represents patients' prior belief that the quality is high. Our estimates of α and p suggest that, although patients generally hold a pessimistic view of the unobserved quality–with a prior belief of only about 38 percent that the quality is high– the private information they receive is highly informative, with an 85% likelihood of revealing the true quality.

The estimated coefficients for the patient, donor, and donor-patient compatibility characteristics $\boldsymbol{\beta}$ are all of the expected signs, have reasonable magnitudes, and are

Figure 5.1: Model Fit–Conditional Acceptance Probability by Sequence Number



consistent with our reduced-form estimates. For example, the patients' utility of accepting a donor increases with higher LAS or a primary blood type match. Meanwhile, it decreases as patients wait longer or as the geographical distance between the patient and the donor increases. Note that the estimated coefficient on the unobserved quality γ is about five times higher than that of the primary blood type match, underscoring the importance of unobserved donor quality in patient acceptance decisions.

To examine the fit of the model, Figure 5.1 plots the predicted acceptance probabilities based on the estimates provided in Table 5. We then utilize these predicted probabilities to simulate acceptance decisions. Figure 5.1 illustrates the comparison between the average predicted acceptance probability and the observed average acceptance probability by sequence number for the first 50 priority positions. Note that these acceptance probabilities are conditional probabilities. For example, if the acceptance probability at sequence number 5 is 0.1, it means that the probability of accepting the offer when the patient is fifth in line is about 10% conditional on it being rejected by all patients with higher priority.

Table 5: Estimation Results

Parameter	Estimate	S.D.
α	0.850	0.039
p	0.383	0.046
γ	4.934	0.005
$\tilde{\boldsymbol{\beta}}$		
Patient Characteristics		
LAS	0.033	0.001
Wait time (month)	-0.024	0.002
BMI	0.145	0.001
Female	-0.053	0.005
Diabetic	0.130	0.244
Previous lung transplant	0.439	0.413
Donor Characteristics		
P/F ratio	0.001	0.000
Age (year)	0.005	0.001
Weight (kg)	-0.070	0.001
Height (cm)	0.071	0.001
History of IV drug use	-0.014	0.344
History of heavy alcohol consumption	-0.034	0.040
Increasing risk status	-0.149	0.111
Donor-Patient Pair Characteristics		
Primary ABO	0.775	0.006
Distance (per NM)	-0.003	0.001
Age difference	0.027	0.001
Height difference	0.076	0.002
Weight difference	-0.064	0.001

*N*ote: This table shows the estimates for our structural model. The standard deviations are obtained through the BHHH estimator.

As evident from Figure 5.1, the model generally fits well with the observed pattern of acceptance, particularly in capturing the sharp decline in the acceptance rate at the initial sequence numbers. For example, the average acceptance rate in the data is approximately 21 percent at sequence number 1, decreasing by half to around 11 percent at sequence number 2–closely resembling our model predictions. The

figure also depicts the actual allocation rate and the simulated allocation rate using these predicted (conditional) acceptance probabilities, which are about 53% and 39%, respectively. The underestimation of the simulated overall allocation rate is expected since the simulation only uses the predicted acceptance probability on the observed match run length, which is truncated. We anticipate that the predicted allocation rate will align more closely with the actual allocation rate when the simulation is based on the full length of the match (see more discussions in Section 6).

6 Counterfactual Experiment

Using our structural estimates, we conduct a series of counterfactual experiments to examine the impact of social learning on allocation outcomes, and its interaction with different allocation policies. Our baseline simulations use the OPTN priority policy as described in Section 2. We then consider two sets of counterfactual experiments: in the first, we change the allocation policy to evaluate different policy proposals; in the second set, we vary the information treatments to better understand the role of social learning.

6.1 Alternative Priority Policies

In this section, we compare two alternative priority policies to the one currently used by the OPTN. The two policies differ in how they order patients in a match run. The first policy, which we call *greedy priority*, offers the organ to patients in descending order of their estimated utility from accepting the organ. The second policy, which we call *reverse greedy priority*, offers the organ to patients in ascending order of their estimated utility from accepting the organ.

