

Improving Organ Procurement Operations

Hammaad Adam, Nikhil Agarwal, Marzyeh Ghassemi *

January 31, 2026

Abstract. We study how decisions made by organ procurement organizations (OPOs)—non-profits that coordinate organ recovery from deceased donors—affect the availability of organs for transplant in the United States. We develop a structural econometric model of a pivotal OPO decision: whether to approach a potential donor’s family to request authorization for donation. Our model conceptualizes this decision in two parts. The OPO first estimates the probabilities of two downstream outcomes: authorization (i.e., family consent) and transplant (i.e., whether the donated organs would be successfully transplanted). It then applies a cost-benefit decision rule that maps these estimates to an approach decision. Our model separately identifies the OPO’s beliefs (i.e., probability estimates), its preferences (i.e., costs / benefits), and the true probabilities of authorization and transplant. We apply our model to a dataset of 35,856 potential donors referred to four OPOs between 2016 and 2021. We find that OPOs missed a substantial number of donation opportunities, recovering organs from only 39% of transplantable donors. Of these missed opportunities, an estimated 16% resulted from conservative preferences, 12% from miscalibrated beliefs, and the remaining 72% from family declines. We conduct a detailed counterfactual policy evaluation to recommend actionable changes to OPO decisions. We find that drastically increasing approach rates would allow OPOs to recover organs from 43% more donors, creating over \$100 million in additional annual societal benefit.

Keywords: Organ transplantation, empirical operations management

*Adam: MIT; work done prior to joining Amazon (current affiliation). Agarwal: Department of Economics, MIT and NBER. Ghassemi: Department of Electrical Engineering and Computer Science & Institute for Medical Engineering and Science, MIT. We thank Carri Chan, Dean Eckles, Lester Mackey, and Rene Bermea for valuable comments and feedback. We are grateful to the leadership of six OPOs—Ginny McBride, Kevin Lee, Diane Brockmeier, Brad Adams, Patti Niles, Kelly Ranum, Janice Whaley, and Brett Gordon—for their pivotal role in enabling access to the data. We also thank data and business intelligence leaders at each OPO for sharing their expertise, including Tom Levanos, Emily Stahlschmidt, Lindsey Speir, Kasper Statz, Naimisha Chakravarthula, Rashida Dupuis, Aninda Dutta, Jennifer Grant, Matilin Rigsby, Jonathan Hewlett, Tyler Strout, Chad Trahan, Andre Romain, Max Prather, Chhavi Solanky, Rachael Wulf, Kevin Paturzo, and Matthew Allen. Finally, we acknowledge Jennifer Erickson, Greg Segal, Tom Pollard, and Ahram Han for their invaluable support and guidance. Support for this article was provided in part by the Robert Wood Johnson Foundation (RWJF). The views expressed here do not necessarily reflect the views of RWJF. This work was also supported, in part, by a National Science Foundation CAREER Award (2339381), an AI2050 Award Early Career Fellowship (G-25-68042), a Gordon & Betty Moore Foundation Award, a Sloan Research Fellow Award (FG-2025-24277), and funding from Takeda Development Center Americas, INC.

1 Introduction

Organ transplantation is a life-saving intervention for patients with advanced diseases. However, transplant-viable organs are a scarce public resource in the United States (US), as there is a severe shortage of deceased donors. In 2024, more than 70,000 patients were added to a national organ transplant waiting list, while fewer than 50,000 received a transplant.¹ As a result of this shortage, over 105,000 patients in the US are currently waiting for an organ transplant.¹

While significant empirical research has focused on the optimal allocation of donated organs to waitlisted patients (e.g., Su and Zenios (2005), Agarwal et al. (2021)), comparatively little attention has been paid to organ procurement, the process by which organs are recovered from deceased donors. In the US, this process is managed by 56 organ procurement organizations (OPOs), each of which is mandated by federal law to conduct procurement in a designated service area. Prior studies have documented wide variation and inefficiencies in OPO performance (Lopez et al. 2023, Johnson et al. 2023, Goldberg et al. 2013, DeRoos et al. 2021), with one estimate suggesting that over 28,000 transplant-viable organs go unrecovered annually (Rosenberg et al. 2020).

In this paper, we study the effect of a specific OPO decision on the availability of deceased donors. We focus on a pivotal step in the procurement process: family approach. When a ventilated patient in a hospital is referred to an OPO as a potential donor, the OPO must decide whether to approach their next-of-kin (NOK) to initiate discussions about donation. To maximize transplant opportunities with their available resources, OPOs aim to approach for potential donors who are most likely to lead to successful transplants. This approach decision represents the largest drop-off in the procurement pipeline: over 85% of referred potential donors are ruled out without an approach (Levan et al. 2022). If even a modest share of those ruled out were viable donors, the missed opportunity for transplantation—and associated social benefit—is significant.

Our work addresses three central questions. First, **how many transplantable donors were referred to an OPO but not procured?** This question captures the number of missed opportunities: potential donors who could have facilitated transplants, but whose families were either not approached or declined to provide consent. Second, **were OPO approach decisions based on accurate predictions of transplant outcomes?** The OPO’s decision is akin to binary classification, as it aims to distinguish between potential donors who would and would not lead to transplants

¹Based on OPTN data as of May 2, 2025

if approached. How accurate were OPOs at this classification task? And third, **what actionable changes to OPO decision-making can most increase the number of transplanted donors?**

To address these research questions, we develop a structural econometric model of the OPO’s decision to approach. We define a potential donor’s *transplant likelihood* as the probability that at least one of their organs would be successfully transplanted, and *authorization likelihood* as the probability that their family would consent to donation if approached. Our model conceptualizes the OPO’s approach decision in two parts. The OPO first forms *beliefs* (i.e., estimates) about these two likelihoods based on all available information. Then, the OPO applies a decision rule that uses these beliefs to determine whether to approach. This decision rule reflects the OPO’s valuation of the three possible outcomes of an approach: no authorization, authorization without a transplant, and a successful transplant. These relative valuations—which we collectively term the OPO’s *preferences*—can be shaped by its costs, regulatory incentives, and other factors.

Our model uses data on approach decisions, authorization outcomes, and transplant outcomes to separately identify the OPO’s beliefs, preferences, and the true likelihoods of transplant and authorization. A central challenge for identification is selection: transplant outcomes are only observed for potential donors who were approached. If the OPO’s decision to approach is influenced by a signal of transplant likelihood that is unobserved in the data, then naive models of transplant likelihood will be biased (Heckman 1979). We address this challenge using an instrument that has an empirically demonstrable effect on the OPO’s approach decision, but we argue is independent of transplant likelihood. This exogenous variation allows us to identify transplant likelihood from approach decisions and transplant outcomes, despite potential selection on unobservables.

Another key empirical challenge is that our model separates the OPO’s beliefs and preferences while allowing beliefs to be incorrect (i.e., differ from true transplant likelihood). Prior econometric models (e.g., Chan et al. (2022), Abaluck et al. (2016)) have addressed this challenge by assuming that the decision-maker has correct beliefs about a subset of parameters that determine the downstream outcome. We take a different approach, instead assuming that the OPO’s beliefs are calibrated to two aggregate, observable signals. Specifically, we assume that the OPO’s beliefs correctly explain the average transplant rate among authorized potential donors and the transplant rate among donors that it was indifferent about approaching (i.e., marginally-approached donors). Both of these metrics are observable to the OPO and directly tied to its financial and regulatory incentives; these calibration conditions are thus interpretable and justifiable. Separating incorrect

beliefs from preferences also requires two additional assumptions: the OPO's preferences (i.e., costs and rewards) do not vary by potential donor and the OPO has correct beliefs about authorization likelihood. We argue that these simplifications are supported by the OPO's regulatory incentives, the structure of the donation process, and their deep experience interacting with donor families.

We apply our model to a dataset that contains detailed clinical and process information for every potential donor referred to four OPOs (35,856 total) over a six-year period (2016–2021). This analysis directly addresses each of our three research questions. First, we find that the OPOs missed a significant number of transplant opportunities. If the OPOs had made an approach for every potential donor and obtained authorization for every approach, they would have recovered over 2.5 times as many transplantable donors. While this upper bound is not practically achievable—some families will decline to provide consent no matter how the OPO approaches—it does highlight the scale of the currently unrealized opportunity.

Second, we find that OPOs were able to predict transplant likelihood with high accuracy. OPOs were particularly accurate at ranking potential donors relative to each other: their beliefs demonstrated an area under the receiver operating characteristic curve (AUROC) of 0.77. However, the maximum achievable AUROC given the OPOs' observed information was 0.87, highlighting room for further improvement. Moreover, despite their ranking accuracy, we find that OPO beliefs underestimated true transplant probability for a large number of potential donors. In fact, if OPO beliefs aligned perfectly with true transplant likelihood, they would have made 21% more approaches, leading to 18% more transplanted donors. These findings highlight that beliefs and preferences both contributed to missed donation opportunities.

Finally, we use our structural model to evaluate a range of policy interventions that can improve OPO operations. The ability to support such counterfactual analyses is a key advantage of our structural approach; by simulating interventions on model parameters, we can assess their downstream effects on transplant outcomes. We evaluate three classes of policies: policies that increase the OPOs' willingness to approach, policies that increase the accuracy of OPO beliefs, and policies that improve the OPOs' ability to obtain authorization from donor families. Overall, we find that there is significant value in a simple policy: approach all potential donors without absolute contraindications. This policy would have led to an estimated 43% gain in transplanted donors, creating over 67,000 additional life years over our study period. Notably, the increase in OPO operating costs required by this policy would have been outweighed by its financial benefit; the increase in donors

would have saved taxpayers \$608 million in Medicare spending while costing only \$378 million. Developing financial and regulatory measures that incentivize OPOs to implement this policy and drastically increase approach rates should be a key policy focus going forward.

The rest of the paper proceeds as follows. In Section 2, we discuss related literature. Section 3 provides background on organ procurement and introduces our data and research questions. Section 4 develops our structural model, then Section 5 discusses the parameter estimates obtained from fitting this model to our dataset. Finally, Section 6 discusses key domain implications of our analysis and Section 7 presents our policy evaluation.

2 Related Work

Organ Allocation and Donation A large body of work in operations research, economics, and computer science has studied the distribution of donated organs to transplant recipients via waitlists and other mechanisms (Su and Zenios 2005, 2006, Bertsimas et al. 2013, Berrevoets et al. 2020, Papalexopoulos et al. 2024, Agarwal et al. 2021, Qin et al. 2021, Agarwal et al. 2025, Kong et al. 2010). This literature focuses on the demand side of organ transplantation, as it considers how already-donated organs should be allocated. In contrast, our work studies the supply-side of transplantation. Prior supply-side research has studied specific shocks that affect organ donation (e.g., opioid overdoses, motorcycle deaths) and their impact on waiting lists (Dickert-Conlin et al. 2024, 2019, 2011). Other work has proposed incentivizing donation by granting registered donors higher priority on transplant waiting lists (Li and Riyanto 2025, Dai et al. 2020). Here, we focus specifically on the effect of OPO decisions on the supply of transplantable organs.

OPO Performance Prior research has highlighted shortcomings in OPO performance and its impact on the supply of donated organs (Rosenberg et al. 2020, Lynch et al. 2022, Doby et al. 2022, 2021, Goldberg et al. 2020). This research has been crucial for government policy, as it led to the development and adoption of the new Final Rule that has tightened regulatory scrutiny on OPOs (Centers for Medicare & Medicaid Services 2021). However, these articles use aggregate metrics to assess OPO performance, and do not specifically focus on the decision to approach. While one existing article has analyzed the effect of increasing approaches on donation rates (Levan et al. 2022), it only adjusted for a few basic covariates and did not account for key empirical challenges

inherent to OPO decisions (i.e., selection, separating beliefs and preferences).

Medical Decision-Making Our model builds on an econometric literature that has studied medical decisions (e.g., ordering a diagnostic test) based on downstream outcomes (e.g., disease occurrence) (Chan et al. 2022, Chandra and Staiger 2020, Abaluck et al. 2016, Currie and MacLeod 2017). Like these studies, we frame the decision as binary classification, separate beliefs and preferences, and identify the model from censored outcomes while accounting for selection on unobservables. Our work extends these prior models by allowing for incorrect beliefs about both observables and unobservables, separating the decision threshold from beliefs about average outcome prevalence, and modeling multiple downstream outcomes (i.e., transplant and authorization).

Our work also contributes to an operations management literature that has used empirical models to study medical decision-making. Significant research has modeled decisions such as intensive care unit admission (Kim et al. 2015, Kc and Terwiesch 2012, Shmueli and Sprung 2005), clinician staffing (Yankovic and Green 2011, Hu et al. 2025), resource scheduling (Shen et al. 2021, Mandelbaum et al. 2020), and patient discharge (Chan et al. 2012, Freeman et al. 2021). Related work has also studied how environmental and cognitive factors (e.g., workload, congestion) affect these decisions and downstream health outcomes (Kim et al. 2020, Kc and Terwiesch 2009, Kim et al. 2024, Kc and Terwiesch 2011, 2012). Like these studies, we empirically model a high-stakes human decision, investigate the role of specific behavioral factors (i.e., beliefs and preferences), and identify how current decisions can be improved to generate better health outcomes.

Structural Estimation in Operations Management A growing body of work in healthcare operations management has used structural estimation to model the behavior of decision-makers and generate managerial insights. For example, Olivares et al. (2008) uses a structural model to identify how hospitals trade off the implicit costs of idle time and schedule overruns when reserving operating room capacity for cardiac surgery. Rath and Rajaram (2022) uses a choice model to study anesthesiologist scheduling and identify a planning policy that reduces overall costs. Bravo et al. (2025) uses structural demand estimation to determine how healthcare facilities can be strategically located to maximize service utilization and access to care. Lee and Zenios (2012) uses a principal-agent model to examine Medicare dialysis reimbursements and design a more optimal pay-for-compliance system. Outside of healthcare, structural estimation has been

used to provide actionable managerial insights in a wide range of settings, including air travel (Li et al. 2014), inventory management (Musalem et al. 2010), bike-share networks (He et al. 2021), employee attrition (Li et al. 2014), and fast-food operations (Allon et al. 2011). Our work complements this literature by using structural estimation to disentangle key behavioral drivers of OPO decisions and identify managerial changes that can increase the supply of organs for transplant.

3 Background, Setting, and Data

3.1 Organ Procurement

In the US, organ procurement from deceased donors is managed by 56 organ procurement organizations (OPOs). Each OPO holds an exclusive government contract to coordinate procurement within a specified donation service area (DSA) (Fleisher et al. 2023). OPOs are regulated under the National Organ Transplant Act of 1984 (United States Congress 1984), with regulatory oversight administered by the Centers for Medicare and Medicaid Services (CMS). OPOs are nonprofit organizations that are financially supported through reimbursements from CMS and fees charged to transplant centers for transplanted organs.

The OPO’s organ donation process proceeds in six stages (Fig. 1)—referral, evaluation, approach, authorization, procurement, and transplant—which we describe below.

- **Referral:** When a ventilated patient in a hospital meets pre-defined clinical triggers, the hospital is required to refer the patient to the local OPO as a potential donor. These clinical triggers are agreed upon between the OPO and the hospital (e.g., Glasgow Coma Scale ≤ 5).
- **Evaluation:** Upon referral, the OPO evaluates the potential donor’s suitability for organ donation. Potential donors with absolute contraindications (e.g., active cancer) are immediately ruled out. Otherwise, the potential donor is evaluated on-site by an OPO coordinator, who confers with hospital staff to collect medical history, test results, and other relevant data.
- **Approach:** If the potential donor is deemed a viable candidate, the OPO approaches their NOK to request their consent for donation. If the potential donor had provided consent during their lifetime (e.g., via a donor registry), the OPO communicates this directive to the NOK.




Figure 1. Stages of the organ procurement process overseen by OPOs. The procurement process has six distinct phases: Referral; Evaluation; Approach; Authorization; Procurement; and Transplant. NOK is Next of Kin; OPO is Organ Procurement Organization; OPTN is Organ Procurement and Transplantation Network.

- **Authorization:** The OPO formally documents authorization (i.e., consent) obtained via the donor registry or from the potential donor's NOK.
- **Procurement:** After authorization, the OPO manages the donor to maintain organ viability and offers organ(s) to potential recipients on the waitlist per the policies of the Organ Procurement and Transplantation Network (OPTN). If an organ is accepted by a recipient's transplant center, the OPO coordinates surgical procurement.
- **Transplant:** Finally, the OPO arranges transport of the organ(s) to the accepting transplant center(s) and transplantation proceeds.

Each of these stages presents a potential point of attrition. For example, a referral may arrive too late for timely intervention, staffing limitations may prevent the OPO from making approaches for viable donors, or ineffective communication during approach may lead families to decline consent. One study estimated that such inefficiencies result in more than 28,000 transplantable organs going unrecovered each year (Rosenberg et al. 2020). Notably, high-performing OPOs recover up to 78% more donors than their low-performing counterparts (Lynch et al. 2022), highlighting the potential impact of improving these processes.

3.2 Data

We study the organ procurement process using data from the electronic health record (EHR) systems of four OPOs, referred to as OPOs A, B, C, and D to preserve anonymity. Our dataset describes all potential donors referred to these OPOs between January 1, 2016 and December 31, 2021. The dataset includes detailed donor-level information, including demographics, clinical measures, and outcomes at each stage of the procurement process. This dataset was curated as part of a broader effort to support transparency in organ procurement through open data (Adam et al. 2026).

3.2.1 Exclusions

Our initial dataset consisted of 13,794 potential donors from OPO A, 28,878 from OPO B, 19,434 from OPO C, and 11,307 from OPO D. We applied two exclusion criteria to this dataset, which reflect our project’s goal of studying discretionary OPO decision-making. First, we excluded any potential donors that were immediately ruled-out upon referral (i.e., without on-site evaluation) due to absolute donation contraindications. These rule-outs were based on pre-specified criteria and thus did not involve a discretionary OPO decision. Finally, we excluded a small number of potential donors who survived their hospitalization, as they did not represent true candidates for deceased donation. We inferred both immediate rule-outs and survival status using logic-based rules applied to the OPO data that were validated by OPO staff. Our final sample included 8,899 potential donors from OPO A, 13,733 from OPO B, 6,477 from OPO C, and 6,747 from OPO D. Table 1 summarizes the outcomes, demographics, and clinical features of our final cohort.

