Organ Availability and Transplant Candidate Behavior

Younghee Son *
Department of Economics
Michigan State University

This version : October 14, 2026

Click here for the latest version

Abstract

Liver transplant candidates decline more than half of the organs offered to them, even though no alternative treatment exists. Each year, roughly one in eight dies or becomes too sick for transplant while waiting. This paper studies how candidates make these decisions. The 2020 Acuity Circles reform changed how candidates are prioritized, both redrawing the geographic boundaries and redefining the medical urgency groups. Using offer-level data from the Scientific Registry of Transplant Recipients, I show that candidates who gained access became more selective: a one-standard-deviation increase in expected offer arrival rates lowers the probability of accepting an offer by 4.7 percentage points, a 12.6 percent decline. To interpret this behavior, I estimate a dynamic discrete choice model of the accept-or-wait decision. Candidates weigh the offer at hand against the chance of a better offer later, knowing they may die or become too sick while waiting. The estimated model implies that the reform raised average welfare at listing by an amount equivalent to roughly doubling donor supply, with only 1.6 percent of candidates made worse off. I then evaluate several counterfactual policy changes, including a widely debated one which would replace the remaining boundaries with a continuous score that incorporates both medical urgency and proximity. I find that prioritizing the sickest candidates raises average welfare but at the cost of wasted organs.

Keywords: Liver allocation, healthcare policy

*I gratefully acknowledge financial support from Stacy Dickert-Conlin's HRSA grant Project No. 9807. I especially thank Todd Elder, Stacy Dickert-Conlin, Hanzhe Zhang and Mat Reeves for their valuable guidance and comments on this paper. The data reported here have been supplied by the Hennepin Healthcare Research Institute (HHRI) as the contractor for the Scientific Registry of Transplant Recipients (SRTR). The interpretation and reporting of these data are the responsibility of the author and in no way should be seen as an official policy of or interpretation by the SRTR or the U.S. Government. All errors are my own.

1. Introduction

In 2023, chronic liver disease and cirrhosis caused 52,274 adult deaths in the United States, making it the ninth leading cause of death (Centers for Disease Control and Prevention, 2025). Once the disease progresses to liver failure, medical therapy can manage complications but cannot restore lost function. Unlike kidney failure, where dialysis can replace lost renal function, no comparable treatment can substitute for the liver. A transplant is therefore the only definitive treatment for liver failure.

Donated livers are far more scarce than the patients who need them. In 2024, 11,458 candidates received a transplant, and 94.7% of these organs came from deceased donors. In the same year, roughly 14,000 new patients joined the waitlist. Because new listings continued to exceed transplants, more than 24,000 candidates were on the waitlist at some point during 2024 (Kwong et al., 2026). Each year, about 12% of candidates die or become too sick to undergo transplant while still waiting.[1] This shortage cannot be resolved through a market because federal law prohibits the sale of human organs.[2] Instead of relying on prices, the United States allocates donated organs through a centralized rule-based system, governed by the Organ Procurement and Transplantation Network (OPTN).

A severe shortage combined with a central allocation system leaves waitlist candidates with limited options. Even so, candidates decline the vast majority of offers they receive. Deceased-donor livers that were eventually transplanted in 2016 generated more than 170,000 individual offers, and only about 4% of these offers were accepted (Wey et al., 2018). Among candidates who later died or were removed from the waitlist, 84% had received at least one liver offer before leaving the list (Lai et al., 2012).[3]

Why would a candidate decline an offer when organs are scarce and transplantation can save lives? This paper argues that the answer lies in forward-looking behavior. Candidates who decline are betting that a better offer will arrive before they become too sick to receive a transplant. Acceptance decisions therefore depend on how many future offers they expect to receive, and that expectation is shaped by the allocation rule. Allocation rules influence outcomes through two channels: they decide who receives offers, and they affect how candidates respond to them. Yet the tools used to evaluate and design allocation rules focus only on the first channel. I study whether candidates adjust their behavior when their access to future offers changes, and what this implies

[1] Source: https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/

[2] The National Organ Transplant Act of 1984 makes the exchange of organs for valuable consideration a federal crime (42 U.S. Code § 274e), punishable by fines of up to $50,000 and prison terms of up to five years.

[3] In my sample, which covers April 2016 to December 2023, 89.4% of deceased-donor livers were declined at the initial offer, and 90.5% of candidates removed from the waitlist because of death or deteriorating health had received at least one offer before removal.

for policy.