In the absence of social learning, the greedy priority policy would maximize the probability that an organ is allocated. Thus, if one only cares about maximizing the overall allocation rate (or allocating the organ to those who value it the most) and ignores social learning, greedy priority seems to be the best policy. This policy, however, is also the "worst" policy if the goal is to minimize the negative inference from

social learning: when an organ is rejected by patients who are most likely to accept it, the inference made by subsequent patients is the most negative; consequently, they are more likely to reject. As a result, when interacting with social learning, whether the net effect of greedy policy on improving the overall allocation rate is positive remains unclear.

Following a similar logic, the reverse greedy priority will minimize the impact of social learning. The negative inference that results from past rejections is the least informative, since these are the rejections of patients who are more likely to reject the organ in the first place. In the rest of this section, we present empirical results examining the effects of both policies, and comparing them with the current OPTN priority policy.

We note first that our simulation of the current policy closely reflects the summary statistics that we find in the data. The allocation rate is about 53%, with an average accepted sequence number of 8.5, meaning that the average position of the patient who accepts, in the sequence of patients who are offered the organ, is 8.5. These values are more consistent with our empirical observations compared to the predictions from our model fit.

Panel (A) of Table 6 presents two efficiency outcomes: the average allocation rate and the average accepted sequence number. Interestingly, we find that the reverse greedy policy increases the current OPTN allocation rate by 2 percentage points, whereas the greedy priority policy decreases it by about 1 percentage point. This finding indicates that the effect of social learning drives the difference when comparing between these alternative priority policies. Although the reverse greedy policy performs the best in improving the overall allocation rate, this improvement comes at the expense of a slower allocation process. For instance, the average accepted sequence number is around 15 under the reverse greedy policy, which is about three times that under the greedy priority policy. These results suggest a trade-off between improving the allocation rate and expediting the allocation process when evaluating different priority policies.

Table 6: Counterfactual 1: Different Priority Policies with Social Learning

Priority Policy	(1) OPTN (Current)	(2) Greedy	(3) Reverse Greedy
Panel (A): Efficiency Outcomes			
Allocation Rate (%)	47.22	46.26	49.33
Accepted Sequence Number	6.91	5.13	15.49
Panel (B): Welfare Outcomes			
Acceptance Utility	2.52	2.75	1.48
Total Acceptance Utility	619.72	662.27	381.03
Acceptance Rate of High Quality (%)	32.63	32.44	32.63
Rejection Rate of Low Quality (%)	44.53	45.30	42.42
Panel (C): Accepted Patient Characteristics			
LAS	59.12	64.24	38.15
Wait time (month)	5.35	3.03	8.27
Primary ABO	1.61	1.60	1.54
Distance (NM)	99.76	155.73	339.75
Provisional Yes	✓	✓	✓
Available Information	one private signal and past rejections		

*N*ote: The variable "Acceptance Utility" represents the average utility of patients who accept the offer. The variable "Total Acceptance Utility" is the sum of all acceptance utilities. In cases where multiple patients accept the same donor, the donor is allocated to the patient with the highest priority according to the corresponding priority policy.

Panel (B) of Table 6 reports welfare outcomes under different priority policies. Patient welfare is represented by the average utility of patients who finally receive the organ.[13] We find that the greedy priority policy achieves the highest average acceptance utility, followed by the current OPTN policy, and then the reverse greedy priority policy. Despite having the lowest allocation rate, the greedy priority policy also achieves the highest total acceptance utility. Therefore, from the perspective of a

[13]This is also known as the ex-ante maximum expected utility. As is well known, the ex-ante maximum expected utility of patient j before the realization of her idiosyncratic utility shocks $(\varepsilon^A_{i,j}, \varepsilon^R_{i,j})$ is

$$\mathbb{E}[\max\{u^A(\tilde{\mathbf{x}}_{i,j}, \theta_i, \varepsilon^A_{i,j}), u^R(\varepsilon^R_{i,j})\}] = c + \log\left(1 + \exp\left(\tilde{\mathbf{x}}'_{i,j}\tilde{\boldsymbol{\beta}} + \gamma\theta_i\right)\right),$$

where c is Euler's constant, $c \approx 0.577$.

social planner aiming to maximize total welfare and expedite the allocation process, the greedy priority policy performs the best among these three priority policies.