3.2.2 Outcomes

Table 1 clearly demonstrates the attrition at each stage of the OPO process. Of the 35,856 referred potential donors in our dataset, only 11,313 had family approaches, 7,018 provided authorization, 5,490 were surgically procured, and 5,164 led to transplants. The largest drop occurs at the approach stage, where 68% of potential donors were not pursued.

Going forward, we focus on three outcomes: whether the potential donor’s family was approached, whether authorization was obtained, and whether at least one donated organ was transplanted. Here, we have combined procurement and transplant into a single outcome. Attrition at the procurement stage is often due to unforeseen circumstances, as the OPO typically does not coordinate surgical

procurement unless the donor's organ(s) have already been accepted by a transplant center. Since 94% of surgically-procured donors resulted in transplants, we do not separate these two outcomes.

3.2.3 Key Questions

We study the OPO's decision to approach a potential donor's family. This step is both consequential and discretionary: it accounts for the largest reduction in the donor pool and is made solely by the OPO. Once an approach occurs, subsequent outcomes such as authorization and transplant depend on external parties (i.e., families and transplant centers). Thus, understanding the OPO's decision to approach is critical for identifying opportunities to increase the number of donors. We investigate three fundamental questions:

1. **How many transplantable donors were referred to the OPO?** We first assess the size of the untapped donor pool. If the OPO could approach all referred potential donors without constraint, how many additional transplants would result? This analysis provides an upper bound on the total opportunity for improvement within the current referral population.
2. **How accurately do OPO approach decisions predict transplant outcomes?** In practice, OPOs cannot make an approach for every referred potential donor, and instead prioritize those who are most likely to lead to successful transplants. Crucially, this assessment is formed under uncertainty: at the time of approach, the OPO does not observe many key predictors of transplant success, such as infection status, full serological testing, and social history. Instead, decisions rely on available information such as age, cause of death, comorbidities, and selected laboratory values. Given this uncertainty, how closely are OPO decisions aligned with true transplantability?
3. **What process or policy changes can increase the number of donors?** While the first two questions assess OPO decisions retrospectively, our final question takes a forward-looking perspective. Which changes to OPO operations would yield the largest gain in transplanted donors?

Table 1. Characteristics of Potential Donors in our Dataset

Variable	All N = 35,856	Approached N = 11,313	Authorized N = 7,018	Transplanted N = 5,164
Age				
Mean (SD)	52.2 (17.9)	43.7 (17.5)	41.7 (16.7)	39.4 (16.6)
Death Type				
Brain Death	6,483 (18.1%)	6,124 (54.1%)	4,855 (69.2%)	4,130 (80.0%)
Cardiac Death	29,373 (81.9%)	5,189 (45.9%)	2,163 (30.8%)	1,034 (20.0%)
First-Person Authorization (FPA)				
Yes	12,292 (34.3%)	4,533 (40.1%)	3,702 (52.8%)	2,692 (52.1%)
No	20,889 (58.3%)	6,190 (54.7%)	3,004 (42.8%)	2,240 (43.4%)
Unknown	2,675 (7.5%)	590 (5.2%)	312 (4.4%)	232 (4.5%)
Race				
White / Caucasian	22,809 (63.6%)	6,845 (60.5%)	4,638 (66.1%)	3,298 (63.9%)
Black / African American	7,247 (20.2%)	2,332 (20.6%)	1,131 (16.1%)	873 (16.9%)
Hispanic	4,998 (13.9%)	1,868 (16.5%)	1,116 (15.9%)	893 (17.3%)
Other / Unknown	802 (2.2%)	268 (2.4%)	133 (1.9%)	100 (1.9%)
Sex				
Male	21,220 (59.2%)	6,758 (59.7%)	4,235 (60.3%)	3,151 (61.0%)
Female	14,636 (40.8%)	4,555 (40.3%)	2,783 (39.7%)	2,013 (39.0%)
Creatinine				
Mean (SD)	1.9 (1.8)	1.5 (1.5)	1.4 (1.4)	1.4 (1.3)
Missing (%)	2,869 (8.0%)	226 (2.0%)	53 (0.8%)	22 (0.4%)
Total Bilirubin				
Mean (SD)	1.3 (3.1)	0.8 (1.9)	0.8 (1.5)	0.7 (1.4)
Missing (%)	7,550 (21.1%)	1,319 (11.7%)	641 (9.1%)	462 (8.9%)

4 Structural Model

The OPO setting presents three empirical challenges: (1) the OPO’s approach decision is driven by both beliefs and preferences, (2) its decision must consider the probability of two outcomes: authorization and transplant, and (3) these outcomes are only observed for potential donors who were approached. It is not possible to address these challenges using a reduced-form model because preferences and beliefs about authorization and transplantation probabilities are latent. We thus propose a structural model that jointly considers the OPO’s approach decision and both downstream outcomes. This structural approach allows us to separate beliefs and preferences while accounting for selection effects. It also allows us to generate actionable managerial insights through counterfactual analyses that simulate interventions on model parameters.

4.1 Observed Outcomes

We consider a population of N potential donors referred to an OPO. For each potential donor i , we observe three binary outcomes: (1) $r_i \in \{0, 1\}$ indicates if the OPO approached the potential donor’s family (2) $a_i \in \{0, 1\}$ indicates whether the family provided consent (or if the potential donor themselves provided consent during their lifetime through a donor registry), and (3) $y_i \in \{0, 1\}$ indicates whether at least one of the donor’s organs was accepted by a transplant center and successfully transplanted in a patient.

Note that a_i and y_i are censored outcomes: a potential donor cannot be authorized unless approached (i.e., $a_i = 0$ if $r_i = 0$), and cannot lead to a transplant unless authorized (i.e., $y_i = 0$ if $a_i = 0$). We define a_i^* and y_i^* as the corresponding uncensored outcomes. Here, a_i^* denotes the authorization outcome if potential donor i were to be approached, and y_i^* denotes the transplant outcome if the potential donor were to be approached and authorized. We only observe a_i^* for approached potential donors (i.e., if $r_i = 1$) and y_i^* for authorized potential donors (i.e., if $a_i = 1$).

4.2 Transplant Outcomes

4.2.1 True Model

We first model the latent outcome of transplantation y_i^* . For each potential donor i , we observe a set of covariates $x_i = [x_i^{(1)}, \dots, x_i^{(d)}]$. These covariates include key donation-relevant variables including

demographics (e.g., age, sex, race), death type (e.g., brain death vs. cardiac death) and laboratory tests (e.g., creatinine). We model y_i^* as a function of these observables,

$$y_i^* = \mathbf{1}[\beta_0^* + x_i^T \beta^* + \epsilon_i > 0]. \quad (1)$$

Here, $\beta_0^* \in \mathbb{R}$ and $\beta^* \in \mathbb{R}^d$ are unknown parameters and ϵ is a latent exogenous index (i.e., $\epsilon_i \perp\!\!\!\perp x_i$).

Now, the OPO can observe key information that affects transplant outcomes but is not observed by the researcher. For example, the OPO often observes a potential donor's neurological state, though this is not noted in the structured data x_i . To capture such information, we assume that the OPO observes a latent signal $\eta_i \perp\!\!\!\perp x_i$ that is correlated with ϵ_i . As is typical in prior work that analyzes medical decisions with censored outcomes (Chan et al. (2022), Freeman et al. (2021)), we assume that η_i and ϵ_i follow a bivariate normal distribution with covariance $\rho^* \in [0, 1]$, and that η_i is an index whose correlation with ϵ_i captures its effect on transplantability. Without loss of generality, we normalize the means and variances of η_i and ϵ_i to 0 and 1.

Assumption 1 η_i and ϵ_i follow a bivariate normal distribution,

$$\begin{bmatrix} \eta_i \\ \epsilon_i \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho^* \\ \rho^* & 1 \end{bmatrix}\right)$$

Together, the parameters $(\beta_0^*, \beta^*, \rho^*)$ capture the true probability that any given potential donor will lead to a transplant given the information observed by the OPO. Under Thm. 1,

$$\mathbb{P}(y_i^* = 1 \mid x_i, \eta_i) = \Phi\left(\frac{\beta_0^* + x_i^T \beta^* + \rho^* \eta_i}{\sqrt{1 - \rho^{*2}}}\right). \quad (2)$$

This expression (derived in App. A.1) clarifies the role of each structural parameter. The intercept β_0^* captures the overall prevalence of donors in the population, as a higher value implies a greater average likelihood of transplant. The vector β^* reflects the marginal effects of the observed covariates x_i on transplant probability. Finally, ρ^* conveys the informativeness of the latent signal η_i relative to the observed covariates. As ρ^* increases, the ratio $\rho^*/\sqrt{1 - \rho^{*2}}$ becomes larger relative to $1/\sqrt{1 - \rho^{*2}}$, increasing the influence of η_i on transplant probability.

4.2.2 OPO Beliefs

The OPO does not directly know $(\beta_0^*, \beta^*, \rho^*)$. Instead, it has formed beliefs about these parameters. These beliefs could be based on past organ accept/reject decisions—which are discrete noisy signals of transplant probability—or on an assessment based on medical knowledge. We explicitly model the OPO’s beliefs as (β_0, β, ρ) , allowing each to deviate from its true value. For example, the OPO may underestimate the overall donor prevalence (i.e., set $\beta_0 < \beta_0^*$) or overestimate the informativeness of its latent signal η_i (i.e., set $\rho > \rho^*$). We assume that the OPO’s perceived probability of transplant is then given by:

$$\hat{\mathbb{P}}_{OPO}(y_i^* = 1 \mid x_i, \eta_i) = \Phi\left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}}\right). \quad (3)$$

Together, the beliefs (β_0, β, ρ) capture the OPO’s ability to accurately predict transplant probability. The closer these beliefs are to their true values, the closer the OPO’s estimate of transplant probability (3) is to the true probability (2), and the more predictive the OPO’s beliefs.

4.3 Authorization Outcomes

4.3.1 True Model

We now consider the authorization outcome a_i^* . We model its probability as

$$\mathbb{P}(a_i^* = 1 \mid x_i, z_i) = \kappa(x_i, z_i) \quad (4)$$

where $z_i \in \{0, 1\}$ is an observed binary feature and κ is an unknown function. Here, z_i is a variable that affects authorization but not transplant (Thm. 2); we later discuss how this instrument allows us to identify key parameters in our model despite the censoring of y_i . We discuss our specific choice of instrument—first person authorization (FPA)—in depth in Section 5.2.

Assumption 2 *There exists an observed binary feature z_i such that*

1. (Relevance) z_i monotonically affects the probability of authorization: $\kappa(x_i, 1) > \kappa(x_i, 0) \forall x_i$.
2. (Exclusion) z_i is independent of the latent indices of transplantation: $z_i \perp\!\!\!\perp \epsilon_i, \eta_i$.

While our model allows the OPO to observe a latent signal of transplant η_i , it restricts the effect of unobservables on authorization (Thm. 3). Specifically, we assume that the OPO does not observe any latent signal of the family’s willingness to consent. This assumption is defensible: OPOs typically do not engage with families until they decide to approach. Further, we assume that authorization and transplantation are conditionally independent outcomes. This restriction is supported by the fact that transplantation is primarily determined by clinical suitability, whereas family consent is influenced by psychosocial factors such as cultural attitudes, religiosity, and prior donor registration (Carola et al. 2023). While some of these variables may be correlated with clinical determinants (e.g., age), their influence is mitigated by our rich covariate set.

Assumption 3 *Authorization is conditionally independent of the latent indices of transplant, $a_i^* \perp \perp (\epsilon_i, \eta_i) \mid x_i, z_i$. Further, the OPO does not observe latent information correlated with $\mathbb{P}(a_i^* = 1)$.*

4.3.2 OPO Beliefs

Our model makes an additional simplification about authorization. While we allow the OPO’s beliefs about transplant probability to be misspecified, we assume that the OPO has accurate beliefs about the authorization probability (Thm. 4). This assumption is grounded in the OPO’s extensive experience interacting with families of potential donors. It is also supported by the structure of the donation process: authorization decisions result from direct conversations between the OPO and the potential donor’s family. Thus, in contrast to transplant decisions—made by transplant physicians further downstream—authorization is shaped by interactions that the OPO both initiates and controls. Moreover, empirical research has documented consistent patterns in consent rates by observable characteristics such as age, race, and death type (Goldberg et al. 2013, Brown et al. 2010, Hulme et al. 2016). These considerations suggest that OPOs are well-positioned to form accurate and calibrated beliefs about authorization.

Assumption 4 *The OPO’s belief about the probability of authorization $\kappa(x_i, z_i)$ is correct for each potential donor i .*

4.4 Approach Decision

We now consider how the OPO uses its beliefs to make its key decision: whether or not to approach the potential donor’s family. Since the OPO has limited resources (i.e., staff and budget), it cannot

attempt to pursue every potential donor. Intuitively, the OPO should prioritize approaching the families of those potential donors who are most likely to lead to successful transplants. Here, we formalize this intuition by considering a simplified version of the OPO’s cost structure.

4.4.1 OPO Utility

Let π denote the total expected benefit of a successful donor (i.e., a donor that leads to a transplant). This benefit can only be realized if the donor is procured, which requires the OPO to (i) approach the donor’s family and obtain authorization, and (ii) recover the donated organ(s). Thus, to realize π , the OPO incurs two costs: an approach cost $C_a\pi$ and a procurement cost $C_p\pi$. The approach cost $C_a\pi$ includes all expenses associated with initiating and conducting family conversations, while the procurement cost $C_p\pi$ includes downstream costs incurred only if authorization is granted (e.g., donor management, logistical coordination, organ placement). For simplicity, we assume that these costs are fixed across donors and express them as a fraction of the total benefit π . The OPO’s overall utility can therefore be expressed as follows:

$$u_i = \begin{cases} 0 & r_i = 0 \\ -C_a\pi & r_i = 1, a_i^* = 0 \\ -C_a\pi - C_p\pi & r_i = 1, a_i^* = 1, y_i^* = 0 \\ \pi - C_a\pi - C_p\pi & r_i = 1, a_i^* = 1, y_i^* = 1 \end{cases} \quad (5)$$

We assume that the OPO’s decision to approach aims to maximize this expected utility. We make two brief notes. First, the costs C_a and C_p convey the OPO’s preferences, as they drive how the OPO values an unsuccessful approach (i.e., which does not lead to transplant) relative to a successful one. Notably, the OPO values an approach that ends in a family decline differently than one that does not lead to a transplant after authorization, as it pays a higher cost in the latter case. Note that C_p and C_a should not be interpreted strictly as financial costs, but rather “as-if” costs that reveal its preferences (i.e., the OPO acts “as-if” these were its costs). Second, a key simplification made by our model is that the OPO receives a fixed reward π for each successful donor. This choice aligns with the OPO’s regulatory incentives: the donation rate² used to assess OPO performance weights each donor equally. In practice, however, the OPO receives a standard acquisition charge (SAC)

²Defined as the number of donors relative to the number of eligible deaths in the OPO’s service area (CMS 2020)

for every organ that is accepted for transplant, and each donor may provide a different number of organs. SACs can also vary by OPO and donor hospital based on factors that influence procurement costs (e.g., distance) (Held et al. 2020, Ozbay et al. 2025). To keep the model tractable, we abstract away from these nuances and assume a constant per-donor benefit.

4.4.2 OPO Optimal Decision Rule

Since the OPO does not observe y_i^* and a_i^* beforehand, it must make decisions based on its expected utility conditional on the information it observes, $\mathbb{E}[u_i | x_i, z_i, \eta_i]$. To maximize this expected utility, the OPO should only approach a given potential donor i if the expected benefit of approach exceeds the expected costs, that is, if $\mathbb{E}[u_i | r_i = 1, x_i, z_i, \eta_i] > \mathbb{E}[u_i | r_i = 0, x_i, z_i, \eta_i]$. Thus,

$$r_i = \mathbf{1} \left[\mathbb{P}(y_i^* = 1 | x_i, \eta_i) > C_p + \frac{C_a}{\kappa(x_i, z_i)} \right]. \quad (6)$$

This rule (derived in App. A.2) indicates that the OPO should only approach for potential donor i if their probability of transplant exceeds a threshold $\tau_i = C_p + C_a/\kappa(x_i, z_i)$. The threshold τ_i is determined by two key terms. First, $C_p + C_a$ defines an absolute floor for transplant probability below which the OPO should never approach. The larger this sum, the higher the relative cost the OPO pays for an unsuccessful approach, and the more certain the OPO needs to be about transplantability to approach. Second, the cost of approach C_a dictates how much the threshold τ_i should be adjusted to account for authorization. Intuitively, as a potential donor's probability of authorization decreases, the less willing the OPO should be to approach, as the more likely the approach is to result in a family decline (resulting in a cost of C_a). This reduced willingness manifests as an increase in the OPO's threshold: the OPO must be even more sure about i 's transplantability to approach if i has low authorization probability. Overall, the costs (C_a, C_p) determine the OPO's willingness to approach (i.e., the total number of approaches it can make while maximizing utility).

Note that the decision rule (6) maximizes the OPO's utility but may not lead to socially optimal outcomes if the OPO overweights its costs relative to the societal benefit created by each additional transplant (e.g., by setting a high C_p that excludes many transplantable donors from being approached). Agarwal et al. (2019) discusses a similar occurrence in kidney exchanges, where optimal choices by individual decision-makers result in suboptimal societal outcomes.

4.4.3 OPO Decision-Rule

To implement the optimal decision rule (6), the OPO must know the true probability of transplant, $\mathbb{P}(y_i^* = 1 \mid x_i, \eta_i)$. However, as discussed in Section 4.3.2, this probability depends on the latent structural parameters $(\beta_0^*, \beta^*, \rho^*)$ that are not directly observed by the OPO. Instead, the OPO must make approach decisions that maximize expected utility $\mathbb{E}[u_i \mid x_i, \eta_i, z_i]$ given its beliefs (β_0, β, ρ) . This yields the following decision rule:

$$r_i = \mathbf{1} \left[\Phi \left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}} \right) > C_p + \frac{C_a}{\kappa(x_i, z_i)} \right] \quad (7)$$

This rule is the same as the optimal rule (6), except that the true probability of transplantation is replaced by the OPO's belief from (3). In summary, the OPO only approaches for potential donor i if this believed probability is greater than the threshold $\tau_i = C_p + C_a / \kappa(x_i, z_i)$. Notably, the threshold τ_i is monotonically shifted by an instrument z_i that is uncorrelated with transplant probability (both true and OPO-believed). This creates variation in approach decisions that is independent of transplant outcomes, which will prove crucial for the identification of our structural parameters.