When a liver becomes available, the system groups candidates by medical urgency and geographic proximity to the donor. Offers move through these groups in batches. On February 4, 2020, the Acuity Circles reform changed both urgency categories and geographic boundaries. Before the reform, proximity was defined by Donation Service Areas (DSAs), and candidates in the donor's DSA were prioritized. The reform replaced DSAs with concentric circles around the donor hospital and split urgency categories more finely. These changes shifted candidates into new groups and changed their position in the offer sequence. The size of the shift differed by urgency and location, and I use this variation to identify how behavior responds. I measure the change in access, ΔAccess, as the change in the probability that a candidate receives a first offer shortly after listing. The analysis uses restricted-use data from the Scientific Registry of Transplant Recipients (SRTR), which record the universe of liver transplant candidates and deceased donors in the United States, and every offer made between them.

I find that candidates respond strongly to changes in future access. A one-standard-deviation increase in ΔAccess lowers the probability of accepting an offer by 4.7 percentage points. Relative to the pre-reform acceptance rate of 37.6%, this represents a 12.6% decline in acceptance. Candidates who gain access therefore become more selective, declining offers they would previously have accepted. Greater selectivity does not undo the gain in access. A one-standard-deviation increase in access lowers overall mortality from listing by about 0.9 to 1.2 percentage points, almost entirely through reductions in waitlist mortality. A back-of-the-envelope calculation implies roughly 76 fewer deaths over the post-reform period among candidates in groups that gained access.

Although the reduced-form estimates show that candidates respond to the allocation rule, they cannot identify how candidates value waiting relative to acceptance, or evaluate rules that have never been implemented. I therefore develop and estimate a dynamic model of acceptance. Candidates receive offers over time and, at each offer, choose whether to accept or to decline and continue waiting. Because each liver is offered simultaneously to a ranked set of candidates, accepting leads to transplant only if no higher-priority candidate also accepts. Waiting carries the risk of deteriorating health or leaving the waitlist before another offer arrives. This is an optimal stopping problem in which the value of waiting depends on offer arrival rates, organ characteristics such as quality and distance, and how priority evolves as health changes. I formulate it in continuous time and estimate it using the conditional choice probability inversion approach of Hotz and Miller (1993) and Arcidiacono et al. (2016). This approach recovers the value of waiting and the value of transplantation without solving the full dynamic program inside the estimation loop.

Because the setting has no prices, I follow Agarwal et al. (2021) and express welfare

changes as the proportional increase in donor supply that would make a candidate equally as well off as the policy does. By this measure, the Acuity Circles reform raises welfare at listing by an amount comparable to roughly doubling donor supply. I decompose this gain and find that it is driven by two large opposing forces. Because circles reach farther than DSAs did, the reform more than triples the rate at which offers arrive. It also increases competition for each liver, but the gain from arrival dominates. The reform is close to a Pareto improvement: only 1.6% of candidates are worse off, and these losses are concentrated in a small number of DSAs that had unusually rich organ supply conditions before the policy change.

Finally, I use the estimated model to ask what the reform's design achieved and what alternative designs could have done. I run two sets of counterfactuals. The first decomposes the new policy into its two components, the distance-based circles and the finer urgency categories. I turn on each component individually while holding the other at its pre-reform setting. I find that the two components are complements. The circles account for most of the welfare gain but also generate the largest share of candidates who are worse off. The finer urgency categories, on their own, are almost purely redistributive and add little to welfare. When combined, however, the two components amplify the welfare gains and shrink the set of candidates left worse off.

I then evaluate a new rule now under consideration by the OPTN, known as "Continuous Distribution". This rule would eliminate eligibility boundaries altogether and rank every candidate by a single score that weights medical urgency against proximity to the donor. I trace out the urgency-proximity frontier by solving the model over a range of weights. Average welfare is highest when the score places all weight on urgency, but this comes with a substantial increase in unused livers. A score that places all weight on proximity produces the largest share of candidates made worse off. Between these extremes, I find a range of weights that improves on Acuity Circles on all three margins: higher average welfare, fewer candidates worse off, and fewer unused livers.

This paper is most closely related to the empirical literature on the demand side of the organ market.[4] Local organ supply affects who joins the waitlist (Dickert-Conlin et al., 2019), and acceptance behavior responds to several features of the decision environment: the quality of the organ and the candidate's own health (Howard, 2002), declines by higher-ranked candidates (Zhang, 2010), and the allocation mechanism itself (Agarwal et al., 2018). Three papers are closest to this one. Choi (2020) shows that liver candidates who gained access under an earlier expansion of regional sharing for the sickest candidates became less likely to accept offers. Bae (2024) studies the analogous reform in kidney allocation and finds that it reduced organ nonuse and pre-transplant mortality; its focus is on organ use and mortality rather than on acceptance decisions. Agarwal et al. (2021)

[4] A large literature studies the supply of transplantable organs (Abadie and Gay, 2006; Dickert-Conlin et al., 2011; Kessler and Roth, 2012; Callison and Levin, 2016).

estimate a dynamic model of kidney acceptance and use it to compare waitlist designs.