To further explore why greedy priority achieves the highest average acceptance utility, we also report the rate of "correct" decisions: accepting high-quality donors and rejecting low-quality ones. To clarify the terminology here, note that it may still be optimal for the patient to reject a high-quality organ, or accept a low-quality organ, since it depends on the donor-patient compatibility for that individual. The donor quality is obviously a key part of the patient decision, but not the only one.[14] We find that, while the acceptance rate of high-quality donors is consistently around 32% across all three policies, the rejection rate of low-quality donors varies from 42% to 45%, with the greedy priority policy achieving the highest rejection rate. Intuitively, as patients progress further down the queue, they become more pessimistic about the organ's quality, regardless of its actual unobserved quality. Consequently, the likelihood of rejecting low-quality donors increases, and for these patients to accept the offer, the utility of accepting must be higher. This could partially contribute to the improvement in overall patient welfare in terms of acceptance utility under the greedy priority policy.

Panel (C) of Table 6 reports the average characteristics of patients who have been allocated an organ. Comparing these characteristics across all three policies, we see that average patient LAS is highest under the greedy priority policy, followed by the current policy, and then the reverse greedy policy. Meanwhile, the current policy has the shortest average distance between the donor and the patient, followed by the greedy and then the reverse greedy policy. This suggests that compared to the current policy, the greedy priority policy places more weight on LAS and less on distance when assigning priority. Therefore, patients with higher LAS, even if they are farther away from the donor (or have only waited a relatively short period), could still benefit from the greedy priority if they are deemed to be more likely to accept

[14]Note that if all patients were homogeneous and followed their private signals, our structural estimates suggest that the acceptance rate for high-quality donors would be around 33%, and the rejection rate for low-quality donors would be around 52%.

the offer.

6.2 Alternative Information Treatments

In this section, we focus on the current OPTN priority policy and vary the amount of information conveyed through the sequence number to understand the effect of social learning. As a baseline, we simulate the current OPTN policy while accounting for social learning. We then consider two informational treatments. The first treatment, referred to as *no social learning*, completely shuts down the (negative) inference from social learning and thus aggregates no private information: each patient receives one private signal and acts as if they are the first in the priority queue. The second treatment, referred to as *information sharing*, aggregates the most amount of private information: each patient receives one private signal and also observes the private signals received by all higher-priority patients.

Table 7 reports the outcomes of these informational treatments. We find that in the no social learning treatment, the overall allocation rate almost doubles, increasing from 47% to about 84%, compared to the baseline where patients learn from past rejections. This significant increase is driven by the fact that only negative information is accumulated through the priority queue, so eliminating it boosts the acceptance rate. However, this increased allocation rate is likely to be an over-acceptance, as most of these additional acceptances come from "incorrect" decisions. For example, in this scenario, the rejection rate of low-quality donors–only about 15%–is much lower than that under the presence of social learning or information sharing. Indeed, average acceptance utility is the lowest in the treatment with no social learning, suggesting a poorer level of match quality due to the lack of information.