4.5 Identification

We have now fully specified a data generating process for each of our observed outcomes, (r_i, a_i, y_i) . The key parameters in our model are the true drivers of transplant quality $(\beta_0^*, \beta^*, \rho^*)$, the OPO beliefs (β_0, β, ρ) , the costs (C_p, C_a) , and the authorization probability $\kappa(x_i, z_i)$. We now discuss how each key parameter in our model is identified from the observed data.

4.5.1 Identification of $\kappa(x_i, z_i)$

The authorization probability $\kappa(x_i, z_i)$ is directly identified from the authorization outcomes a_i . Though a_i is only observed if $r_i = 1$, this selection does not depend on unobservables (Thm. 3). Thus, standard binary choice results (Matzkin 1992) establish that $\kappa(x_i, z_i)$ is non-parametrically identified from the observed a_i .

4.5.2 Identification of $\beta_0^*, \beta^*, \rho^*$

Next, we discuss identification of the true drivers of transplant quality. The key challenge for identification is censoring: y_i^* is only observed if $r_i = 1$ and $a_i^* = 1$. Since the approach mechanism r_i is driven by observables x_i and unobservables η_i —both of which are correlated with y_i^* — $(\beta_0^*, \beta^*, \rho^*)$ cannot be identified from y_i alone. This is an example of sample selection, where outcomes are only observed for a non-random subset of potential donors.

We use the Heckman selection framework (Heckman 1979, Van de Ven and Van Praag 1981) to identify $(\beta_0^*, \beta^*, \rho^*)$. Overall, these parameters are identified by the instrument z_i that affects the OPO's approach, but not transplantation (Thm. 2). To apply the Heckman framework, we first express the OPO's decision rule (7) in partially linear form,

$$r_i = \mathbf{1}[\tilde{\beta}_0 - s(\tau_i) + x_i^T \tilde{\beta} + \eta_i > 0], \quad (8)$$

where $\tilde{\beta}_0 = \beta_0/\rho$, $\tilde{\beta} = \beta/\rho$, and the function $s(\tau_i) = \frac{\sqrt{1-\rho^2}}{\rho} \Phi^{-1}(\tau_i)$. Together, (1) and (8) define a canonical Heckman biprobit model (Van de Ven and Van Praag 1981) with r_i as the selection mechanism and $y_i^*|r_i = 1$ as the censored outcome. Identification in this model requires a variable that is included in the selection equation, but excluded from the outcome equation (Wooldridge 2010). This variation is provided by the instrument z_i , which has a monotonically decreasing effect on the threshold τ_i (and thus an increasing effect on $\mathbb{E}[r_i]$), but no effect on y_i^* . This establishes that $(\beta^*, \beta_0^*, \rho^*)$ is identified from the distribution of $y_i^*|r_i = 1, x_i$. Consequently, $(\beta^*, \beta_0^*, \rho^*)$ is also identified by the distribution of $y_i^*|r_i = 1, a_i^* = 1, x_i$, as $a_i^* \perp\!\!\!\perp y_i^*|x_i$ (Thm. 3). This establishes identification, as the latter distribution is directly identified from the censored outcomes y_i .

4.5.3 Identification of $(C_p, C_a, \beta_0, \beta, \rho)$

Finally, we consider the OPO's belief parameters (β_0, β, ρ) and cost parameters (C_p, C_a) , which jointly determine its approach decision r_i . As shown in (8), the OPO's decision rule for r_i can be expressed as a partially linear binary choice model. Standard results establish that the approach probability $\Pr(r_i = 1 | x_i, z_i)$ is nonparametrically identified (Matzkin 1992). Further, our instrument z_i shifts the threshold $C_p + C_a/\kappa(x_i, z_i)$ without affecting the OPO's beliefs. As a result, $\kappa(x_i, z_i)$ is not collinear with x_i , and we can non-parametrically separate the coefficient $\tilde{\beta} = \beta/\rho$ from the function $\tilde{s}(x_i, z_i) = \tilde{\beta}_0 - s(C_p + C_a/\kappa_i)$. Finally, we separate each component in the function using

variation in the authorization probabilities $\kappa(x_i, z_i)$. For each $\kappa(x_i, z_i)$ we have identified a function $\tilde{s}(x_i, z_i)$ that depends on four unknowns: $(C_p, C_a, \beta_0, \rho)$. Thus, if $\kappa(x_i, z_i)$ takes on more than four values, we obtain an overidentified nonlinear system, from which the parameters can be recovered via nonlinear least squares.

4.5.4 Calibration of OPO Beliefs

The previous identification argument for $(C_p, C_a, \beta_0, \beta, \rho)$ relies on the parametric form of $\tilde{s}$. A core challenge for non-parametric identification is that we allow the OPO's beliefs to be misspecified: (β_0, β, ρ) may differ from the true data-generating parameters $(\beta_0^*, \beta^*, \rho^*)$. Without further restrictions, we cannot use the observed transplant outcomes y_i to identify these beliefs. However, given the OPO's extensive experience in organ procurement, it is implausible that the OPO's beliefs are entirely inaccurate (e.g., uncorrelated with y_i). Prior work has thus overcome this challenge by assuming that the decision-maker has correct beliefs about one (or more) of the true structural parameters. For example, Chandra and Staiger (2020) assumes knowledge of both β_0^* and β^* , Abaluck et al. (2016) assumes knowledge of β_0^* , and Chan et al. (2022) assumes knowledge of ρ^* .

Our model imposes milder restrictions on belief accuracy. Rather than assuming that the OPO knows a latent structural parameter, we require only that its beliefs are consistent with two aggregate metrics of its procurement performance. Specifically, we assume that the OPO's beliefs are calibrated to: (1) the average transplant rate conditional on authorization, $\mathbb{E}[y_i^* \mid a_i = 1]$, and (2) the average transplant rate among “marginally approached” potential donors (i.e., potential donors whose estimated transplant probability lies just above the threshold). Each of these conditions implies a moment restriction that regularizes the OPO's beliefs using limited information from y_i (see App. A.3). While these assumptions are not strictly necessary for identification, they reduce our model's reliance on functional form assumptions and improve performance in finite samples.

Assumption 5 *Let $\mathbb{P}(A)$ denote the true probability of event A , and let $\hat{\mathbb{P}}^{OPO}(A)$ denote the OPO's estimated probability under its belief parameters (β_0, β, ρ) . We assume the following:*

1. (Aggregate Calibration) *The OPO's belief about the average transplant probability among authorized potential donors is correct,*

$$\hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid a_i = 1) = \mathbb{P}(y_i^* = 1 \mid a_i = 1).$$

2. (*Near-Threshold Calibration*) Let $\hat{p}_i = \hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid x_i, \eta_i)$. The OPO’s belief about the average transplant probability for “marginally approached” potential donors (i.e., with $\hat{p}_i \approx \tau_i$) is correct,

$$\hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid \hat{p}_i \in M_i^\delta) = \mathbb{P}(y_i^* = 1 \mid \hat{p}_i \in M_i^\delta)$$

where $M_i^\delta = [\tau_i, \tau_i + \delta]$ for some small $\delta > 0$.

The first condition is intuitive: it ensures that the OPO’s beliefs correctly predict the average transplant rate among authorized potential donors. This metric is both observable and utility-relevant, as it captures how often the OPO incurs the cost C_p without realizing the benefit π . The second condition is more nuanced but still behaviorally motivated. It asserts that the OPO’s beliefs are, on average, well-calibrated for *marginally approached* potential donors: those whose estimated transplant probability $\hat{p}_i$ lies just above the decision threshold τ_i . These potential donors are uniquely relevant to the OPO’s utility, since incorrect beliefs near the threshold are more likely to result in incorrect decisions. For example, if the OPO overestimates transplant probability for potential donors near the threshold, it would approach lower-quality donors than it intends to, and would observe that these potential donors result in a lower-than-expected average transplant rate. In contrast, mistaken beliefs further away from the threshold are less consequential, as the OPO’s decision with correct beliefs may remain unchanged. Overall, the near-threshold calibration condition reflects the idea that the OPO has the greatest incentive to be accurate where errors are most consequential. It also implies a form of local optimality: if the OPO observes that marginal donors consistently underperform expectations, it can improve utility by slightly raising its threshold. Together, our two calibration conditions separate the OPO’s beliefs from its preferences by assuming a baseline level of accuracy grounded in real-world incentives. That said, if these conditions are violated in practice, we note that our model may overstate the accuracy of OPO beliefs.

4.6 Estimation

Our estimation proceeds in two stages. We first estimate the authorization probabilities $\kappa(x_i, z_i)$ using an L1-regularized, nonparametrically-calibrated probit regression of the observed outcomes a_i on the covariates (x_i, z_i) . This estimation is described in depth in App. B.2. We note that our final predictions are generated using 10-fold cross-fitting, so the prediction $\hat{\kappa}(x_i, z_i)$ for each potential

donor i is generated from a model that was not trained on its authorization outcome a_i .

We then use Bayesian inference to fit our model of OPO behavior to the observed data and estimated $\hat{\kappa}(x_i, z_i)$. We define an uninformative prior for each of our model parameters $(\beta_0^*, \beta^*, \rho^*, \beta_0, \beta, \rho, C_a, C_p)$, then infer their posterior distributions using Markov Chain Monte Carlo (MCMC). Our calibration conditions (Thm. 5) are incorporated as penalty terms in the model’s log likelihood. We implemented our model in Stan (Stan Development Team 2024), using the No-U-Turn Sampler (Hoffman et al. 2014) for MCMC inference with 4 chains of 3,000 draws each. The first half of each chain is discarded as burn-in; the remaining 6,000 draws can be interpreted as samples from the true joint posterior distribution of our structural parameters. We computed mean estimates, standard errors, and 95% credible intervals for each parameter using the mean, standard deviation, and quantiles of these samples. This estimation procedure is asymptotically equivalent to maximum likelihood based on the Bernstein-von-Mises theorem (Theorem 10.1 in Van der Vaart (2000)). We provide more details on our estimation procedure—including priors, full model equations, and convergence diagnostics—in App. B.3.

To estimate our model using data from multiple OPOs, we use two specifications. The *unpooled* specification separately estimates the model for each OPO, allowing all structural parameters to vary. This approach allows for maximum heterogeneity between OPOs. In contrast, the *pooled costs* specification assumes that the costs (C_p, C_a) are constant across OPOs. This specification reflects that OPOs are likely to have similar costs (at least relative to their per-donor benefit π), as they all operate under the same reimbursement structure. Pooling these cost parameters across OPOs allows us to estimate them with higher precision. Crucially, both specifications allow for important heterogeneity in non-cost parameters; for example, the drivers of transplant likelihood $(\beta_0^*, \beta^*, \rho^*)$ can vary based on the preferences of the transplant centers closest to each OPO.³ We find that both specifications yield similar results; we present the more-parsimonious pooled specification in the main text and include the unpooled specification as a robustness check in App. C.2.

5 Empirical Results

We now present the results of fitting our structural model to the data from OPOs A-D (described in Section 3.2). We begin by describing our key covariates and choice of instrument. We then present

³Since organs are offered to nearby transplant centers (i.e., within 250 miles of the donor) before further-away centers (OPTN 2025)

Table 2. FPA instrument relevance. We display coefficients, heteroskedasticity-consistent standard errors, and F-values of a probit regression of approach / authorization on FPA, adjusting for all covariates.

Outcome	Coef (SE)	F-statistic
Authorization	0.65 (0.023)	829.5
Approach	0.07 (0.019)	14.1

the results of model estimation, highlighting key parameter values and their interpretation. In the sections that follow, we use these estimates to evaluate the efficiency of OPO decisions and explore the broader policy implications of our findings.

5.1 Covariates

Our analysis included a rich set of observed covariates. Following prior work on consent for organ donation (Goldberg et al. 2013, Siminoff et al. 2001, Levan et al. 2022), our model of authorization included demographic covariates such as age, race, sex, and brain death status. Our model of transplant probability incorporated a broader set of clinical factors, including cause of death, comorbidities (e.g., diabetes, hypertension), laboratory test results (e.g., creatinine, total bilirubin), and active treatments (e.g., dialysis, vasopressors). We accounted for these variables using well-established specifications from the clinical literature; for example, our linear-spline specification for creatinine matches that used in the Kidney Donor Risk Index (Rao et al. 2009). We also included fixed effects for year to account for trends in authorization and transplant probability over time. A full description of covariates and specifications is provided in App. B.1.

5.2 Instrument

Our identification strategy relies on an instrument z_i that affects authorization but not transplant outcomes. We use first-person authorization (FPA) as our instrument. FPA is a binary variable that indicates whether a potential donor consented to organ donation during their lifetime by enrolling in a donor registry (e.g., through a driver’s license application). FPA is a strong and legally binding predictor of authorization: in most states, individuals who provide FPA are considered to have given consent for organ donation that cannot be revoked by their next-of-kin (Iltis and Denny 2025).⁴ We

⁴In practice, OPOs may defer to family members who strongly object to donation, and so FPA does not guarantee authorization.

Table 3. Mean predicted authorization probabilities by demographic group.

Covariate	Group	Authorization Probability			
		OPO A	OPO B	OPO C	OPO D
Age	[0,18)	54.3%	39.3%	52.1%	48.9%
	[18,25)	63.5%	62.3%	65.6%	61.4%
	[25,35)	60.4%	53.9%	65%	68.2%
	[35,45)	58%	50.8%	63.2%	61.7%
	[45,55)	52.1%	44.2%	57.8%	57.8%
	[55,65)	40.4%	34.2%	46.9%	49.4%
	[65,75)	30.1%	17.5%	28.7%	47.6%
	[75,80)	25.5%	11.3%	26.6%	45.8%
Death Type	Brain Death	75.8%	81%	76.3%	75.6%
	Not Brain Death	36.7%	25.5%	42.1%	50.9%
Race	White / Caucasian	47.1%	43.2%	54.5%	57.6%
	Black / African American	33%	25.4%	29.4%	36.5%
	Hispanic	49.8%	34.4%	55.8%	52.4%
	Other / Unknown	35.6%	29.1%	33.7%	49.3%

confirm empirically that FPA has a strong effect on authorization outcomes for each OPO (Table 2). In addition, FPA also has a strong effect on the OPO’s decision to approach, supporting our model’s assertion that the OPO adjusts its approach decisions to account for authorization. The significant effect of FPA on approach and authorization—evidenced by large F-statistics (> 10) and positive coefficients—supports the instrument’s relevance (as required by Thm. 2).

Notably, FPA is typically determined well before clinical signs of imminent death. Most individuals register as organ donors through routine interactions with state agencies: for example, when obtaining a driver’s license or registering a vehicle at the Department of Motor Vehicles (DMV) (Dageforde et al. 2020, Feeley et al. 2017). In contrast, transplant suitability depends heavily on clinical factors present at the time of death, such as cause of death, neurological status, infection, and organ function. Given this temporal gap, it is unlikely that FPA status is systematically correlated with the unobserved clinical signals that drive transplant probability. Thus, FPA arguably satisfies the exclusion restriction and is relevant for the approach decision.

Table 4. Estimated drivers of transplant probability and OPO beliefs (select coefficients). We display mean estimated coefficients and 95% credible intervals (in parentheses) across 6,000 MCMC draws. Bolded and starred coefficients have credible intervals that exclude zero.

I. OPO A and B

Variable	OPO A		OPO B	
	OPO	True	OPO	True
Strength of Latent Signal (ρ, ρ^*)	0.52* (0.44, 0.61)	0.04* (0.03, 0.06)	0.49* (0.42, 0.57)	0.04* (0.03, 0.06)
Intercept (β_0, β_0^*)	-0.12 (-0.33, 0.07)	0.31 (-0.21, 0.83)	-0.22* (-0.41, -0.05)	0.73* (0.20, 1.32)
Brain Death (vs. Cardiac Death)	1.21* (1.02, 1.42)	1.42* (1.17, 1.67)	1.29* (1.10, 1.49)	1.24* (1.05, 1.43)
Age	-0.19* (-0.25, -0.14)	-0.44* (-0.63, -0.26)	-0.08* (-0.12, -0.04)	-0.41* (-0.59, -0.24)
Age $\times \mathbf{1}\{\text{Age} > 50\}$	-0.33* (-0.44, -0.23)	-0.17 (-0.56, 0.22)	-0.20* (-0.28, -0.13)	-0.29 (-0.60, 0.03)
Diabetes	0.02 (-0.03, 0.08)	-0.35* (-0.61, -0.10)	-0.04* (-0.07, -0.01)	-0.19* (-0.38, -0.01)
Sepsis	-0.33* (-0.42, -0.25)	-0.55* (-0.93, -0.18)	-0.23* (-0.30, -0.17)	-0.33* (-0.65, -0.003)
Creatinine	-0.29* (-0.38, -0.21)	-0.19 (-0.51, 0.11)	-0.13* (-0.18, -0.08)	-0.27* (-0.54, -0.01)
log(Total Bilirubin)	-0.10* (-0.13, -0.07)	-0.14* (-0.28, -0.01)	-0.07* (-0.09, -0.05)	-0.18* (-0.31, -0.06)

II. OPO C and D

Variable	OPO C		OPO D	
	OPO	True	OPO	True
Strength of Latent Signal (ρ, ρ^*)	0.60* (0.51, 0.69)	0.05* (0.04, 0.07)	0.43* (0.33, 0.54)	0.04* (0.03, 0.06)
Intercept (β_0, β_0^*)	0.19 (-0.06, 0.43)	0.62 (-0.02, 1.35)	-0.74* (-1.09, -0.44)	-0.45 (-1.08, 0.14)
Brain Death (vs. Cardiac Death)	1.23* (1.05, 1.41)	1.25* (1.01, 1.49)	1.18* (0.91, 1.48)	1.56* (1.26, 1.89)
Age	-0.03 (-0.09, 0.03)	-0.24* (-0.44, -0.05)	0.03 (-0.02, 0.08)	-0.30* (-0.53, -0.08)
Age $\times \mathbf{1}\{\text{Age} > 50\}$	-0.61* (-0.76, -0.47)	-0.57* (-1.06, -0.10)	-0.51* (-0.69, -0.37)	-0.34 (-0.83, 0.13)
Diabetes	-0.09* (-0.16, -0.02)	-0.22 (-0.51, 0.07)	-0.04 (-0.10, 0.01)	-0.29 (-0.61, 0.05)
Sepsis	-0.38* (-0.51, -0.26)	-0.36 (-0.84, 0.14)	-0.13* (-0.23, -0.04)	-0.004 (-0.48, 0.49)
Creatinine	-0.31* (-0.41, -0.22)	-0.58* (-0.94, -0.24)	-0.16* (-0.25, -0.09)	-0.23 (-0.56, 0.09)
log(Total Bilirubin)	-0.07* (-0.11, -0.04)	-0.17* (-0.34, -0.01)	-0.09* (-0.13, -0.05)	-0.24* (-0.45, -0.05)

Table 5. Estimated OPO costs and approach thresholds. We present mean estimated parameters and 95% credible intervals (in parentheses) across 6,000 MCMC draws.