I build on these papers in three ways. First, I measure how access changes for every combination of urgency and location, so the acceptance response is identified across the entire waitlist—among candidates who lost access as well as those who gained it. Second, I estimate how the same change in access affects both acceptance and mortality, so the mortality estimate captures the net effect of more offers and more selective decisions. Third, I estimate a dynamic model of acceptance, following Agarwal et al. (2021), and use it to evaluate rules that have never been implemented. I differ from them by studying livers, where the incentive to wait is less obvious: kidney candidates accumulate priority with time on the waitlist and can remain on dialysis, while liver candidates have neither. The counterfactuals also provide evidence on the rule the OPTN is considering before it is adopted. Ashlagi and Roth (2026) note that no evidence on the efficiency of such rules yet exists. More broadly, the paper relates to work on how agents behave when scarce goods are centrally allocated by rule rather than by price—kidneys (Agarwal et al., 2021), hearts (Sweat, 2026), public housing (Waldinger, 2021), child care (Kennes et al., 2014), and school seats (Abdulkadiroğlu et al., 2017). In nearly all of these settings an outside option exists. Livers have none, so the cost of waiting is highest in this setting.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and the Acuity Circles reform. Section 3 describes the data. Section 4 presents the empirical strategy and results. Section 5 develops and estimates the dynamic model and uses it to evaluate the reform and the alternative rules that have never been implemented. Section 6 concludes.

2. Institutional Detail

This section describes the institutional features that determine how livers are allocated and how candidates make acceptance decisions.

2.1. Regulatory Framework and Organizations

In 1984, the U.S. Congress enacted the National Organ Transplant Act (NOTA) to address the organ shortage and improve the allocation of donated organs. The law established the Organ Procurement and Transplantation Network (OPTN) to operate a national system for matching organs with potential recipients. The OPTN maintains the national waitlist, allocates deceased-donor organs, and sets allocation policy. It is overseen by the Health Resources and Services Administration (HRSA) and operated through federal contracts. From 1986 through the end of my sample period, a single contractor, the United Network for Organ Sharing (UNOS), a private non-profit organization, operated the OPTN. Beginning in 2025, under the Securing the U.S. OPTN Act of 2023,

HRSA moved to a multi-vendor model: an OPTN Board of Directors independent of any contractor was established, several functions were brought in-house or awarded to other contractors, and UNOS continues to maintain the national matching system under a transitional contract. NOTA also established the Scientific Registry of Transplant Recipients (SRTR), which collects and reports transplant data for research and policy evaluation.

2.2. Transplant Process

To receive a deceased-donor transplant, a patient must first register at a transplant center. A candidate may register at multiple centers, a practice known as multi-listing, to increase the probability of receiving an offer.[5] There are 146 active liver transplant centers in the United States.[6] Each center evaluates candidates for listing and reports their medical information to the OPTN. The center then updates each candidate's medical status and activity over time.

When a liver becomes available, the OPTN matching system screens the waitlist for medically compatible candidates. Compatibility depends on blood type, body size, and other clinical criteria, together with any minimum requirements specified by the candidate.[7] The system then ranks compatible candidates according to the allocation policy, which depends primarily on medical urgency and geography, and sends out offers. Because organs must be transplanted quickly, offers are sent to multiple candidates at the same time. If several candidates accept, the highest-priority acceptor receives the transplant. Because a response is required within an hour, the transplant team may decide whether to accept without the candidate's knowledge. Throughout the paper, "candidate" refers to the candidate and the team deciding on the candidate's behalf.[8]

2.3. Liver Allocation Rule

Allocation rules differ across organ types.[9] In liver allocation, the key determinants are medical urgency and distance between the donor and the recipient.[10]

[5] About 4% of the candidates listed for liver transplants are on multiple lists (Brown et al., 2020).

[6] The United States has about 250 active transplant centers across all organ types, although the number varies by organ. See https://optn.transplant.hrsa.gov/about/search-membership/ for details.

[7] Candidates must specify acceptable donor weight and may set additional criteria, such as tolerance for antigen mismatch or willingness to accept specific donor conditions.

[8] The data do not distinguish between decisions made by the candidate and by the transplant center.

[9] For example, kidney and lung allocation also depend on waiting time and expected survival benefit. In liver allocation, waiting time serves only as a tie-breaker among candidates with the same MELD score, and it counts only time spent at the candidate's current or a higher score. See https://optn.transplant.hrsa.gov/patients/about-transplantation/how-organ-allocation-works/.