Furthermore, we find that having more information, whether negative or not, accelerates patients' decision-making. This finding is reflected in the average accepted sequence number: it is about 2.5 when patients share their private information, increases to 6.9 when this information remains private but patients can infer additional information from past rejections, and further increases to 7.3 when there is no

Table 7: Counterfactual 2: Different Information Treatments under the OPTN Priority policy

Information Treatment	(1) Social Learning	(2) No Social Learning	(3) Information Sharing
Panel (A): Efficiency Outcomes			
Allocation Rate (%)	47.22	83.69	48.94
Accepted Sequence Number	6.91	7.33	2.49
Panel (B): Welfare Outcomes			
Acceptance Utility	2.52	1.83	2.71
Total Acceptance Utility	619.72	798.60	692.09
Acceptance Rate of High Quality (%)	32.63	39.92	40.31
Rejection Rate of Low Quality (%)	44.53	15.36	50.48
Panel (C): Accepted Patient Characteristics			
LAS	59.12	52.63	59.13
Wait Time (Month)	5.35	7.02	6.49
Primary ABO	1.61	1.57	1.61
Distance (NM)	99.76	94.82	77.39
Provisional Yes	✓	✓	✓
Priority policy	OPTN Priority (Current)		

additional information beyond their own private signals.

7 Discussion

8 Conclusion

References

Agarwal, N., I. Ashlagi, M. A. Rees, P. Somaini, and D. Waldinger (2021): "Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys," *Econometrica*, 89, 37–76.

Agarwal, N., I. Ashlagi, P. Somaini, and D. Waldinger (2018): "Dynamic incentives in wait list mechanisms," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 108, 341–347.

Banerjee, A. V. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107, 797–817.

Benvenuto, L. J. and M. Aversa (2019): "Should We Accept This Offer? When Assessing Donor Lungs, Don't Rely on Others," *Annals of the American Thoracic Society*, 16, 304–305.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 992–1026.

Cipriani, M. and A. Guarino (2014): "Estimating a structural model of herd behavior in financial markets," *American Economic Review*, 104, 224–51.

Cohen, J., J. Shults, D. Goldberg, P. Abt, D. Sawinski, and P. Reese (2018a): "Kidney allograft offers: Predictors of turndown and the impact of late organ acceptance on allograft survival," *American Journal of Transplantation*, 18, 391–401.

——— (2018b): "Kidney transplant outcomes: position in the match-run does not seem to matter beyond other donor risk factors," *American Journal of Transplantation*, 18, 1577–1578.

De Mel, S., K. Munshi, S. Reiche, H. Sabourian, et al. (2020): "Herding in Quality Assessment: An Application to Organ Transplantation," Tech. rep., Faculty of Economics, University of Cambridge.

Gilroy, R., K. Cmunt, M. Kevin Myer, C. Sonnenday, L. Preczewski, T. Stevens, H. Aguiar, and L. Swanson (2019): "Disrupting the Status Quo," 2016 National Critical Issues Forum.

Goldberg, D. S., B. French, J. D. Lewis, F. I. Scott, R. Mamtani, R. Gilroy, S. D. Halpern, and P. L. Abt (2016): "Liver transplant center variability in accepting organ offers and its impact on patient survival," *Journal of hepatology*, 64, 843–851.

Hackmann, M., R. A. English, and K. W. Kizer (2022): "Improving Procurement, Acceptance, and Use of Deceased Donor Organs," in *Realizing the Promise of Equity in the Organ Transplantation System*, National Academies Press (US).

Harhay, M. O., R. Porcher, G. Thabut, M. J. Crowther, T. DiSanto, S. Rubin, Z. Penfil, Z. Bing, J. D. Christie, J. M. Diamond, et al. (2019): "Donor lung sequence number and survival after lung transplantation in the United States," *Annals of the American Thoracic Society*, 16, 313–320.

Muñoz-Rodriguez, E. (2021): "Incentive-Compatible Triaging in Deceased Donor Transplantation," *Available at SSRN 3805693*.

Newham, M. and R. Midjord (2018): "Herd behavior in FDA committees: A structural approach," DIW Berlin Discussion Paper No. 1744.

OPTN (2019): "Organ Procurement and Transplantation Network Policies," Tech. rep., OPTN.

Shi, P. and J. Yin (2022): "Eliminating Waste in Cadaveric Organ Allocation," *USC Marshall School of Business Research Paper Sponsored by iORB*.

Zhang, J. (2010): "The sound of silence: Observational learning in the US kidney market," *Marketing Science*, 29, 315–335.