	Mean	95% CI
C_p	0.404	(0.335, 0.463)
C_a	0.002	(0.000, 0.003)
$\min \tau_i$ ($= C_p + C_a$)	0.406	(0.337, 0.465)
$\max \tau_i$ ($= C_p + C_a / \min \hat{k}_i$)	0.507	(0.413, 0.592)

5.3 Parameter Estimates

We now present our estimates of authorization probability $\kappa(x_i, z_i)$ and structural parameters $(\beta_0^*, \beta^*, \rho^*, \beta_0, \beta, \rho, C_a, C_p)$, along with brief interpretations. In the following sections, we use these estimates to examine key domain-relevant findings and explore their implications for policy.

5.3.1 Authorization Model

The average predicted probability of authorization conditional on approach was 63% for OPO A, 58% for OPO B, 67% for OPO C, and 68% for OPO D, closely matching previously reported estimates for other OPOs (Levan et al. 2022). We display mean predicted authorization rates by demographic group in Table 3. Consistent with prior work (Goldberg et al. 2013), we find that authorization rates were lower among potential donors who were non-White, over the age of 55, or did not experience brain death. Overall, our authorization model demonstrates strong predictive performance, achieving a held-out AUROC of 0.83 across all OPOs.

5.3.2 True Transplant Probability and OPO Beliefs

We present estimated coefficients for the drivers of transplant probability $(\beta_0^*, \beta^*, \rho^*)$ in Table 4 (select variables) and Tables 8 and 9 (full). Our estimates are consistent with prior findings (Rao et al. 2009, Cohen et al. 2017, Collett et al. 2017): brain death was a strong positive predictor of transplant suitability, while older age, diabetes, sepsis, elevated creatinine, and elevated bilirubin were associated with lower transplant probability. These tables also display the OPO beliefs (β_0, β, ρ) alongside their true counterparts. We find that OPO beliefs generally align with the true coefficients. In the following section, we investigate the extent to which these beliefs lead to accurate predictions of transplant outcomes.




Figure 2. Illustration of the OPO’s decision. We plot a histogram of OPO-estimated transplant probability $\hat{\mathbb{P}}^{OPO}$ for all potential donors in our dataset. For each potential donor, we generate $\hat{\mathbb{P}}^{OPO}$ using our mean estimates of (β_0, β, ρ) and a random sample from the posterior of η_i . OPOs only approach potential donors for whom $\hat{\mathbb{P}}^{OPO}$ exceeds the threshold $\tau_i = C_p + C_a / \kappa(x_i, z_i)$. The left facet plots the histogram for potential donors whose families were not approached, while the right facet plots the corresponding histogram for approached potential donors. For comparison, we plot the minimum threshold under which the OPO would never approach, $C_p + C_a$ (dashed line).

5.3.3 OPO Costs

Finally, we present our estimates of the OPOs’ cost parameters C_p and C_a . Recall that these cost estimates do not necessarily reflect the underlying operational costs, but rather are “as-if” costs that reveal the OPO’s preferences. We estimate $C_p = 0.404$ and $C_a = 0.002$, which together imply a minimum transplant probability threshold of $\tau_i = 0.406$ to justify approach. OPOs raised this threshold substantially for potential donors with low authorization probability. For example, among potential donors with the lowest estimated authorization probability ($\kappa_i = 0.02$), the implied threshold was as high as $\tau_i = 0.507$. We visualize the decision rule implied by these estimated costs and beliefs in Fig. 2.

6 Key Findings and Domain Implications

We now use our estimated structural model to address two key research questions about organ procurement. First, we quantify the total number of transplantable donors within the referred population, independent of current OPO operations. Second, we assess the quality of OPO approach decisions by evaluating the predictive accuracy of their beliefs. These analyses offer important domain-relevant insights and lay the foundation for counterfactual policy analyses in Section 7.

6.1 How Many Transplantable Donors are Referred to an OPO?




Figure 3. Total Number of Transplantable Donors. We plot the additional number of donors the OPOs would procure if they (1) approached for all potential donors (“Approach All”), and (2) obtained authorization from all approached families (“No Family Declines”). We plot the mean numbers of transplantable donors across all four OPOs; error bars indicate 95% credible intervals.

We first estimate the total number of transplantable donors in the population of potential donors referred to the OPOs. This number represents the full opportunity available to OPOs, that is, the maximum number of donors that could be recovered without any practical constraints. Specifically, realizing this opportunity would require OPOs to (1) approach all donors, and (2) obtain authorization from every approached family. Note that this upper bound is not practically achievable; for example, an OPO only has a limited influence on the family’s decision and some families will decline consent even with the best approach practices. However, this estimate provides a useful benchmark to assess the scope of potential improvements.

We estimate the total number of transplantable donors using our model’s estimates of true transplant probability, $\mathbb{P}(y_i^* = 1 \mid x_i, \eta_i)$. For a single draw of $(\beta_0^*, \beta^*, \rho^*)$, we simulate a value of η_i for each potential donor from its distribution conditional on approach. We then compute the implied probabilities $\mathbb{P}(y_i^* = 1 \mid x_i, \eta_i)$ and simulate transplant outcomes y_i^* . Repeating this process across all 6,000 posterior draws, we obtain a distribution of the total number of transplantable donors $\sum_i y_i^*$. We report the mean estimate and 95% credible interval from these simulations. This approach accounts for uncertainty in both donor-level outcomes (i.e., from ϵ_i) and parameter estimation. Further, unlike prior estimates of this opportunity set (Lee et al. 2022), our analysis accounts for

selection on the unobservable η_i due to the censoring of transplant outcomes .

Fig. 3 visualizes the opportunity set for all OPOs combined. Overall, we estimate that if the OPOs approached all potential donors and secured authorization from all families, they would have recovered transplantable organs from 13,103 donors—more than 2.5 times the actual number of recovered donors (5,164). Approaching all potential donors would capture approximately 28% of this total opportunity, with the remaining missed donors attributable to authorization. These estimates highlight the substantial scope for improving organ procurement. In Section 7, we explore a range of policy interventions that could help OPOs capture some of this unrealized opportunity.

6.2 How Accurate Are OPO Beliefs?

Next, we use our model to evaluate the accuracy of OPO beliefs. We split our evaluation into three subquestions: (1) how closely do OPO approach decisions align with “optimal” decisions (i.e., decisions informed by correct beliefs), (2) how accurately do OPO beliefs predict transplant outcomes, and (3) do OPO beliefs exhibit systematic miscalibration?

6.2.1 Alignment with Optimal Rule

We first assess how different OPO approach decisions would be with correct beliefs of transplant probability. Recall equation (6), which established that the OPO’s optimal approach rule—given fixed costs C_p and C_a —is to approach for a potential donor if their transplant probability $\mathbb{P}(y_i^* | x_i, \eta_i)$ exceeds the threshold τ_i . Since an OPO uses its believed probabilities $\hat{\mathbb{P}}^{OPO}(y_i^* | x_i, \eta_i)$ instead of the true probabilities, it can make two errors:

1. *Missed Opportunities*: potential donors who were not approached, but would have been if the OPO had correctly estimated their transplant probability. For these potential donors,

$$\hat{\mathbb{P}}^{OPO}(y_i^* | x_i, \eta_i) < \tau_i < \mathbb{P}(y_i^* | x_i, \eta_i).$$

2. *Low-Benefit Approaches*: potential donors who were approached, but would not have been if the OPO had correctly estimated their transplant probability. For these potential donors,

$$\mathbb{P}(y_i^* | x_i, \eta_i) < \tau_i < \hat{\mathbb{P}}^{OPO}(y_i^* | x_i, \eta_i).$$




Figure 4. Evaluating the accuracy of OPO Beliefs. (a) OPO decision errors. We plot a histogram of the true transplant probability implied by our estimated $(\beta_0^*, \beta^*, \rho^*)$ for potential donors who were not approached (top facet) and approached (bottom facet) across all four OPOs. The dashed line conveys the minimum approach threshold, $C_p + C_a$. Missed Opportunities are potential donors with high transplant probability who were not approached, while Low-Benefit Approaches are those with low transplant probability who were approached. (b) ROC curves that capture the predictive accuracy of OPO beliefs (aggregated across all OPOs). We compare the predictive accuracy of OPO beliefs (orange squares) versus true transplant probability (blue diamonds). We plot the rate of true positives (i.e., transplantable donors who were correctly predicted as transplantable) versus false positives (i.e., untransplantable donors who were incorrectly predicted as transplantable) for a range of classification thresholds; the area under the curve (AUROC) summarizes predictive accuracy.

We use our model to quantify the frequency of each error. We find that the OPOs’ beliefs led to 4,705 missed opportunities (95% CI: [3476, 6346]) and 2,323 low-benefit approaches (95% CI: [1782, 2833]). Said another way, 21% of all approaches were low-benefit, while 19% of non-approaches were missed opportunities. These results indicate that there is substantial value in correcting OPO beliefs on transplant probability.

6.2.2 Predictive Accuracy of OPO Beliefs

While the preceding analysis shows that OPO approach decisions often diverged from the optimal rule, it does not explain why. One potential reason is if the OPO’s predicted probabilities $\hat{\mathbb{P}}^{\text{OPO}}$ poorly discriminated between transplantable and untransplantable donors. To assess this explanation, we characterize the predictive accuracy of OPO beliefs using AUROC, a standard binary classification metric (Fawcett 2006). In our setting, AUROC measures the probability that a randomly selected

Table 6. Estimated AUROC by OPO. We display the mean posterior AUROC for each OPO’s beliefs along with 95% credible intervals.

	OPO A	OPO B	OPO C	OPO D
AUROC: OPO Beliefs	0.78 (0.75, 0.81)	0.75 (0.73, 0.77)	0.78 (0.75, 0.81)	0.74 (0.69, 0.79)
AUROC: Correct Beliefs	0.86 (0.83, 0.89)	0.86 (0.84, 0.88)	0.89 (0.86, 0.91)	0.85 (0.81, 0.89)

transplantable donor (i.e., with $y_i^* = 1$) is assigned a higher predicted probability $\hat{\mathbb{P}}^{\text{OPO}}$ than a randomly selected untransplantable donor (i.e., with $y_i^* = 0$). It thus summarizes how accurately the OPO ranks potential donors by transplant probability: an AUROC of 1.0 indicates perfect ranking, while an AUROC of 0.5 indicates no predictive discrimination (i.e., random guessing). We compute the AUROC for each MCMC draw by comparing the posterior samples of $\hat{\mathbb{P}}^{\text{OPO}}$ to true outcomes y_i^* simulated from the posterior samples of true transplant probability. We report the mean AUROC with 95% credible intervals for each OPO separately, as well as in aggregate across all OPOs.

Overall, we find that OPO beliefs were accurate at predicting transplant outcomes. The four OPOs in aggregate achieved an AUROC of 0.77 (95% CI: [0.75, 0.78]). For comparison, using the true probabilities $\mathbb{P}(y_i^* = 1 \mid x_i, \eta_i)$ to classify potential donors yielded an AUROC of 0.87 (95% CI: [0.85, 0.88]), which reflects the maximum achievable accuracy given observed information. There was some variation in predictive accuracy between OPOs (Table 6), though all four OPOs showed similar trends. Overall, this analysis indicates that OPO beliefs were predictive; AUROCs ≥ 0.74 are highly accurate considering the limited amount of information the OPO has before approach. However, the gap between the OPO and optimal AUROCs suggest that correcting the OPO’s ranking of potential donors could lead to further improvements in decision quality. We evaluate the impact of such improvement in our policy evaluation in Section 7.

6.2.3 Calibration of OPO Beliefs

While AUROC is a valuable metric, it evaluates only the ranking accuracy of predicted probabilities, and not how close these predictions are to the true probabilities. To provide a fuller picture of predictive accuracy, we assess the calibration of OPO beliefs. Specifically, we bin the true transplant probabilities into equal-width intervals (0.1 increments from 0 to 1) and examine the miscalibration of beliefs within each bin. Here, we define miscalibration as the difference between the true transplant probability and the OPO’s belief. The closer miscalibration is to zero, the better




Figure 5. Miscalibration of OPO Beliefs. We plot the distribution of OPO belief miscalibration by true transplant probability. Miscalibration (y-axis) is defined as the true probability minus the OPO’s belief. Potential donors are grouped into ten bins by true probability; within each bin, we show a boxplot of miscalibration, averaged across MCMC draws. The solid line marks the median, the box spans the interquartile range, and the whiskers indicate the 2.5th and 97.5th percentiles. The dashed grey line indicates perfect calibration (i.e., belief equals truth).

the OPO’s beliefs align with transplant probability.

Fig. 5 plots the distribution of the OPO’s miscalibration for all potential donors within each true probability bin (averaged across MCMC draws). The plot reveals a clear pattern: OPOs often underestimated transplant probability for high-quality donors. In bins where the true probability exceeded 0.5, OPOs underestimated transplant probability for 75% of potential donors (i.e., 75% of the miscalibration distribution is above zero). OPO beliefs were better calibrated for low-likelihood donors; for example, for true probability below 0.3, the median miscalibration was close to 0.

These findings add important nuance to our assessment of predictive accuracy. While OPOs accurately ranked potential donors, they systematically underestimated transplant probability. This underestimation is not uniform across potential donors, but rather most pronounced for higher-quality donors with true transplant probability above 0.5. Overall, our calibration analysis indicates that the deviations from the optimal rule identified in Section 6.2.1 stemmed not only from ranking errors, but also from pessimistic beliefs about transplantability. In Section 7, we evaluate the relative impact of correcting each of these inaccuracies.

7 Evaluating Alternative Policies

Finally, we use our structural model to evaluate the impact of potential changes to OPO operations. A key advantage of structural modeling is its ability to support counterfactual policy analysis: by simulating interventions on key model parameters, we can assess their downstream effects on transplant outcomes. We focus on three key aspects of OPO operations: (1) willingness to approach, (2) beliefs about transplant probability, and (3) ability to obtain authorization from families.

7.1 Policies

7.1.1 Increase Approaches

We first assess the impact of increasing the OPO’s willingness to approach. Holding the OPO’s beliefs and authorization ability constant, we increase its willingness to approach by lowering the procurement cost C_p . We consider two policies:

Policy 1: Approach 20% More: The OPO reduces its procurement cost C_p enough to increase its total number of approaches by 20%.

Policy 2: Approach All: The OPO approaches the families of all potential donors. This policy reflects an upper bound on the gains achievable solely by expanding approach rates. It corresponds to a situation where the OPO disregards the cost of an unsuccessful approach (i.e., acts as if $C_p = C_a = 0$).

7.1.2 Correct Beliefs

Next, we hold the OPO’s costs and authorization ability fixed and correct its beliefs about transplant probability. Our earlier evaluation of OPO accuracy suggested two sources of inefficiency: mis-ranking of potential donors and underestimation of the overall prevalence of transplantable donors. To disentangle these effects, we evaluate two policies:

Policy 3: Correct Beliefs (Fixed Approach): The OPO corrects its beliefs about transplant probability but maintains the same number of approaches. This policy isolates the gains from improved donor ranking without changing the perceived donor prevalence. We implement this policy by assuming the OPO knows the true $(\beta_0^*, \beta^*, \rho^*)$ but increases C_p to keep the number of approaches fixed.

Policy 4: Correct Beliefs (Increased Approach): The OPO corrects its beliefs to $(\beta_0^*, \beta^*, \rho^*)$ and allows the number of approaches to vary accordingly. If more potential donors now exceed the existing thresholds, the OPO approaches more. This policy captures gains from correcting both the OPO’s donor ranking *and* belief of donor prevalence.

7.1.3 Improve Authorization

Finally, we evaluate the impact of improving the OPO’s ability to obtain family authorization. In our structural model, the probability of authorization was treated as a fixed quantity. However, in practice, the OPO plays an active role in influencing authorization outcomes. For example, the timing of the approach, the information communicated to the family about the donation process, and the training of the coordinator making the request are all modifiable factors known to affect consent rates (Simpkin et al. 2009). We evaluate the following policies:

Policy 5: Improve Authorization (5%): The OPO raises its mean authorization rate by 5 percentage points. We implement this policy by shifting each potential donor’s inverse authorization probability $\Phi^{-1}(\kappa_i)$ by a positive constant, which raises the mean authorization rate conditional on approach from 62% to 67%.

Policy 6: Improve Authorization (Equity): The OPO eliminates racial disparities in authorization rates. Prior evidence (Goldberg et al. 2013, Levan et al. 2025, Rodrigue et al. 2006) and our estimates of κ_i both document lower authorization rates among Black and Hispanic families relative to White families. We consider a scenario in which the OPO raises the authorization rate from non-White families to match that of White families (while still maintaining variation by sex, age, brain death, and FPA status).

7.2 Results and Discussion

Table 7 summarizes the results of our policy evaluation. We primarily evaluate each policy on the number of additional transplanted donors, as each additional donor creates nearly 31 additional life-years for waitlisted patients (Schnitzler et al. 2005). We also display the number of additional approaches and authorizations required by each policy. Recall that in our model, the OPO pays $C_a\pi$ for each approach and $C_p\pi$ for attempting to procure each authorized donor; these two metrics thus summarize the additional implementation costs for the OPO.

Table 7. Results of counterfactual policy evaluation. The Status Quo row reports the observed number of donors, approaches, and authorizations; subsequent rows report the incremental Δ relative to this baseline. 95% CI's represent Bayesian credible intervals.