[10] See https://optn.transplant.hrsa.gov/policies-bylaws/policies/ for details.

2.3.1. Medical Urgency

Candidates are ranked by medical urgency using either a Status 1A designation or a MELD score. Status 1A is reserved for the most critical cases, in which candidates face imminent death without a transplant. To qualify, a candidate must demonstrate a life expectancy of fewer than seven days and satisfy specific medical criteria.[11] Fewer than 50 candidates hold Status 1A nationwide at any given time.[12] All other candidates are ranked by the Model for End-stage Liver Disease (MELD) score, which measures the severity of illness. Scores range from 6 to 40, and a higher score indicates a higher probability of death within 90 days without a transplant (Singal and Kamath, 2013). A score of 40 corresponds to an 89% mortality risk, while a score of 6 corresponds to a risk of about 3%.[13] The score is computed from the following formula and rounded to the nearest integer.

$$\begin{aligned} MELD =& 1.33(Female) + 4.56\ln(bilirubin) + 0.82(137 - sodium) \\ & - 0.24(137 - sodium)\ln(bilirubin) + 9.09\ln(INR) + 11.14\ln(creatinine) \\ & + 1.85(3.5 - albumin) - 1.83(3.5 - albumin)\ln(creatinine) + 6. \end{aligned}$$

A candidate may request an exception when the calculated MELD score does not appropriately reflect their clinical severity. Some diagnoses, such as Hepatocellular Carcinoma, qualify automatically and are assigned an exception score.[14] Allocation uses the higher of the calculated score and the exception score.

2.3.2. Geographic Proximity

A recovered liver must be transplanted within 8-12 hours to preserve its viability. During transport, the organ is preserved on ice in a cold solution. The duration between removal from the donor and implantation in the recipient is known as cold ischemia time (CIT), and it strongly affects outcomes. Longer CIT significantly increases the risk of graft failure and prolongs hospitalization (Lozanovski et al., 2020; Pan et al., 2018). Because of these constraints, distance between the hospital where the liver is recovered (donor hospital) and the candidate's transplant center plays an important role in allocation.

[11] Qualifying conditions include fulminant liver failure, anhepatic status, primary non-function within seven days of transplant, hepatic artery thrombosis within seven days, and acute decompensated Wilson's disease, as defined in OPTN policy.

[12] https://optn.transplant.hrsa.gov/policies-bylaws/a-closer-look/liver/

[13] https://www.mayoclinic.org/medical-professionals/transplant-medicine/calculators/meld-score-and-90-day-mortality-rate-for-alcoholic-hepatitis/itt-20434719

[14] During the sample period, 27.5% of candidates applied for an exception, including automatically eligible cases, and 93.4% of those requests were approved.

2.4. Acuity Circles

NOTA also created regional Organ Procurement Organizations (OPOs), which coordinate the recovery, preservation, and transport of organs. Each OPO serves a designated geographic area called a Donation Service Area (DSA), defined as a group of counties.[15] The OPTN also groups DSAs into 11 regions (Figure 1). Figure 2 shows the boundaries of DSAs, which are generally aligned with state borders but not identical. From 1989 to 2020, liver allocation relied primarily on these geographic units. A recovered organ was first offered within the donor hospital's DSA, then within the broader region, and finally nationwide.

These geographic units were not designed to equalize access. They were drawn largely for administrative convenience (Axelrod et al., 2015). As a result, the DSA system produced large disparities across areas. Transplant rates differed by as much as a factor of 20 across DSAs (Yeh et al., 2011). For candidates with MELD scores between 38 and 39, the 90-day waitlist mortality rate ranged from 14 percent to 82 percent across DSAs (Massie et al., 2011). These disparities were inconsistent with federal regulation. In 2000, the Department of Health and Human Services issued the OPTN Final Rule, which provides more detailed regulations for the operation of the OPTN.[16] The rule states that place of residence and place of listing should not be major determinants of access to transplantation.

The Acuity Circles (AC) policy, implemented on February 4, 2020, aimed to reduce these disparities. The reform made two main changes. First, it replaced DSAs with concentric circles of 150, 250, and 500 nautical miles centered on the donor hospital, so that priority depends on distance rather than administrative boundaries. Second, the reform introduced finer groupings of medical urgency. Under the previous system, candidates were grouped into broad tiers based on Status 1A and MELD ranges. The new system divides candidates into narrower MELD bands, such as 37 to 40, 33 to 36, 29 to 32, 15 to 28, and below 15. Within each band, priority expands outward across distance-based circles centered on the donor hospital.