Policy Name	Donors		Approaches		Authorizations	
	Δ	95% CI	Δ	95% CI	Δ	95% CI
0. Status Quo	5,164		11,313		7,018	
	+		+		+	
1. Approach 20% More	370	(332, 410)	2,263	(2260, 2265)	905	(857, 952)
2. Approach All	2,195	(1980, 2427)	24,543	(24543, 24543)	8,757	(8629, 8885)
3. Correct Beliefs (Fixed Approach)	496	(428, 567)	0	(-2, 2)	253	(135, 365)
4. Correct Beliefs (Increased Approach)	955	(626, 1321)	2,382	(670, 4566)	1,271	(555, 2124)
5. Improve Authorization (5%)	252	(174, 331)	0	(0, 0)	537	(453, 618)
6. Improve Authorization (Equity)	394	(315, 472)	0	(0, 0)	699	(617, 783)

We now discuss these results in detail. Our discussion focuses not only on quantitative metrics, but also potential paths and barriers to implementing the evaluated policies.

7.2.1 Increase Approaches

A central theme of our analysis is that incentivizing OPOs to approach more can substantially increase the number of donors. The “Approach All” and “Approach 20% More” policies both lead to significant gains, increasing the number of transplanted donors by 43% and 7% respectively. The associated societal benefit is extremely large; since the average donor saves 30.8 life-years for waitlisted patients (Schnitzler et al. 2005), “Approach All” would create an estimated 67,606 additional life-years. An added benefit of this policy is its operational simplicity. While correcting beliefs and improving authorization rates require interventions with uncertain success (as discussed below), “Approach All” only requires the OPO to identify absolute contraindications (e.g., active cancer; recall that our dataset already excludes any such referrals). Moreover, it eliminates any effects of inaccurate beliefs: if the OPO approaches all potential donors, its estimates of transplant probability do not affect its decisions.

The key barrier to implementing “Approach All” is cost, as it requires large increases in the number of approaches and authorizations. However, a simple calculation indicates that these increased costs would be outweighed by the policy’s financial benefits alone. Based on prior estimates in McCormick et al. (2022), a single organ donor saves taxpayers approximately \$277,000 in Medicare spending.⁵ Increasing the number of donors by 2,195 would thus generate \$608 million in savings.

⁵This estimate is based on taxpayer savings from kidney transplantation. From McCormick et al. (2022), a single

To quantify the associated costs in dollars, we can use our structural estimates of the proportional costs of procurement and approach (C_p, C_a) and an estimate of the per-donor OPO reward π . One ballpark estimate for π is \$105,455, which results from multiplying the average organ yield per-donor in our dataset⁶ with the standard acquisition charges reported in Ozbay et al. (2025). Combining these estimates, each additional approach costs $C_a\pi = \$190$ and each additional procurement costs $C_p\pi = \$42,374$. Thus, the additional 24,543 approaches and 8,757 authorizations required by “Approach All” increase OPO costs by \$378 million, which is only 62% of the \$608 million saved by Medicare. Framed another way, each additional approach would cost an OPO \$15,000 on average, but result in \$24,000 in taxpayer savings.

A key question remains: how can policymakers incentivize OPOs to drastically increase approach rates? One way to achieve this is through additional Medicare reimbursements. Currently, OPOs receive standard acquisition charges (SACs) from transplant centers only for transplanted organs. This creates a financial disincentive for OPOs to approach marginal donors (i.e., with low transplant probability). If instead, Medicare reimbursed all costs incurred during procurement, the OPO would be incentivized to approach more broadly. Our analysis suggests that such a policy is financially viable: for the four OPOs in our dataset, Medicare could have reimbursed \$190 per approach and \$42,000 per authorization, and still have generated net savings of nearly \$40 million per year. However, it is worth highlighting that a version of this reimbursement structure already exists. For kidneys specifically, Medicare reimburses OPOs for all procurement costs in excess of their received SACs using end-of-year reconciliation payments (Rosenberg et al. 2020, Ozbay et al. 2025). It is thus unclear whether extending this policy to other organs or supplementing it to add direct payments per-approach and per-authorization would be an effective tool. Further assessing the financial viability and potential impact of such policies is a key direction for future work.

Another path that can incentivize OPOs to increase approach rates is regulatory scrutiny. Recall that C_p does not necessarily reflect a true financial cost, but rather is an “as-if” cost implied by OPO decisions (i.e., OPOs act as if this was their procurement cost). These costs could be artificially inflated by a principal-agent problem (Ross 1973): OPO incentives may not be aligned with those of the system designer (i.e., CMS) or downstream beneficiaries (i.e., waitlisted patients), as historically

kidney transplant requires 1.5 kidneys and yields \$416,000 in savings; thus, the per-kidney benefit is $\approx \$277,000$. To be conservative, assume that each additional donor provides exactly one kidney (instead of the average of 3.3 organs / donor observed in our data).

⁶The average donor in our dataset provided 1.5 kidneys, 0.77 livers, 0.51 lungs, and 0.32 hearts.

there have been few repercussions for poor procurement. Because of these misaligned incentives, an OPO may set C_p higher than its actual financial costs would necessitate. The new OPO Final Rule (Centers for Medicare & Medicaid Services 2021) seeks to address this problem by penalizing OPOs with low donation rates with de-certification (i.e., revoking their license to conduct procurement). Although the Final Rule does not directly target approach rates, it incentivizes OPOs to increase the number of recovered donors, which may indirectly encourage more approaches. Since our data precedes the implementation of this policy, we cannot directly evaluate its impact. However, if the Final Rule fails to substantially increase donations, our results indicate that future reforms should explicitly encourage higher approach rates in addition to the current donation metrics.

We conclude this section with a note of caution. Increasing approach rates is only an advisable strategy if the additional approaches are both effective and ethical. For example, late or non-compassionate approaches may understandably lead families to decline consent (Siminoff et al. 2013, Pengel et al. 2025). More concerningly, there have been recent reports of OPOs compromising ethical standards, attempting to procure patients despite signs of increasing consciousness (Rosenthal 2025, Rosenthal and Tate 2025). Thus, as policymakers consider ways to incentivize more approaches, it is imperative that they implement and enforce safeguards that protect the safety and rights of potential donors and their families.

7.2.2 Correct Beliefs

Next, we evaluate the impact of correcting the OPO's beliefs. We find that only correcting the OPO's ranking (i.e., keeping approaches fixed) would have increased the number of transplanted donors by 10%. However, to fully realize the impact of correct beliefs, the OPO must also correct its underestimation of transplant probability: this would have led to a 21% increase in approaches and an 18% increase in transplanted donors.

While correcting both aspects of an OPO's beliefs is impactful, each requires a different intervention. Correcting average probability underestimation is more straightforward: it simply requires an increase in the overall approach rate. Conceptually, increasing the OPO's estimate of β_0 has a very similar effect on the approach decision as reducing the OPO's cost C_p (i.e., both increase the average approach rate). Thus, the policy interventions from the previous section will partially realize the benefits of this correction, even if they do not explicitly target beliefs.

In contrast, correcting the misranking of potential donors is more difficult. One potential avenue is

the development of a decision-support tool that provides real-time estimates of transplant probability (using our estimated β_0^* and β^*). However, implementing decision-support tools in clinical settings is notoriously challenging. For example, a recent evaluation of a widely deployed sepsis prediction model found that it performed no better than random chance at predicting sepsis onset in time to influence treatment decisions (Kamran et al. 2024). Developing an effective decision-support system would require overcoming a range of human, technological, and organizational barriers (Meunier et al. 2023). While developing and deploying such tools is beyond the scope of this paper, it represents a clear direction for future research. In the meantime, we recommend that immediate policy efforts target approach rates: not only are such policies easier to implement, but they will also implicitly realize some of the benefit of correcting beliefs.

7.2.3 Improving Authorization

Finally, we evaluate the impact of improving authorization rates from approached families. We find that uniformly raising the authorization rate by 5 percentage points yields the smallest gains of any policy (a 5% increase in transplanted donors). However, eliminating racial gaps in authorization yields notably larger gains, increasing the number of transplanted donors by 8%. These gains are particularly cost-effective, as they require no additional approaches. This finding suggests that interventions focused explicitly on improving OPO interactions with minority families can yield large benefits. Of course, reducing racial gaps in consent is a complex challenge: these disparities have persisted for decades (Rodrigue et al. 2006, Goldberg et al. 2013, Levan et al. 2025) and partially stem from broader medical mistrust that cannot be resolved by OPOs alone (Siminoff et al. 2006). However, there are still concrete actions OPOs can take to improve consent rates from minority families. Prior work highlights community outreach, culturally targeted interventions, improved staff training, and racially concordant approaches as promising strategies (Zaramo et al. 2008, Hong et al. 1994, Baughn et al. 2010, Andrews et al. 2012, Council et al. 2022). Determining which of these strategies is most effective is beyond the scope of this paper; however, our results highlight the potential impact of such interventions.

8 Conclusion

In this article, we have investigated how decisions made by organ procurement organizations affect the availability of transplantable organs in the US. Using a novel econometric model and data on four OPOs, we studied the decision to approach a potential donor’s family to obtain their consent. Our analysis identified several structural barriers that have hindered organ procurement, including conservative approach thresholds, miscalibrated OPO beliefs, and limited ability to obtain family consent. Through counterfactual policy evaluations, we identified several changes to OPO processes that can increase the supply of donor organs. To our knowledge, our work is the first analysis of OPO behavior that uses rigorous empirical tools and granular donor-level data.

Our work has several limitations that suggest directions for future research. First, our model makes several assumptions, including the bivariate normality of latent indices, correct beliefs of authorization likelihood, and calibrated beliefs of transplant likelihood. Relaxing our assumptions while still separating beliefs and preferences and controlling for sample selection is a key future direction. Second, our model is identified by an instrument that affects the OPO’s decision to approach, but not transplant likelihood. While we present strong evidence that our FPA instrument satisfies the first condition on relevance, the latter exclusion restriction is untestable. Third, due to data availability, our analysis focuses only on four OPOs over a six-year span (2016-2021). Expanding our analysis to a larger number of OPOs and more recent years is an interesting future direction, especially given recent regulatory changes like the new Final Rule (CMS 2020). Finally, while our policy evaluation highlights specific processes to target, it does not determine how these changes should be operationalized. Future work can build upon our analysis by evaluating specific policy implementations, focusing both on internal managerial changes and federal regulations.

References

- Abaluck J, Agha L, Kabrhel C, Raja A, Venkatesh A (2016) The determinants of productivity in medical testing: Intensity and allocation of care. *American Economic Review* 106(12):3730–3764.
- Adam H, Pollard T, Suriyakumar V, Moody B, Adams JN, Erickson J, Segal G, Wadsworth M, Wilson A, Ghassemi M (2026) Organ retrieval and collection of health information for donation: The orchid dataset. *Scientific Data* .

- Agarwal N, Ashlagi I, Azevedo E, Featherstone CR, Karaduman Ö (2019) Market failure in kidney exchange. *American Economic Review* 109(11):4026–4070.
- Agarwal N, Ashlagi I, Rees MA, Somaini P, Waldinger D (2021) Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica* 89(1):37–76.
- Agarwal N, Hodgson C, Somaini P (2025) Choices and outcomes in assignment mechanisms: The allocation of deceased donor kidneys. *Econometrica* 93(2):395–438.
- Aki V, Andrew G, Daniel S, Bob C, Paul-Christian B (2020) Rank-normalization, folding, and localization: An improved rhat for assessing convergence of mcmc. *Bayesian Analysis* 6:16.
- Albert JH, Chib S (1993) Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association* 88(422):669–679.
- Allon G, Federgruen A, Pierson M (2011) How much is a reduction of your customers’ wait worth? an empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing & Service Operations Management* 13(4):489–507.
- Andrews AM, Zhang N, Magee JC, Chapman R, Langford AT, Resnicow K (2012) Increasing donor designation through black churches: results of a randomized trial. *Progress in transplantation* 22(2):161–167.
- Baughn D, Auerbach SM, Siminoff LA (2010) Roles of sex and ethnicity in procurement coordinator—family communication during the organ donation discussion. *Progress in Transplantation* 20(3):247–255.
- Berrevoets J, Jordon J, Bica I, van der Schaar M, et al. (2020) Organite: Optimal transplant donor organ offering using an individual treatment effect. *Advances in neural information processing systems* 33:20037–20050.
- Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.
- Bravo F, Gandhi A, Hu J, Long EF (2025) Closer to home: A structural estimate-then-optimize approach to improve access to healthcare services. *Management Science* .
- Brown CV, Foulkrod KH, Dworaczyk S, Thompson K, Elliot E, Cooper H, Coopwood B (2010) Barriers to obtaining family consent for potential organ donors. *Journal of Trauma and Acute Care Surgery* 68(2):447–451.
- Carola V, Morale C, Vincenzo C, Cecchi V, Errico L, Nicolais G (2023) Organ donation: psychosocial factors of the decision-making process. *Frontiers in psychology* 14:1111328.
- Centers for Medicare & Medicaid Services (2021) Medicare and medicaid programs: Organ procurement organizations conditions for coverage: Revisions to the outcome measure requirements for or-

- gan procurement organization; final rule URL <https://www.cms.gov/files/document/112020-opo-final-rule-cms-3380-f.pdf>.
- Chan CW, Farias VF, Bambos N, Escobar GJ (2012) Optimizing intensive care unit discharge decisions with patient readmissions. *Operations research* 60(6):1323–1341.
- Chan DC, Gentzkow M, Yu C (2022) Selection with variation in diagnostic skill: Evidence from radiologists. *The Quarterly Journal of Economics* 137(2):729–783.
- Chandra A, Staiger DO (2020) Identifying sources of inefficiency in healthcare. *The Quarterly Journal of Economics* 135(2):785–843.
- CMS (2020) Organ procurement organizations. conditions for coverage: Revisions to the outcome measure requirements for organ procurement organizations; final rule. in 42 cfr part 486, washington, dc, 2020.
- Cohen JB, Eddinger KC, Locke JE, Forde KA, Reese PP, Sawinski DL (2017) Survival benefit of transplantation with a deceased diabetic donor kidney compared with remaining on the waitlist. *Clinical Journal of the American Society of Nephrology* 12(6):974–982.
- Collett D, Friend PJ, Watson CJ (2017) Factors associated with short-and long-term liver graft survival in the united kingdom: development of a uk donor liver index. *Transplantation* 101(4):786–792.
- Council NR, et al. (2022) *Realizing the promise of equity in the organ transplantation system*.
- Currie J, MacLeod WB (2017) Diagnosing expertise: Human capital, decision making, and performance among physicians. *Journal of labor economics* 35(1):1–43.
- Dageforde LA, Muren W, Chang SH, Vachharajani N, Brockmeier D, Yu J, Anderson B, Shenoy S, Lin Y, Khan A, et al. (2020) Organ procurement organization run department of motor vehicle registration and drivers licensing offices leads to increased organ donor first person authorization registrations. *Transplantation* 104(2):343–348.
- Dai T, Zheng R, Sycara K (2020) Jumping the line, charitably: Analysis and remedy of donor-priority rule. *Management Science* 66(2):622–641.
- DeRoos LJ, Zhou Y, Marrero WJ, Tapper EB, Sonnenday CJ, Lavieri MS, Hutton DW, Parikh ND (2021) Assessment of national organ donation rates and organ procurement organization metrics. *JAMA surgery* 156(2):173–180.
- Dickert-Conlin S, Elder T, Lemont B, Teltser K (2024) Opioids and organs: How overdoses affect the supply and demand for organ transplants. *American Journal of Health Economics* 10(4):505–538.
- Dickert-Conlin S, Elder T, Moore B (2011) Donorcycles: Motorcycle helmet laws and the supply of organ donors. *The Journal of Law and Economics* 54(4):907–935.

- Dickert-Conlin S, Elder T, Teltser K (2019) Allocating scarce organs: How a change in supply affects transplant waiting lists and transplant recipients. *American Economic Journal: Applied Economics* 11(4):210–239.
- Doby BL, Ross-Driscoll K, Shuck M, Wadsworth M, Durand CM, Lynch RJ (2021) Public discourse and policy change: Absence of harm from increased oversight and transparency in opo performance. *American Journal of Transplantation* 21(8):2646–2652.
- Doby BL, Ross-Driscoll K, Yu S, Godwin M, Lee KJ, Lynch RJ (2022) Examining utilization of kidneys as a function of procurement performance. *American Journal of Transplantation* .
- Fawcett T (2006) An introduction to roc analysis. *Pattern recognition letters* 27(8):861–874.
- Feeley TH, Anker AE, Evans M, Reynolds-Tylus T (2017) A department of motor vehicle–based intervention to promote organ donor registrations in new york state. *Progress in Transplantation* 27(3):273–280.
- Fleisher L, Moody-Williams J, Blum J (2023) The transplant eco-system: the role of data in cms oversight of the organ procurement organizations. *Centers for Medicare & Medicaid Services*. <https://www.cms.gov/blog/transplant-eco-system-role-data-cms-oversight> .
- Freeman M, Robinson S, Scholtes S (2021) Gatekeeping, fast and slow: An empirical study of referral errors in the emergency department. *Management Science* 67(7):4209–4232.
- Goldberg DS, Doby B, Lynch R (2020) Addressing critiques of the proposed cms metric of organ procurement organ performance: more data isn’t better. *Transplantation* 104(8):1662–1667.
- Goldberg DS, Halpern SD, Reese PP (2013) Deceased organ donation consent rates among racial and ethnic minorities and older potential donors. *Critical care medicine* 41(2):496–505.
- He P, Zheng F, Belavina E, Girotra K (2021) Customer preference and station network in the london bike-share system. *Management Science* 67(3):1392–1412.
- Heckman JJ (1979) Sample selection bias as a specification error. *Econometrica: Journal of the econometric society* 153–161.
- Held PJ, Bragg-Gresham JL, Peters T, Chertow GM, McCormick F, Roberts JP (2020) The cost of procuring deceased donor kidneys: Evidence from opo cost reports 2013-2017. *American Journal of Transplantation* 20(4):1087–1094.
- Hoffman MD, Gelman A, et al. (2014) The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.* 15(1):1593–1623.
- Hong BA, Kappel DF, Whitlock MJ, Parks-Thomas T, Freedman B (1994) Using race-specific community programs to increase organ donation among blacks. *American Journal of Public Health* 84(2).

- Hu Y, Chan CW, Dong J (2025) Prediction-driven surge planning with application to emergency department nurse staffing. *Management Science* 71(3):2079–2126.
- Hulme W, Allen J, Manara AR, Murphy PG, Gardiner D, Poppitt E (2016) Factors influencing the family consent rate for organ donation in the uk. *Anaesthesia* 71(9):1053–1063.
- Iltis AS, Denny B (2025) First-person authorization and family objections to organ donation. *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine*, jhaf008 (Oxford University Press US).
- Johnson W, Kraft K, Chotai P, Lynch R, Dittus RS, Goldberg D, Ye F, Doby B, Schaubel DE, Shah MB, Karp SJ (2023) Variability in Organ Procurement Organization Performance by Individual Hospital in the United States. *JAMA Surgery* 158(4):404–409, ISSN 2168-6254, URL <http://dx.doi.org/10.1001/jamasurg.2022.7853>.
- Kamran F, Tjandra D, Heiler A, Virzi J, Singh K, King JE, Valley TS, Wiens J (2024) Evaluation of sepsis prediction models before onset of treatment. *Nejm Ai* 1(3):AIoa2300032.
- Kc DS, Terwiesch C (2009) Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management science* 55(9):1486–1498.
- Kc DS, Terwiesch C (2011) The effects of focus on performance: Evidence from california hospitals. *Management Science* 57(11):1897–1912.
- Kc DS, Terwiesch C (2012) An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing & Service Operations Management* 14(1):50–65.
- Kim SH, Chan CW, Olivares M, Escobar G (2015) Icu admission control: An empirical study of capacity allocation and its implication for patient outcomes. *Management Science* 61(1):19–38.
- Kim SH, Tong J, Peden C (2020) Admission control biases in hospital unit capacity management: How occupancy information hurdles and decision noise impact utilization. *Management Science* 66(11):5151–5170.
- Kim SH, Zheng F, Brown J (2024) Identifying the bottleneck unit: Impact of congestion spillover in hospital inpatient unit network. *Management Science* 70(7):4200–4218.
- Kim WR, Mannalithara A, Heimbach JK, Kamath PS, Asrani SK, Biggins SW, Wood NL, Gentry SE, Kwong AJ (2021) Meld 3.0: the model for end-stage liver disease updated for the modern era. *Gastroenterology* 161(6):1887–1895.
- Kong N, Schaefer AJ, Hunsaker B, Roberts MS (2010) Maximizing the efficiency of the us liver allocation system through region design. *Management Science* 56(12):2111–2122.

- Lee C, Tsai C, Adler A, et al. (2022) Emerging frontier in heart transplantation: donation after circulatory death. american college of cardiology expert analysis, november 21, 2022.
- Lee DK, Zenios SA (2012) An evidence-based incentive system for medicare's end-stage renal disease program. *Management science* 58(6):1092–1105.
- Levan ML, Massie AB, Trahan C, Hewlett J, Strout T, Klitenic SB, Vanterpool KB, Segev DL, Adams BL, Niles P (2022) Maximizing the use of potential donors through increased rates of family approach for authorization. *American Journal of Transplantation* 22(12):2834–2841.
- Levan ML, Terlizzi K, Rigsby M, Klitenic S, Hewlett J, Adams BL, Barnes J, Funk G, Segev DL, Massie AB (2025) Revisiting racial/ethnic disparities in the deceased organ donor referral process. *American Journal of Transplantation* .
- Li J, Granados N, Netessine S (2014) Are consumers strategic? structural estimation from the air-travel industry. *Management Science* 60(9):2114–2137.
- Li M, Riyanto YE (2025) Incentivizing organ donation under different priority rules: The role of information. *Management Science* 71(2):1418–1437.
- Lopez R, Mohan S, Schold JD (2023) Population characteristics and organ procurement organization performance metrics. *JAMA Network Open* 6(10):e2336749–e2336749.
- Lynch RJ, Doby BL, Goldberg DS, Lee KJ, Cimeno A, Karp SJ (2022) Procurement characteristics of high- and low-performing opos as seen in optn/srtr data. *American Journal of Transplantation* 22(2):455–463.
- Mandelbaum A, Momčilović P, Trichakis N, Kadish S, Leib R, Bunnell CA (2020) Data-driven appointment-scheduling under uncertainty: The case of an infusion unit in a cancer center. *Management Science* 66(1):243–270.
- Matzkin RL (1992) Nonparametric and distribution-free estimation of the binary threshold crossing and the binary choice models. *Econometrica: Journal of the Econometric Society* 239–270.
- McCormick F, Held PJ, Chertow GM, Peters TG, Roberts JP (2022) Projecting the economic impact of compensating living kidney donors in the united states: cost-benefit analysis demonstrates substantial patient and societal gains. *Value in Health* 25(12):2028–2033.
- Meunier PY, Raynaud C, Guimaraes E, Gueyffier F, Letrilliart L (2023) Barriers and facilitators to the use of clinical decision support systems in primary care: a mixed-methods systematic review. *The Annals of Family Medicine* 21(1):57–69.
- Musalem A, Olivares M, Bradlow ET, Terwiesch C, Corsten D (2010) Structural estimation of the effect of out-of-stocks. *Management Science* 56(7):1180–1197.

- Olivares M, Terwiesch C, Cassorla L (2008) Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Science* 54(1):41–55.
- OPTN (2018) A guide to calculating and interpreting the kidney donor profile index (kdpi). URL https://optn.transplant.hrsa.gov/media/j34dm4mv/kdpi_guide.pdf.
- OPTN (2025) Optn policies effective as of june 26, 2025. URL https://optn.transplant.hrsa.gov/media/eavh5bf3/optn_policies.pdf, accessed on Jul 14, 2025.
- Ozbay E, Rava A, Zur E (2025) Lost in transplantation: Cost shifting in publicly subsidized programs. *Available at SSRN 5187593*.
- Papalexopoulos T, Alcorn J, Bertsimas D, Goff R, Stewart D, Trichakis N (2024) Reshaping national organ allocation policy. *Operations Research* 72(4):1475–1486.
- Pengel L, Mazarello-Paes V, Paredes-Zapata D, Oniscu G, Gaglianone CG, Zhu L, Wang Y, Dhanda N, Tocher J, Aviles L (2025) The experiences of clinical staff approaching families for organ donation consent: A systematic review and thematic synthesis of qualitative studies. *Transplantation Reviews* 39(3):100935.
- Platt J, et al. (1999) Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers* 10(3):61–74.
- Qin Y, Imrie F, Hüyük A, Jarrett D, van der Schaar M, et al. (2021) Closing the loop in medical decision support by understanding clinical decision-making: A case study on organ transplantation. *Advances in Neural Information Processing Systems* 34.
- Rao PS, Schaubel DE, Guidinger MK, Andreoni KA, Wolfe RA, Merion RM, Port FK, Sung RS (2009) A comprehensive risk quantification score for deceased donor kidneys: the kidney donor risk index. *Transplantation* 88(2):231–236.
- Rath S, Rajaram K (2022) Staff planning for hospitals with implicit cost estimation and stochastic optimization. *Production and Operations Management* 31(3):1271–1289.
- Rodrigue JR, Cornell DL, Howard R (2006) Organ donation decision: comparison of donor and nondonor families. *American journal of Transplantation* 6(1):190–198.
- Rosenberg P, Ciccarone M, Seeman B, Washer D, Martin G, Tse J, Segal G (2020) Transforming organ donation in america. *Bridgespan*.
- Rosenthal BM (2025) Doctors were preparing to remove their organs. then they woke up. *The New York Times*. Accessed on Aug 12, 2025 at <https://www.nytimes.com/2025/06/06/us/kentucky-organ-donations.html>.
- Rosenthal BM, Tate J (2025) A push for more organ transplants is putting donors at risk. *The New*

- York Times*. Accessed on Aug 12, 2025 at <https://www.nytimes.com/2025/07/20/us/organ-transplants-donors-alive.html>.
- Ross SA (1973) The economic theory of agency: The principal's problem. *The American economic review* 63(2):134–139.
- Schnitzler MA, Whiting JF, Brennan DC, Lentine KL, Desai NM, Chapman W, Kalo Z, et al. (2005) The life-years saved by a deceased organ donor. *American Journal of Transplantation* 5(9):2289–2296.
- Shen Y, Chan CW, Zheng F, Argenziano M, Kurlansky P (2021) The impact of surgeon daily workload and its implications for operating roomscheduling. *Production and Operations Management* 10591478251398334.
- Shmueli A, Sprung CL (2005) Assessing the in-hospital survival benefits of intensive care. *International journal of technology assessment in health care* 21(1):66–72.
- Siminoff LA, Agyemang AA, Traino HM (2013) Consent to organ donation: a review. *Progress in Transplantation* 23(1):99–104.
- Siminoff LA, Burant CJ, Ibrahim SA (2006) Racial disparities in preferences and perceptions regarding organ donation. *Journal of general internal medicine* 21(9):995–1000.
- Siminoff LA, Gordon N, Hewlett J, Arnold RM (2001) Factors influencing families' consent for donation of solid organs for transplantation. *Jama* 286(1):71–77.
- Simpkin AL, Robertson LC, Barber VS, Young JD (2009) Modifiable factors influencing relatives' decision to offer organ donation: systematic review. *Bmj* 338.
- Stan Development Team (2024) Stan modeling language users guide and reference manual, version 2.32. URL https://mc-stan.org/docs/2_36/reference-manual/index.html.
- Su X, Zenios SA (2005) Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations research* 53(3):443–455.
- Su X, Zenios SA (2006) Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management science* 52(11):1647–1660.
- United States Congress (1984) S.2048 - national organ transplant act .
- Van de Ven WP, Van Praag BM (1981) The demand for deductibles in private health insurance: A probit model with sample selection. *Journal of econometrics* 17(2):229–252.
- Van der Vaart AW (2000) *Asymptotic statistics*, volume 3 (Cambridge university press).
- Wooldridge JM (2010) *Econometric analysis of cross section and panel data* (MIT press).

- Yankovic N, Green LV (2011) Identifying good nursing levels: A queuing approach. *Operations research* 59(4):942–955.
- Zaramo C, Morton T, Yoo J, Bowen G, Modlin C (2008) Culturally competent methods to promote organ donation rates among african-americans using venues of the bureau of motor vehicles. *Transplantation proceedings*, volume 40, 1001–1004 (Elsevier).

A Theoretical Details

A.1 Deriving Transplant Probability

We derive equation (2) for $\mathbb{P}(y_i^* = 1 \mid x_i, \eta_i)$, the true probability of transplant conditional on the information observed by the OPO. By Thm. 1, (η_i, ϵ_i) follows a bivariate normal distribution with unit variance and correlation ρ^* . By the well-known conditional property of the bivariate normal, we know that $\epsilon_i \mid \eta_i \sim \mathcal{N}(\rho^* \eta_i, 1 - \rho^{*2})$. Thus,

$$\begin{aligned}
 \mathbb{P}(y_i^* = 1 \mid x_i, \eta_i) &= \mathbb{P}(\beta_0^* + x_i^T \beta^* + \epsilon_i > 0 \mid x_i, \eta_i) \\
 &= \mathbb{P}(\epsilon_i > -\beta_0^* - x_i^T \beta^* \mid x_i, \eta_i) \\
 &= 1 - \mathbb{P}(\epsilon_i \leq -\beta_0^* - x_i^T \beta^* \mid x_i, \eta_i) \\
 &= 1 - F_{\epsilon_i \mid \eta_i}(-\beta_0^* - x_i^T \beta^*) \\
 &= 1 - \Phi\left(\frac{-\beta_0^* - x_i^T \beta^* - \rho^* \eta_i}{\sqrt{1 - \rho^{*2}}}\right) \\
 &= \Phi\left(\frac{\beta_0^* + x_i^T \beta^* + \rho^* \eta_i}{\sqrt{1 - \rho^{*2}}}\right).
 \end{aligned}$$

where $F_{\epsilon_i \mid \eta_i}$ is the cumulative distribution function (CDF) of ϵ_i conditional on η_i and the last step follows from the symmetry of the normal distribution.

A.2 Deriving the Optimal Rule

To derive the optimal decision rule (6), we consider the expected utility of an approach conditional on the OPO's observed information (x_i, z_i, η_i) . Using the law of iterated expectations with the OPO's utility in (5),

$$\begin{aligned}
 \mathbb{E}[u_i \mid r_i = 1, x_i, z_i, \eta_i] &= -\pi C_a \\
 &\quad - \pi C_p \mathbb{P}(a_i^* = 1, y_i^* = 0 \mid x_i, \eta_i, z_i) \\
 &\quad + \pi(1 - C_p) \mathbb{P}(a_i^* = 1, y_i^* = 1 \mid x_i, \eta_i, z_i)
 \end{aligned}$$

By Thm. 2 and Thm. 3, $\mathbb{P}(a_i^*, y_i^* | x_i, z_i, \eta_i) = \mathbb{P}(a_i^* | x_i, z_i) \mathbb{P}(y_i^* | x_i, \eta_i)$. Thus, the expression simplifies to

$$\begin{aligned} \mathbb{E}[u_i | r_i = 1, x_i, z_i, \eta_i] \\ = \pi \left[-C_a - C_p \kappa(x_i, z_i) + \kappa(x_i, z_i) \mathbb{P}(y_i^* = 1 | x_i, \eta_i) \right] \end{aligned}$$

The OPO should approach if this expected utility exceeds that of a non-approach, which is just 0 (from (5)). Setting $\mathbb{E}[u_i | r_i = 1, x_i, z_i, \eta_i] > 0$ and algebraically simplifying yields the decision rule (6).

A.3 Calibration Moment Conditions

Each calibration condition in Thm. 5 implies a moment condition involving our structural parameters. Here, we derive these conditions.

A.4 Aggregate Calibration

Recall the aggregate calibration condition,

$$\hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid a_i = 1) = \mathbb{P}(y_i^* = 1 \mid a_i = 1).$$

Observe that the right hand-side of this equation can directly be identified from data as the average proportion of authorized donors that result in transplants, $\sum_i y_i / \sum_i a_i$. Thus, we only need to express the left hand side in terms of our structural parameters. We write,

$$\begin{aligned} & \hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid a_i = 1) \\ &= \mathbb{E}[\hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid a_i = 1, x_i, z_i, \eta_i) \mid a_i = 1] \\ & \quad \text{(law of iterated expectation)} \\ &= \mathbb{E}[\hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid r_i = 1, a_i^* = 1, x_i, z_i, \eta_i) \mid a_i = 1] \\ & \quad \text{(since } a_i = r_i a_i^*) \\ &= \mathbb{E}[\hat{\mathbb{P}}^{OPO}(y_i^* = 1 \mid x_i, z_i, \eta_i) \mid a_i = 1] \\ & \quad \text{(since } a_i^* \perp\!\!\!\perp y_i^* \mid x_i, z_i, \\ & \quad \text{and } r_i \text{ is deterministic conditional on } x_i, z_i, \eta_i) \\ &= \mathbb{E}\left[\Phi\left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}}\right) \mid a_i = 1\right] \\ & \quad \text{(from equation (3))} \\ &= \mathbb{E}\left[\Phi\left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}}\right) \mid a_i = 1\right]. \end{aligned}$$

Since this term is an expectation over x_i and η_i , it is a function only of the unknown parameters (β_0, β, ρ) . Thus, if the distributions of x_i and η_i conditional on $a_i = 1$ are identified, then expectations over these distributions are also identified, and our calibration condition yields a well-defined

moment condition. We first note that $x_i \mid a_i = 1$ is directly identified from the observed distribution of covariates for authorized potential donors. Thus, all that remains is to establish that the distribution of $\eta_i \mid a_i = 1$ is identified. This follows directly from the expression of the OPO's approach rule as a binary choice model. Recall from (8) that we can express the OPO's decision rule as

$$r_i = \mathbf{1}[m(x_i, z_i) + \eta_i > 0]$$

for a partially-linear function m . From standard binary choice theory (Matzkin 1992), we know that m is identified from the observed outcomes r_i . Thus, for some constant b , the conditional CDF of η_i is

$$\begin{aligned} \mathbb{P}(\eta_i \leq b \mid a_i = 1) &= \mathbb{P}(\eta_i \leq b \mid r_i = 1, a_i^* = 1) \\ &= \mathbb{P}(\eta_i \leq b \mid r_i = 1) \\ &\quad (\text{since } \eta_i \perp\!\!\!\perp a_i^* \text{ by Thm. 3}) \\ &= \mathbb{P}(\eta_i \leq b \mid \eta_i > -m(x_i, z_i)). \end{aligned}$$

Since η_i is normally distributed, this defines a truncated normal random variable with mean 0, variance 1, and a lower truncation bound of $-m(x_i, z_i)$. Thus, since $m(x_i, z_i)$ is identified, the distribution of $\eta_i \mid a_i = 1$ is identified.

In summary, we can express our calibration condition as

$$\mathbb{E} \left[\Phi \left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}} \right) \mid a_i = 1 \right] = \frac{\sum_{i=1}^N y_i}{\sum_{i=1}^N a_i}$$

Given fixed values of (β_0, β, ρ) , we can approximate the left hand side by sampling from the (identified) distributions of $x_i, \eta_i \mid a_i = 1$. Thus, this defines a moment condition with our structural parameters.

A.5 Near-Threshold Calibration

The near-threshold calibration condition identifies a linear combination of the cost parameters (C_a, C_p) . Recall that we defined the OPO's believed probability of transplant as $\hat{p}_i = \hat{\mathbb{P}}^{OPO}(y_i^* \mid x_i, \eta_i)$, the OPO's approach threshold $\tau_i = C_p + \frac{C_a}{\kappa(x_i, z_i)}$, and the set of marginally approached donors

as $\mathcal{M} = \{i : \hat{p}_i \approx \tau_i, r_i = 1\}$. The near-threshold calibration condition implies that

$$\mathbb{E}[y_i \mid i \in \mathcal{M}] = C_p + C_a \cdot \mathbb{E}\left[\frac{1}{\kappa(x_i, z_i)} \mid i \in \mathcal{M}\right].$$

Because y_i is observed and $\kappa(x_i, z_i)$ is known, this provides a linear moment condition in (C_p, C_a) , provided $\mathcal{M}$ is identified. Although $\hat{p}_i$ is unobserved, we can identify $\mathcal{M}$ using a limit argument. Consider potential donors whose observables indicate that they were extremely unlikely to be approached (i.e., $\mathbb{P}(r_i = 1 \mid x_i, z_i) \approx 0$), but who were approached anyway (i.e., $r_i = 1$). Given their low predicted probability of approach, it is likely that the unobservable η_i for each potential donor was just large enough to push them over the threshold. More formally, if $r_i = 1$ and $\mathbb{P}(r_i = 1 \mid x_i, z_i) \rightarrow 0$, we expect that $\hat{p}_i \rightarrow \tau_i$. Thus, $\mathcal{M}$ is approximated by the set $\{i : r_i = 1, \Pr(r_i = 1 \mid x_i, z_i) \approx 0\}$. The near-threshold calibration condition thus identifies a linear moment in (C_p, C_a) .