[15] At the time of the Acuity Circles reform, there were 58 DSAs. By mid-2026, this number had decreased to 55, due to mergers of OPOs: CTOP and MAOB in January 2021, DCTC and MDPC in January 2023, KYDA and OHOV in October 2024. The decertification of FLMP in September 2025 changed which OPO operates the DSA but not the number of DSAs. The Centers for Medicare & Medicaid Services (CMS) may grant a hospital a waiver to work with an OPO other than the one designated for its area if the waiver is expected to increase organ donation and ensure equitable treatment of patients in both service areas (Social Security Act §1138(a)(2); 42 C.F.R. §486.308(e)–(f)).

[16] CFR 42, Chapter 1, Subchapter K, Part 121

3. Data

This section describes the data sources, sample construction, and key variables used in the analysis.

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

I construct a dataset of all adult whole-liver offers made between April 2016 and December 2023.[17] Each observation represents a single offer to a candidate. I assemble the dataset by merging offer records with donor and candidate information using a unique donor and candidate identifier.

The primary source is the Potential Transplant Recipient (PTR) file. This file records the offer process, known as the matchrun, including how candidates are screened for compatibility, ranked under the allocation rule, and how they respond. Responses take four values: yes, no, bypass, and provisional yes. Because livers must be placed quickly, offers are sent to multiple candidates simultaneously, so lower-ranked candidates may not receive the liver even if they accept. A provisional yes indicates that a candidate has evaluated the offer and is interested in accepting the liver or receiving more information. I treat both yes and provisional yes as acceptance, and no as decline. A bypass occurs when the OPO skips higher-ranked candidates and offers the liver to a lower-ranked, but closer candidate who is more likely to accept, in order to limit cold ischemia time and avoid nonuse. Because bypassed candidates are not observed choosing, I exclude these observations.

I link the PTR file to two donor datasets. The first records the disposition of deceased-donor organs, which allows me to restrict the sample to livers that were recovered for transplant.[18] The second contains donor characteristics, including demographics and medical history. I then construct geographic measures, including GPS coordinates and zip codes, for donor hospitals. Between April 2016 and December 2023, 2,424 donor hospitals recovered at least one deceased-donor liver for transplant. Because distance between the donor hospital and transplant center affects allocation under the Acuity Circles policy, I compute the distance between every donor hospital and transplant center using GPS coordinates.[19] The SRTR does not report transplant center coordinates, so I

[17] I exclude livers from pediatric donors and segmented liver transplants because they follow distinct allocation rules.

[18] The data also include organs that were not recovered due to authorization issues and those recovered but not used for transplant.

[19] I use the "geodist" module in Stata, specifying the "sphere" option to calculate the distance between

collect these from publicly available addresses. Figure 3 maps the liver transplant centers and donor hospitals in the sample. Finally, I merge candidate-level data. These include demographics, such as age, race, educational attainment, and payment methods, as well as listing information and acceptance criteria.

The final sample contains 8,610,991 offers of deceased-donor whole livers. Table 1 reports descriptive statistics for donors, candidates and the offers made between them. The descriptive statistics and the empirical evidence in Section 4 use the full sample. To reduce computation time, the structural estimation in Section 5 uses a 20% random subsample of donors, together with all offers of their livers.

Panel A reports offer-level outcomes by medical urgency at the time of the offer. The distribution of offers across MELD tiers changes clearly after the Acuity Circles reform. The share of offers to high-urgency candidates declines, while the share of offers to low-urgency candidates increases. For example, the share of offers to Status 1A candidates falls from 0.98% to 0.52%, and the share to MELD 29 to 32 falls from 12.34% to 5.22%. In contrast, the share to MELD 6 to 14 candidates rises from 12.30% to 27.54%. The overall offer acceptance rate declines slightly from 37.56% to 35.82%, but acceptance rates by MELD tier move more sharply. The acceptance rate of Status 1A candidates falls from 20.57% to 12.30%, while the rate for MELD 6 to 14 candidates rises from 17.40% to 25.91%. Panel B reports candidate-level outcomes. The time from listing to first offer falls across all medical urgency groups after the reform. The average wait decreases from 14.86 days to 6.60 days. The largest reductions occur among lower priority candidates. For example, the wait for candidates with MELD 6 to 14 declines from 28.88 to 14.16 days. Panel C shows donor-level outcomes. The number of offers per donor increases sharply, from 86.56 before the reform to 165.24 after it. This increase reflects broader geographic sharing, which expands the pool of potential recipients and accelerates the pace of offer circulation.

Together, these patterns indicate that the Acuity Circles reform affected both the distribution of offers and candidate behavior. A decline in the share of offers to high-MELD candidates does not imply that their access fell. Because the number of offers per donor roughly doubled, high-urgency candidates may receive as many offers as before or more. Greater offer frequency reduces the cost of declining a given offer. As a result, these candidates may become more selective. This behavior can shift organs toward lower-MELD candidates, who had limited access under the previous system. These observations motivate the analysis that follows.