Abaluck et al. (2016) uses a similar near-threshold technique to identify physician decision thresholds for ordering medical tests. Their model directly assumes that physicians have true knowledge of their intercept β_0^* , which identifies thresholds using a calibration condition that is similar to ours. However in our setting, assuming knowledge of β_0^* prevents us from distinguishing between two contributors to potential OPO under-performance: (1) incorrect beliefs of the average rate of transplantable donors, and (2) overly-high OPO thresholds. We thus do not assume knowledge of β_0^* , instead directly assuming the weaker near-threshold calibration condition.

B Empirical Details

B.1 Covariates and Specification

Here, we discuss the covariates and specifications used in our structural model. Our set of x_i included the following donation-relevant features:

- **Referral Year.** We include a fixed effect for each year with 2016 as the reference level.
- **Age.** We control for age using a piece-wise linear specification. Our specification matches that used in the Kidney Donor Profile Index (KDPI) (OPTN 2018), a clinically-validated measure of the quality of a donor kidney. We control for three age-related variables: (1) a linear control for age for all potential donors, (2) a linear control for age greater than 50, and (3) a linear control for age less than 18.
- **Sex.** We control for sex using a binary indicator for female (1) or male (0).
- **Race.** We control for race as a categorical variable with four levels: White / Caucasian (reference level), Black / African American, Hispanic, and Other / Unknown.
- **Brain Death.** We include a binary indicator for whether the potential donor suffered brain death (1) or cardiac death (0).
- **Cause of Death.** We control for cause of death as a categorical variable with five levels: Anoxia (reference level), CNS Tumor, CVA/Stroke, Head Trauma, Other / Unknown. Note that we do not control for cause of death for OPO A, as it was typically not recorded unless a potential donor's family was approached.
- **Laboratory test results.** We control for the following lab tests of kidney / liver function. Note that these values were determined using both structured data and free text notes written by OPO staff during donor evaluation.
 - **Creatinine.** Serum creatinine conveys kidney function and is a crucial marker for organ donation. Our piecewise-linear specification matches that used in the KDPI (OPTN 2018): (1) a linear control for Creatinine, and (2) a linear control for Creatinine > 1.5 mg/dl. In line with KDPI, we cap high values of creatinine at 8 mg/dl (OPTN 2018). In

addition to structured data, creatinine values were extracted from free-text OPO notes using regular expressions. This rules-based approach looked for mentions of creatinine keywords (e.g., “cr”, “creat”, “creatinine”) followed by a number (e.g., “Cr 1.5”).

- **Total Bilirubin** Total bilirubin is an important indicator of liver function. In line with the Model for End Stage Liver Disease (MELD)—a clinically-validated measure of the severity of liver disease—we controlled for the natural log of total bilirubin (Kim et al. 2021). In addition to structured data, bilirubin values were extracted from free-text OPO notes using regular expressions. This rules-based approach looked for mentions of bilirubin keywords (e.g., “TBili”, “Total Bili”) followed by a number (e.g., “TBili 1.5”).
- **Aspartate aminotransferase (AST) and alanine aminotransferase (ALT).** AST and ALT are important markers of liver function. Because these variables were highly correlated and right skewed with a long tail, we controlled for the log of the maximum of AST and ALT. In line with clinical expertise, we also added binary indicators for high AST or ALT (≥ 100) and extremely high AST or ALT (≥ 1000). In addition to structured data, AST and ALT values were extracted from free-text OPO notes using regular expressions. This rules-based approach looked for mentions of the keywords “AST” and “ALT” followed by a number (e.g., “AST 43”).
- **Missingness Indicators.** Lab test results were occasionally not recorded in either the structured data or free-text notes. In our dataset, 8% of potential donors were missing creatinine, 21% were missing total bilirubin, and 20% were missing AST and ALT. To control for non-random patterns, we added binary indicators for missingness of each of these variables, then zero-imputed the main value. For example, if creatinine was missing, we set the creatinine missing indicator to 1 and assigned the main creatinine features as 0. We add these missingness indicators for all three labs.
- **Comorbidities and Treatments.** Finally, we controlled for a number of donation-relevant conditions and treatments. The set of relevant conditions (and associated keywords) were identified in collaboration with a transplant physician. Since these variables were not recorded in the structured data, they were imputed using free-text progress notes made by OPO staff to describe a potential donor’s condition. Each variable was imputed by using regular

expressions (regex) to identify the keywords listed below. All variables are binary indicators with 1 indicating presence of the condition / treatment and 0 indicating the absence.

- **Diabetes:** “diabetes”, “diabetic”, “dm”
- **Hypertension:** “htn”, “hypertension”
- **Sepsis:** “sepsis”, “septic”
- **End-stage Renal Disease (ESRD):** “esrd”
- **Chronic Kidney Disease (CKD):** “ckd”
- **Liver cirrhosis:** “cirrhosis”
- **Chronic Obstructive Pulmonary Disease (COPD):** “copd”
- **Atrial fibrillation:** “afib”
- **Continuous renal replacement therapy (CRRT):** “crrt”, “cvvh”
- **Dialysis:** “dialysis”, “on hd”
- **Extracorporeal Membrane Oxygenation(ECMO):** “ecmo”
- **Use of vasopressors:** We infer vasopressor usage with two conditions: (1) mention of one of the following vasopressors: vasopressin (“vaso”, “vasopress”), norepinephrine (“levo”, “levophed”, “norepinephrine”, “norepi”), epinephrine (“epi”, “epinephrine”), dopamine (“dopa”, “dopamine”), or phenylephrine (“neo”, “neosyneph”, “neo-syneph”), and (2) no mention of the phrase “no pressors”

B.2 Estimating Authorization Probability

Data. To train our models of authorization, we focused only on the set of potential donors whose families were approached. In line with prior work (e.g., Goldberg et al. (2013), Levan et al. (2022)), our model of authorization used the following variables: FPA, age, race, sex, brain death, and referral year. We controlled for age as a categorical variable with six levels: 0-17, 18-34, 35-44, 45-54, 55-64, and 65+.

Modeling Approach. We use a two-step approach to estimate authorization probability. This approach requires (1) a fixed set of training observations, and (2) a regularization hyper-parameter λ . Given these two components, we predict authorization probability in two steps:

1. **Logistic Regression** We first use 70% of training observations to fit a logistic regression of authorization outcomes on the observed covariates. This logistic regression is regularized using the L1 penalty with strength λ .
2. **Calibration** We then use the remaining 30% training observations to calibrate predictions. We fit a generalized additive model (GAM) of authorization outcomes to the predicted probabilities from the logistic regression. The GAM models the predicted probabilities as a spline with 4 knots. This approach is similar to Platt scaling (Platt et al. 1999), except that we use a GAM for calibration instead of a logistic transform.

We can then predict authorization probability for a given test data point (i.e., *not* in the training set) by first passing it through the logistic regression, then passing the logistic regression output through the calibration-GAM. This two-step approach yields authorization probabilities that are not only predictive, but also well-calibrated. Such calibration is important, since the predicted probabilities are inputs into our structural model.

Generating Predictions. We estimate authorization probability for all potential donors as follows:

1. **Pick Regularization Strength.** First, we identify the best L1-regularization parameter λ using 10-fold cross-validation. For a given λ , we split the data into ten subsets (i.e., “folds”). To obtain predicted authorization probabilities for each fold, we use a calibrated model (as described above) trained on the remaining nine folds. Once we have obtained predictions for all ten folds, we evaluate accuracy using AUROC. We repeat this process for 30 values of the regularization parameter between 0.0005 and 0.02, and choose the λ that maximizes AUROC.
2. **Predict Using Best Parameter.** We then use our modeling approach with the best λ to generate final estimates of authorization probability. For potential donors who were approached (i.e., used to train the models of authorization), we generate predictions using 10-fold cross-validation. This approach ensures that the predicted probability for any potential donor comes from a model that did not use data from that potential donor, avoiding information leakage. For potential donors who were not approached (i.e., were not used to estimate authorization models), we predict authorization probabilities using each of the 10 cross-validation models, then take the mean as the final estimates.

B.3 Bayesian Structural Estimation

Here, we provide the full details for the Bayesian estimation of our structural model. Recall that our goal is to estimate the true transplant drivers $(\beta_0^*, \beta^*, \rho^*)$, the OPO's beliefs (β_0, β, ρ) , and the cost parameters (C_p, C_a) from the observed outcomes (r_i, a_i, y_i) , covariates x_i , instrument z_i , and estimated authorization probabilities $\hat{\kappa}_i$.

B.3.1 Priors

We define uninformative priors over our model parameters. First consider the cost of procurement C_p . Given the OPO's decision rule (7), it must be the case that $C_p < 1$, otherwise no potential donors would be approached. Similarly, $C_a < \min_{i:r_i=1} \hat{\kappa}_i$ (i.e., the minimum authorization probability among approached potential donors), otherwise (7) would be incompatible with the observed data. Further, both $C_a, C_p \geq 0$ by definition. We impose uniform priors for both cost parameters over these defined ranges,

$$\begin{aligned} C_p &\sim \text{Unif}[0, 1] \\ C_p &\sim \text{Unif}[0, \min_{i:r_i=1} \hat{\kappa}_i] \end{aligned}$$

Next, we consider the correlation parameters. By definition, $\rho^*, \rho \in [0, 1]$. Thus, for both these distributions, we use a standard normal prior that is truncated between 0 and 1,

$$\begin{aligned} \rho^* &\sim \text{TruncNorm}_{[0,1]}(0, 1) \\ \rho &\sim \text{TruncNorm}_{[0,1]}(0, 1) \end{aligned}$$

Finally, we discuss the true drivers (β_0^*, β^*) and the OPO beliefs (β_0, β) . Recall that the OPO's decision-rule can be expressed in partially linear form (8), if we define $\tilde{\beta} = \beta/\rho$ and $\tilde{\beta}_0 = \beta_0/\rho$. Similarly, the true transplant probability (2) can also be written more succinctly by defining $\tilde{\beta}^* = \beta^*/\rho^*$ and $\tilde{\beta}_0^* = \beta_0^*/\rho^*$. For convergence, it is beneficial to define our Bayesian model in terms of $(\tilde{\beta}_0, \tilde{\beta}, \tilde{\beta}_0^*, \tilde{\beta}^*)$ instead of $(\beta_0, \beta, \beta_0^*, \beta^*)$. This linear transformation is only for estimation and

without loss of generality. We define uninformative normal priors on the transformed parameters,

$$\tilde{\beta}_0 \sim \mathcal{N}(0, 100)$$

$$\tilde{\beta} \sim \mathcal{N}(0, 100)$$

$$\tilde{\beta}_0^* \sim \mathcal{N}(0, 100)$$

$$\tilde{\beta}^* \sim \mathcal{N}(0, 100)$$

Overall, all priors used in our model are diffuse and have a negligible impact on our results.

B.3.2 Likelihood

Our model fits the likelihood of two observed variables: r_i and y_i . We first discuss the approach outcome r_i . Recall from (8) that the OPO's approach rule can be written in partially linear form. To translate this into our Bayesian model, we use the data augmentation technique from Albert and Chib (1993). Define

$$f_i = \frac{\beta_0}{\rho} - \frac{\sqrt{1-\rho^2}}{\rho} \Phi^{-1} \left(C_p + \frac{C_a}{\hat{\kappa}_i} \right) + x_i^T \frac{\beta}{\rho}.$$

We can then define an auxiliary variable for each potential donor, $q_i = f_i + \eta_i$. Since $\eta_i \sim \mathcal{N}(0, 1)$, q_i is also normally distributed with variance 1 and mean f_i . Further, from (8), $q_i > 0$ for approached potential donors and $q_i < 0$ for unapproached potential donors. Thus, we can sample q_i conditional on r_i from a truncated normal distribution,

$$q_i \sim \begin{cases} \mathcal{N}^+(f(x_i), 1) & \text{if } r_i = 1 \\ \mathcal{N}^-(f(x_i), 1) & \text{if } r_i = 0 \end{cases} \quad (9)$$

where $\mathcal{N}^+$ is the standard normal truncated on $[0, \infty]$ and $\mathcal{N}^-$ is the standard normal truncated on $[-\infty, 0]$. This augmented approach is equivalent to the standard probit binary choice model which samples r_i as a Bernoulli variable with probability $\Phi(f_i)$ (Albert and Chib 1993).

The key benefit of sampling q_i instead of r_i is that we can use the resulting sample of q_i in the

censored second stage. Specifically, we can set $\eta_i = q_i - f(x_i)$ and sample

$$y_i \sim \text{Bernoulli}\left(a_i \Phi\left(\frac{\beta_0^* + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^{*2}}}\right)\right). \quad (10)$$

This models the censored outcomes y_i , accounting for selection on both the observables x_i and unobservables η_i . Overall, (9) and (10) fully define our Heckman-style model.

B.3.3 Additional Moment Conditions

As discussed in App. A.3, our calibration conditions on the OPO's beliefs (Thm. 5) imply two moment conditions. We implement these moment conditions as penalty terms in our Bayesian model. We first consider the aggregate calibration condition. This can be implemented as

$$\frac{1}{\sum_{i=1}^N a_i} \sum_{i: a_i=1}^N \Phi\left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}}\right) \sim \mathcal{N}\left(\frac{\sum_{i=1}^N y_i}{\sum_{i=1}^N a_i}, \lambda_{cal}^2\right) \quad (11)$$

Here, the left-hand side expresses the average probability that an authorized potential donor leads to a transplant implied by the model parameters (as discussed in App. A.3). The sampling statement forces this average to be close to the empirically observed average, $\sum y_i / \sum a_i$. λ_{cal} controls how strictly this penalty is enforced: the smaller the λ_{cal} , the closer the model implied average must be to the empirical one. This approach can be thought of as akin to regularization. In our empirical work, we find that setting $\lambda_{cal} = 0.001$ is sufficient to ensure aggregate calibration.

Finally, we discuss the implementation of our near-threshold calibration condition. For near-threshold calibration, the model-implied OPO transplant probability near the approach threshold (i.e., $C_p + C_a/\kappa_i$) must on average be close to the true transplant probability. For each potential donor, we define distance from the threshold as

$$d_i = \Phi\left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}}\right) - C_p - \frac{C_a}{\kappa_i}$$

Then, we need to ensure calibration for those potential donors for whom d_i is just greater than 0. One way to do so would be to define some small $\delta > 0$, then ensure average calibration for all potential donors with $d_i \in [0, \delta]$. We adopt a smooth relaxation of this strategy. Specifically, we define calibration weights $w_i = r_i \exp(-d_i^2/\lambda_{tr}^2)$ for a small constant λ_{tr} . Then, w_i is large for all

potential donors with d_i close to 0, and decays as d_i grows larger. Further, the weights are zero for any potential donor who was not approached, as these potential donors do not contribute to the condition. We then enforce the moment condition as

$$\frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i \hat{\kappa}_i \Phi\left(\frac{\beta_0 + x_i^T \beta + \rho \eta_i}{\sqrt{1 - \rho^2}}\right) \sim \mathcal{N}\left(\frac{\sum_{i=1}^N w_i y_i}{\sum_{i=1}^N w_i}, \lambda_{cal}^2\right) \quad (12)$$

Effectively, this condition enforces that conditional on approach, the model-implied weighted average OPO belief is close to the weighted average of true transplant outcomes. In our empirical work, we set $\lambda_{tr} = 0.02$, which results in decreasing positive weights for $d_i \in [0, 0.05]$ and near-zero weights for $d_i \notin [0, 0.05]$. As before, we set $\lambda_{cal} = 0.001$.

B.3.4 Posterior Inference

We implemented our model in Stan, and used the NUTS MCMC sampler to approximate the posterior distribution of our parameters. We use 4 chains of 3,000 samples each. Within each chain, the first 1,400 samples are used for warm-up (i.e., burn-in), and the remaining are used for posterior inference. We validated convergence according to existing best practices (Aki et al. 2020) by confirming that the blended R-hat was smaller than 1.01 and the bulk effective sample size exceeded 400 for each parameter.

C Additional Results

C.1 Full Model Coefficients

In Table 4 in the main text, we presented a subset of parameter estimates. We present additional coefficients for our structural model in Table 8 and Table 9.

Table 8. Full coefficients (β, β^*) for OPOs A & B. We present mean estimated coefficients and 95% credible intervals across 6,000 MCMC draws. Bolded and starred coefficients have credible intervals that exclude 0.

Variable	OPO A		OPO B	
	OPO	True	OPO	True
Year: 2017	-0.04 (-0.12, 0.03)	0.21 (-0.11, 0.55)	-0.14* (-0.20, -0.08)	-0.08 (-0.35, 0.18)
Year: 2018	0.03 (-0.04, 0.11)	0.17 (-0.15, 0.51)	-0.06* (-0.11, -0.003)	-0.23 (-0.50, 0.03)
Year: 2019	0.13* (0.06, 0.20)	0.41* (0.11, 0.73)	0.15* (0.09, 0.20)	0.07 (-0.20, 0.33)
Year: 2020	0.23* (0.15, 0.31)	0.39* (0.10, 0.69)	0.19* (0.14, 0.26)	-0.09 (-0.36, 0.17)
Year: 2021	0.14* (0.05, 0.23)	0.24 (-0.11, 0.62)	0.14* (0.09, 0.20)	-0.27* (-0.53, -0.02)
Brain Death (vs. Cardiac Death)	1.21* (1.02, 1.42)	1.42* (1.17, 1.67)	1.29* (1.10, 1.49)	1.24* (1.05, 1.43)
Age	-0.19* (-0.25, -0.14)	-0.44* (-0.63, -0.26)	-0.08* (-0.12, -0.04)	-0.41* (-0.59, -0.24)
Age $\times$ $\mathbf{1}\{\text{Age} > 50\}$	-0.33* (-0.44, -0.23)	-0.17 (-0.56, 0.22)	-0.20* (-0.28, -0.13)	-0.29 (-0.60, 0.03)
Age $\times$ $\mathbf{1}\{\text{Age} < 18\}$	0.57* (0.41, 0.75)	0.51 (-0.26, 1.23)	0.51* (0.38, 0.64)	0.56 (-0.13, 1.23)
Sex: Female (vs. Male)	-0.04 (-0.08, 0.000)	-0.13 (-0.33, 0.07)	-0.003 (-0.03, 0.03)	-0.02 (-0.18, 0.14)
Race: Black / African American (vs. White)	-0.08* (-0.14, -0.03)	0.16 (-0.13, 0.47)	-0.06* (-0.10, -0.02)	0.11 (-0.11, 0.36)
Race: Hispanic (vs. White)	-0.21* (-0.37, -0.05)	0.01 (-0.53, 0.60)	-0.05* (-0.09, -0.01)	0.01 (-0.18, 0.20)
Race: Other / Unknown (vs. White)	-0.12 (-0.26, 0.01)	0.001 (-0.55, 0.64)	-0.04 (-0.13, 0.04)	-0.25 (-0.66, 0.22)
Cause of Death: CNS / Tumor (vs. Anoxia)			0.005 (-0.23, 0.22)	-0.32 (-0.98, 0.35)
Cause of Death: CVA/Stroke (vs. Anoxia)			0.15* (0.10, 0.20)	0.08 (-0.14, 0.29)
Cause of Death: Head Trauma (vs. Anoxia)			0.12* (0.07, 0.17)	0.02 (-0.22, 0.25)
Cause of Death: Other / Unknown (vs. Anoxia)			-0.10* (-0.14, -0.06)	-0.64* (-0.90, -0.38)
Diabetes	0.02 (-0.03, 0.08)	-0.35* (-0.61, -0.10)	-0.04* (-0.07, -0.01)	-0.19* (-0.38, -0.01)
Hypertension	-0.03 (-0.08, 0.02)	-0.04 (-0.26, 0.19)	-0.01 (-0.05, 0.02)	-0.05 (-0.24, 0.13)
Sepsis	-0.33* (-0.42, -0.25)	-0.55* (-0.93, -0.18)	-0.23* (-0.30, -0.17)	-0.33* (-0.65, -0.003)
End-Stage Renal Disease	-0.07 (-0.21, 0.07)	0.01 (-0.52, 0.54)	-0.07 (-0.16, 0.01)	0.09 (-0.34, 0.54)
Chronic Kidney Disease	-0.07 (-0.16, 0.02)	0.09 (-0.33, 0.54)	-0.10* (-0.16, -0.04)	0.01 (-0.32, 0.37)
Cirrhosis	0.08 (-0.003, 0.18)	0.18 (-0.26, 0.66)	-0.17* (-0.25, -0.10)	-0.11 (-0.54, 0.35)
COPD	-0.01 (-0.07, 0.05)	-0.04 (-0.32, 0.27)	-0.01 (-0.05, 0.04)	-0.10 (-0.36, 0.17)
Atrial fibrillation	-0.10* (-0.19, -0.02)	0.02 (-0.37, 0.44)	-0.04 (-0.10, 0.02)	0.05 (-0.28, 0.42)
Dialysis	0.04 (-0.04, 0.12)	-0.03 (-0.42, 0.35)	-0.02 (-0.09, 0.04)	-0.22 (-0.55, 0.11)
CRRT	-0.18* (-0.28, -0.10)	-0.003 (-0.48, 0.47)	-0.17* (-0.25, -0.10)	-0.21 (-0.58, 0.18)
ECMO	-0.12* (-0.25, -0.01)	0.07 (-0.40, 0.62)	-0.12* (-0.23, -0.01)	-0.04 (-0.58, 0.53)
On Vasopressors	0.02 (-0.02, 0.07)	-0.15 (-0.36, 0.05)	0.05* (0.02, 0.08)	-0.16 (-0.35, 0.02)
Creatinine	-0.29* (-0.38, -0.21)	-0.19 (-0.51, 0.11)	-0.13* (-0.18, -0.08)	-0.27* (-0.54, -0.01)
Creatinine $\times$ $\mathbf{1}\{\text{Creatinine} > 1.5\}$	0.21* (0.13, 0.31)	-0.04 (-0.37, 0.32)	0.11* (0.06, 0.17)	0.18 (-0.11, 0.49)
log(Total Bilirubin)	-0.10* (-0.13, -0.07)	-0.14* (-0.28, -0.01)	-0.07* (-0.09, -0.05)	-0.18* (-0.31, -0.06)
log(max(AST, ALT))	-0.003 (-0.04, 0.03)	-0.08 (-0.20, 0.04)	-0.03* (-0.05, -0.001)	-0.05 (-0.17, 0.06)
$\mathbf{1}\{\text{max(AST, ALT)} > 100\}$	0.04 (-0.03, 0.12)	0.12 (-0.19, 0.42)	0.06* (0.01, 0.12)	0.05 (-0.21, 0.30)
$\mathbf{1}\{\text{max(AST, ALT)} > 1000\}$	-0.14* (-0.27, -0.02)	-0.42* (-0.79, -0.05)	-0.12* (-0.20, -0.04)	-0.34 (-0.71, 0.03)
Creatinine Missing	-0.59* (-0.75, -0.45)	-0.64* (-1.23, -0.08)	-0.08 (-0.18, 0.02)	-1.05* (-1.79, -0.39)
Total Bilirubin Missing	0.07 (-0.02, 0.17)	-0.61* (-1.06, -0.18)	-0.06 (-0.14, 0.02)	-0.08 (-0.54, 0.39)
AST and ALT Missing	-0.20* (-0.38, -0.03)	-0.07 (-0.60, 0.45)	-0.12 (-0.26, 0.01)	-0.16 (-0.73, 0.39)

Table 9. Full coefficients (β, β^*) for OPOs C and D. We present mean estimated coefficients and 95% credible intervals across 6,000 MCMC draws. Bolded and starred coefficients have credible intervals that exclude 0.

Variable	OPO C		OPO D	
	OPO	True	OPO	True
Year: 2017	-0.02 (-0.11, 0.06)	-0.09 (-0.42, 0.24)	0.08 (-0.02, 0.18)	-0.24 (-0.65, 0.19)
Year: 2018	-0.06 (-0.15, 0.03)	-0.06 (-0.38, 0.26)	0.18* (0.08, 0.29)	0.11 (-0.31, 0.55)
Year: 2019	-0.05 (-0.14, 0.04)	0.09 (-0.24, 0.43)	0.26* (0.16, 0.39)	0.12 (-0.28, 0.53)
Year: 2020	-0.05 (-0.13, 0.04)	0.03 (-0.30, 0.35)	0.43* (0.31, 0.59)	0.22 (-0.16, 0.63)
Year: 2021	-0.09* (-0.19, -0.01)	-0.04 (-0.37, 0.29)	0.61* (0.45, 0.80)	0.39* (0.01, 0.80)
Brain Death (vs. Cardiac Death)	1.23* (1.05, 1.41)	1.25* (1.01, 1.49)	1.18* (0.91, 1.48)	1.56* (1.26, 1.89)
Age	-0.03 (-0.09, 0.03)	-0.24* (-0.44, -0.05)	0.03 (-0.02, 0.08)	-0.30* (-0.53, -0.08)
Age $\times$ $\mathbf{1}\{\text{Age} > 50\}$	-0.61* (-0.76, -0.47)	-0.57* (-1.06, -0.10)	-0.51* (-0.69, -0.37)	-0.34 (-0.83, 0.13)
Age $\times$ $\mathbf{1}\{\text{Age} < 18\}$	0.41* (0.21, 0.62)	0.77 (-0.04, 1.52)	0.12 (-0.05, 0.29)	0.27 (-0.54, 1.07)
Sex: Female (vs. Male)	-0.06* (-0.11, -0.004)	-0.28* (-0.50, -0.06)	-0.04 (-0.09, 0.005)	-0.15 (-0.40, 0.11)
Race: Black / African American (vs. White)	-0.17* (-0.25, -0.10)	0.17 (-0.19, 0.58)	-0.12* (-0.19, -0.06)	0.06 (-0.31, 0.46)
Race: Hispanic (vs. White)	-0.08* (-0.15, -0.01)	0.06 (-0.20, 0.33)	-0.08 (-0.22, 0.05)	0.12 (-0.46, 0.73)
Race: Other / Unknown (vs. White)	-0.04 (-0.21, 0.14)	0.67* (0.02, 1.45)	-0.04 (-0.24, 0.15)	-0.40 (-1.14, 0.30)
Cause of Death: CNS / Tumor (vs. Anoxia)	-0.16 (-0.54, 0.23)	-0.23 (-1.00, 0.59)	-0.80* (-1.75, -0.07)	0.06 (-0.77, 0.90)
Cause of Death: CVA/Stroke (vs. Anoxia)	0.10* (0.02, 0.18)	-0.11 (-0.42, 0.19)	0.11* (0.05, 0.19)	0.15 (-0.17, 0.49)
Cause of Death: Head Trauma (vs. Anoxia)	0.09* (0.01, 0.17)	0.58* (0.27, 0.91)	0.04 (-0.02, 0.11)	0.27 (-0.05, 0.61)
Cause of Death: Other / Unknown (vs. Anoxia)	-0.14* (-0.21, -0.06)	-1.23* (-1.64, -0.84)	-0.36* (-0.48, -0.25)	-0.19 (-0.72, 0.34)
Diabetes	-0.09* (-0.16, -0.02)	-0.22 (-0.51, 0.07)	-0.04 (-0.10, 0.01)	-0.29 (-0.61, 0.05)
Hypertension	-0.03 (-0.09, 0.03)	0.12 (-0.11, 0.36)	0.01 (-0.04, 0.06)	0.03 (-0.25, 0.31)
Sepsis	-0.38* (-0.51, -0.26)	-0.36 (-0.84, 0.14)	-0.13* (-0.23, -0.04)	-0.004 (-0.48, 0.49)
End-Stage Renal Disease	-0.24* (-0.46, -0.03)	0.35 (-0.42, 1.17)	0.06 (-0.10, 0.22)	-0.17 (-0.87, 0.51)
Chronic Kidney Disease	-0.17* (-0.31, -0.03)	-0.19 (-0.82, 0.49)	-0.01 (-0.11, 0.09)	0.24 (-0.28, 0.80)
Cirrhosis	-0.27* (-0.42, -0.13)	0.09 (-0.51, 0.76)	-0.18* (-0.34, -0.05)	-0.04 (-0.62, 0.58)
COPD	-0.14* (-0.23, -0.05)	0.01 (-0.37, 0.44)	0.003 (-0.06, 0.06)	-0.12 (-0.44, 0.22)
Atrial fibrillation	-0.01 (-0.13, 0.11)	0.23 (-0.25, 0.73)	-0.11* (-0.21, -0.02)	0.14 (-0.34, 0.67)
Dialysis	-0.14* (-0.29, -0.002)	-0.32 (-0.87, 0.27)	0.04 (-0.06, 0.14)	0.14 (-0.37, 0.69)
CRRT	-0.33* (-0.48, -0.19)	-0.53 (-1.06, 0.01)	-0.16* (-0.28, -0.06)	-0.46 (-1.00, 0.09)
ECMO	0.09 (-0.12, 0.29)	0.06 (-0.61, 0.80)	-0.09 (-0.38, 0.17)	0.07 (-0.73, 0.87)
On Vasopressors	0.03 (-0.02, 0.09)	0.04 (-0.20, 0.28)	0.06* (0.01, 0.12)	0.33* (0.04, 0.61)
Creatinine	-0.31* (-0.41, -0.22)	-0.58* (-0.94, -0.24)	-0.16* (-0.25, -0.09)	-0.23 (-0.56, 0.09)
Creatinine $\times$ $\mathbf{1}\{\text{Creatinine} > 1.5\}$	0.24* (0.14, 0.35)	0.40* (0.02, 0.80)	0.11* (0.03, 0.21)	0.09 (-0.28, 0.47)
log(Total Bilirubin)	-0.07* (-0.11, -0.04)	-0.17* (-0.34, -0.01)	-0.09* (-0.13, -0.05)	-0.24* (-0.45, -0.05)
log(max(AST, ALT))	-0.03 (-0.09, 0.02)	0.01 (-0.14, 0.15)	0.01 (-0.03, 0.05)	-0.08 (-0.23, 0.06)
$\mathbf{1}\{\text{max(AST, ALT)} > 100\}$	0.08 (-0.02, 0.19)	-0.06 (-0.40, 0.28)	0.04 (-0.05, 0.13)	0.17 (-0.19, 0.54)
$\mathbf{1}\{\text{max(AST, ALT)} > 1000\}$	-0.17* (-0.33, -0.02)	-0.26 (-0.71, 0.21)	-0.21* (-0.37, -0.07)	0.14 (-0.38, 0.71)
Creatinine Missing	-0.44* (-0.60, -0.30)	-0.81* (-1.58, -0.12)	-0.35* (-0.51, -0.20)	-0.14 (-0.82, 0.50)
Total Bilirubin Missing	-0.03 (-0.13, 0.07)	-0.25 (-0.82, 0.33)	0.02 (-0.10, 0.14)	-0.11 (-0.65, 0.46)
AST and ALT Missing	-0.22 (-0.45, 0.003)	-0.39 (-1.07, 0.26)	-0.03 (-0.23, 0.17)	-0.42 (-1.08, 0.19)

C.2 Robustness Check

In the main text, we presented results from a model that assumed that the costs C_p and C_a were constant across all OPOs. Here, we present a robustness check that relaxes this assumption by allowing these costs to vary by OPO. Table 10 reports separate estimates of the costs and approach thresholds for each OPO. While the estimates are broadly similar, there is some evidence of heterogeneity. For instance, OPO D appears to have a lower procurement cost ($C_p = 0.28$) compared to the other OPOs (which all have $C_p \approx 0.42$). Similarly, OPO B has the lowest C_a , suggesting it adjusts its approach threshold less in response to authorization probability than the other OPOs. However, these differences are not sharply estimated: the 95% credible intervals are wide and generally include zero. As a result, we cannot draw strong conclusions about variation in costs across OPOs.

Importantly, both the pooled and unpooled specifications yield nearly identical results and domain implications. To demonstrate this, we replicate key analyses from the main text using the unpooled model. Fig. 6 shows the total donation opportunity (as in Section 6.1), Table 11 and Fig. 7 display the AUROC and calibration of OPO beliefs (as in Section 6.2), and Table 12 presents our policy analysis results (as in Section 7). All these results are highly consistent with the pooled model. While we omit full coefficient tables for brevity, we note that the parameter estimates are also similar across specifications. This robustness check confirms that pooling costs across OPOs does not meaningfully affect our findings or policy implications.

Table 10. Estimated OPO costs and approach thresholds (Unpooled). We present mean estimated parameters and 95% credible intervals (in parentheses) across 6,000 MCMC draws.

	OPO A	OPO B	OPO C	OPO D
C_{proc}	0.42 (0.30, 0.52)	0.42 (0.29, 0.53)	0.43 (0.26, 0.56)	0.28 (0.14, 0.42)
C_{appr}	0.01 (0.01, 0.02)	0.001 (0.000, 0.003)	0.01 (0.000, 0.02)	0.004 (0.000, 0.01)
$\min \tau_i$	0.43 (0.32, 0.53)	0.42 (0.29, 0.53)	0.44 (0.27, 0.56)	0.29 (0.15, 0.42)
$\max \tau_i$	0.64 (0.52, 0.73)	0.49 (0.36, 0.59)	0.53 (0.37, 0.65)	0.41 (0.21, 0.64)




Figure 6. Total Number of Transplantable Donors (unpooled costs). We plot the additional number of donors the OPOs would procure if they (1) approached for all potential donors (“Approaching All”), and (2) obtained authorization from all approached families (“No Family Declines”). We plot the mean numbers of transplantable donors across all four OPOs; error bars indicate 95% credible intervals.

Table 11. Estimated AUROC by OPO (unpooled costs). We display the mean posterior AUROC for each OPO’s beliefs along with 95% credible intervals.

	Combined	OPO A	OPO B	OPO C	OPO D
AUROC: OPO Beliefs	0.76 (0.74, 0.78)	0.78 (0.75, 0.81)	0.75 (0.73, 0.77)	0.78 (0.75, 0.81)	0.74 (0.69, 0.79)
AUROC: Correct Beliefs	0.87 (0.85, 0.88)	0.86 (0.83, 0.89)	0.86 (0.84, 0.88)	0.89 (0.86, 0.91)	0.85 (0.81, 0.89)

Table 12. Counterfactual policy evaluation (unpooled costs). The Status Quo row reports the observed number of donors, approaches, and authorizations; subsequent rows report the incremental Δ relative to this baseline. 95% CI’s represent Bayesian credible intervals.

Policy Name	Donors		Approaches		Authorizations	
	Δ	95% CI	Δ	95% CI	Δ	95% CI
0. Status Quo	5,164 +		11,313 +		7,018 +	
1. Approach 20% More	370	(328, 412)	2,263	(2260, 2265)	892	(830, 954)
2. Approach All	2,196	(1988, 2418)	24,543	(24543, 24543)	8,757	(8629, 8885)
3. Correct Beliefs (Fixed Approach)	488	(380, 578)	0	(-2, 2)	314	(119, 478)
4. Correct Beliefs (Increased Approach)	957	(602, 1303)	2,454	(532, 4574)	1,402	(578, 2302)
5. Improve Authorization (5%)	251	(172, 329)	0	(0, 0)	538	(452, 621)
6. Improve Authorization (Equity)	393	(314, 475)	0	(0, 0)	699	(617, 783)




Figure 7. Miscalibration of OPO beliefs (unpooled costs). We plot the distribution of OPO belief miscalibration by true transplant probability. Miscalibration (y-axis) is defined as the true probability minus the OPO's belief. Potential donors are grouped into ten bins by true probability; within each bin, we show a boxplot of miscalibration, averaged across MCMC draws. The solid line marks the median, the box spans the interquartile range, and the whiskers indicate the 2.5th and 97.5th percentiles. The dashed grey line indicates perfect calibration (i.e., belief equals truth).