---

two points on a sphere.

4. Empirical Evidence

4.1. Empirical Strategy

My empirical strategy combines a measure of each candidate's exposure to the reform with the sharp timing of its implementation. I first construct a treatment intensity measure, ΔAccess, by simulating the allocation process under the old and new rules. I then compare how behavior and outcomes change after the reform across candidates whose access changed by different amounts.

4.1.1. Treatment Intensity

The Acuity Circles reform reallocates livers across candidates who differ in both location and medical urgency. Centers located near more donor hospitals under the new distance-based system may gain access, while centers with fewer nearby donors may experience smaller gains or losses. The introduction of finer MELD categories generates additional variation across candidates within the same center. As a result, the policy changes access along two dimensions, geography and medical urgency.

To quantify these changes, I define treatment intensity, ΔAccess, as the change in the probability that a candidate receives a first offer within a fixed period after listing. The measure is defined at the MELD-tier-by-DSA level:

$$\Delta\text{Access}_i = P_{AC}(T_i < x \mid MELD_i, DSA_i) - P_{DSA}(T_i < x \mid MELD_i, DSA_i). \tag{4.1}$$

Here, T_i is the number of days from listing to the first offer, $MELD_i$ is the candidate's MELD tier, and DSA_i is the DSA of the listing center.

I construct $P_{AC}(\cdot)$ and $P_{DSA}(\cdot)$ by simulating matchruns under each policy using pre-reform donor and candidate data. For each donor, the simulation produces two ranked candidate lists, one under the DSA rule and one under the AC rule. Because livers must be placed quickly, offers are issued in batches. I therefore classify a candidate as having received an offer if the candidate's simulated rank falls within the first batch. During the analysis period, about 94% of livers reach a final disposition within the first batch. In the observed data, the 95th percentile of the first-batch size is 296 candidates before the reform and 1,478 candidates under the new rule.

When constructing the ranks, I use the candidate's MELD score at the time of the offer when it is available in the matchrun data. When it is not available, I use the MELD score at listing. This substitution is unlikely to introduce substantial error because MELD scores change little over short horizons.

The simulation produces a binary outcome Y_i that equals one if candidate i receives a first offer within x days. I estimate the conditional probabilities in Equation (4.1) by

running separate linear probability models under each policy.[20] This procedure produces predicted probabilities for each of the 312 MELD-DSA cells, from which I calculate ΔAccess.[21]

I use a 22-day horizon for the baseline.[22] The estimated treatment intensity varies widely across cells, as shown in Figure 4. Access rises in 42.6% of the 312 MELD-DSA cells. The change differs by both MELD tier and geography. On average, candidates with MELD below 15 lose access, while candidates in the MELD 37 to 40 tier gain the most (Figure 5a). The effects also vary within MELD tiers across DSAs. Even among high-priority candidates, some areas gain substantially while others experience small losses (Figure 5b). The policy therefore generates meaningful two-dimensional variation in access.

4.1.2. Specification

I use the timing of the policy implementation to estimate how changes in access affect candidate behavior and outcomes. The baseline specification is:

$$Y_{igct} = \alpha + \gamma \Delta\text{Access}_{gc(i)} + \delta \left(\Delta\text{Access}_{gc(i)} \times \text{Post}_t\right) + \lambda_{c(i)} + \lambda_{t(i)} + \lambda_{g(i)} + \epsilon_{igct}. \quad (4.2)$$

The outcome variable Y_{igct} is defined at multiple levels depending on the specification. At the offer level, it is an indicator for whether candidate i accepts an offer. At the candidate or recipient level, it captures outcomes such as mortality, or the quality of the accepted liver. At the aggregate level, it represents center-level outcomes, such as registrations or multi-listing. $Post_t$ equals one after the reform. $\Delta\text{Access}_{gc(i)}$ is the predicted change in the probability of receiving a first offer within 22 days. The parameter of interest, δ, captures how outcomes change in cells with larger access gains after the reform. The specification includes center, MELD-tier, and year fixed effects. Standard

[20] The specification is

$$Y_i = \alpha + \sum_m \beta_m MELD_m + \sum_d \gamma_d DSA_d + \sum_{m,d} \delta_{m,d}(MELD_m \times DSA_d) + u_i,$$

where $MELD_m$ and DSA_d are indicator variables for MELD tier and DSA, respectively, and x ranges from 1 to 180 days. The outcome variable Y_i is defined as follows:

$$Y_i = \begin{cases} 1 & \text{if } T_i < x \text{ and not censored,} \\ 0 & \text{if } T_i \geq x \text{ and not censored,} \\ 0 & \text{if (days of listing)} \geq x \text{ and censored,} \\ \text{excluded} & \text{if (days of listing)} < x \text{ and censored.} \end{cases}$$

[21] There are six MELD tiers and 52 DSAs. During the analysis period, four DSAs (CTOP, MAOB, DCTC, MDPC) underwent mergers. I separate these merged DSAs into their original DSAs and treat them as distinct entities. Five DSAs without a liver transplant center are excluded.

[22] I use 22 days because the ΔAccess_i measure stabilizes after 21 days, with less than 5% variation thereafter.

errors are clustered at the MELD-by-DSA level.

To examine treatment effect dynamics, I also estimate an event-study specification:

$$Y_{igct} = \alpha + \gamma \Delta\text{Access}_{gc(i)} + \sum_{\tau=-4(\neq -1)}^{3} \delta_\tau \times \mathbf{1}\{t = \tau\} \times \Delta\text{Access}_{gc(i)} + \lambda_{c(i)} + \lambda_{g(i)} + \epsilon_{igct}. \quad (4.3)$$

Here, t indexes event time in years relative to the reform, and $\tau = -1$ is the omitted reference period. The coefficients δ_τ trace the relationship between predicted access gains and outcomes in each event year. I again cluster standard errors at the MELD-tier-by-DSA level.

4.2. Results

4.2.1. Candidate Behavior

(a) Offer acceptance. Table 2 reports estimates with offer acceptance as the dependent variable. Column (1) gives the average effect across all MELD tiers. The coefficient on Post × ΔAccess is -2.193 and is statistically significant. To convey magnitude, here and throughout, I scale estimates by a one-standard-deviation change in ΔAccess, which is 2.16 percentage points across candidates listed before the reform. A one-standard-deviation increase in ΔAccess lowers the probability of accepting an offer by 4.7 percentage points, a 12.6% decline relative to the pre-policy acceptance rate of 37.56%. This pattern is consistent with forward-looking behavior. When candidates expect better access to future offers, waiting becomes more valuable. Candidates therefore become more selective about the offers they accept.

Columns (2) to (7) examine heterogeneity across MELD tiers. The aggregate effect is driven largely by candidates with MELD 29 to 32. For this group, the coefficient is −6.024 and statistically significant, so a one-standard-deviation increase in access lowers acceptance by 13.0 percentage points for this tier. For the Status 1A candidates, the interaction term is omitted because ΔAccess varies little across DSAs. Status 1A candidates already receive top priority under both allocation rules, so their simulated change in access is small.

Figure 6a reports the corresponding event-study estimates. The coefficients become negative after the reform and grow in magnitude over time. This pattern suggests that candidates became more selective gradually as the policy took effect.

(b) Monthly registration. Table 3 examines new monthly waitlist registrations, aggregated to the center-by-MELD-group level. The dependent variable is a count. Each coefficient therefore gives the change in registrations per center, MELD tier, and month for a one-unit change in ΔAccess. All specifications include year, month-of-year, and

center fixed effects. Column (1) also includes MELD-tier fixed effects.

In Column (1), the coefficient on Post$\times\Delta$Access is 7.287. A one-standard-deviation increase in access implies 0.158 additional registrations per center–MELD group per month, a 13.2% increase relative to the pre-policy mean of 1.195. This result suggests that patients, referring physicians, or both initiate listings more often at centers experiencing larger improvements in access, conditional on medical urgency.

Columns (2) to (7) report heterogeneity across MELD tiers. The estimates are positive for MELD 6 to 14, 15 to 28, and 33 to 36, and negative for 29 to 32. Only the estimate for MELD 33 to 36 is statistically significant. For that group, a one-standard-deviation increase in access implies 0.109 additional registrations per month. As with acceptance, the Status 1A interaction is omitted because ΔAccess has little cross-DSA variation for this group. Figure 6b shows that registrations rise immediately and persistently in cells with larger access gains.

(c) Multi-listing. Table 4 examines multi-listing, which occurs when a candidate registers at an additional center while already active elsewhere.[23] The dependent variable is the annual count of new multi-listing registrations at the center-by-MELD-group level.

Column (1) reports a coefficient of 1.828, corresponding to 0.04 additional multi-listing registrations per year for a one-standard-deviation increase in access. The estimate is positive but not statistically distinguishable from zero. The tier-specific estimates in Columns (2) to (7) are also imprecise. The estimate for MELD 6 to 14 is positive but has a large standard error. The estimates for the other groups are negative and statistically insignificant. Figure 6c shows no systematic change after the reform. I find no evidence that changes in access affected multi-listing behavior, either on average or within any MELD subgroup.

4.2.2. Mortality, Survival, and Post-Transplant Outcomes

The behavioral responses above are intermediate outcomes. Whether improved access ultimately benefits candidates depends on how these responses affect survival and the types of transplants performed. I study four outcomes: overall mortality from listing, waitlist mortality, post-transplant graft failure, and the quality of transplanted livers.

Each outcome is measured over a fixed horizon. To ensure comparability, I restrict the sample in each specification to listings or transplants for which the full outcome horizon is observed within the sample period. The estimates can therefore be interpreted as event probabilities by the stated horizon.

[23] Multi-listing registrations account for 4.95% of registrations in the sample.

(a) Overall mortality from listing. Whether greater access benefits candidates is, at its core, a question about survival. First, I examine overall mortality from listing, which counts deaths regardless of whether they occur on the waitlist or after transplant. Because this outcome is defined for all listed candidates and does not condition on receiving a transplant, it is not directly affected by changes in the composition of waitlist or transplant recipients.

The dependent variable equals one if the candidate dies within the specified horizon following listing. Table 5 shows that the coefficients on Post $\times$ ΔAccess are negative and statistically significant at the 1% level at every horizon. The estimates range from -0.420 at three months to -0.540 at one year. A one-standard-deviation increase in predicted access reduces overall mortality by 0.91 percentage points at three months and 1.17 percentage points at one year, corresponding to 11.9% and 7.0% of the pre-policy means of 7.6% and 16.7%.

Figure 7a reports the corresponding event-study estimates at the three-month horizon. Patterns at longer horizons are similar (Appendix A). The effect emerges in the implementation year and grows over time. This gradual deepening mirrors the increase in selectivity documented in Figure 6a, suggesting that candidates and their centers adjusted to the new offer environment progressively rather than immediately at implementation.

(b) Waitlist mortality. To understand where the reduction in overall mortality arises, I next focus on deaths that occur before transplantation. The waitlist mortality measure equals one if a candidate is removed from the waitlist due to death or clinical deterioration within the specified horizon. Transplantation and other removals are treated as survival on the list. This outcome isolates the pre-transplant component of mortality.

Table 6 shows a similar pattern to the overall results. The coefficients are negative and statistically significant at the 1% level at all horizons, ranging from -0.391 at three months to -0.492 at one year. A one-standard-deviation increase in access reduces waitlist mortality by about 0.85 to 1.06 percentage points. These effects correspond to 13.6% and 8.4% of the pre-policy means of 6.2% and 12.6%. The magnitudes closely track those for overall mortality, indicating that nearly all of the survival gains operate through reduced deaths on the waitlist. The event-study estimates in Figure 7b show the same temporal patterns as overall mortality. The effects appear at implementation and grow over time, paralleling the dynamics in Figure 7a.

(c) Post-transplant graft failure. Table 7 reports estimates for an indicator equal to one if the liver fails within the specified horizon following transplant. Graft failure is defined as graft loss, re-transplantation, or recipient death.

In contrast to the mortality results, the coefficients on Post $\times$ ΔAccess are positive and statistically significant. The estimates increase from 0.224 at three months to 0.338 at

one year. Thus, a one-standard-deviation increase in predicted access raises graft failure by 0.48 to 0.73 percentage points, which represents roughly 8 to 9% of the pre-policy means of 5.3% to 8.9%.

This result likely reflects a change in the composition of transplant recipients rather than a deterioration in transplant quality. Graft failure is observed only for candidates who receive a transplant, and the reform changes who reaches transplantation. These estimates therefore describe changes in recipient cohorts rather than the effect of access on a fixed population. In groups that gained access, candidates who would have remained on the waitlist under the previous system now receive transplants. Within a MELD tier, these marginal recipients may have a poorer underlying prognosis than the average pre-reform recipient. This compositional shift can raise average graft failure among recipients even as overall mortality from listing declines.

The increase in graft failure therefore does not contradict the reduction in mortality. Graft failure and overall mortality measure different outcomes. First, graft failure is a composite outcome that includes non-fatal events, such as graft loss followed by re-transplantation, in addition to death. Second, graft failure is measured from the date of transplant, whereas overall mortality is measured from the date of listing. Third, graft failure is defined only for transplant recipients, while overall mortality is defined for the full population of listed candidates. For these reasons, graft failure and waitlist mortality do not mechanically add up to overall mortality.

Figure 7c reports the corresponding event-study estimates. The post-reform coefficients are positive in every year, matching the sign of the pooled estimates in Table 7, but they are individually imprecise. Each year-specific coefficient is estimated from only part of the post-period variation, so an effect of the modest size found in the table is not statistically distinguishable from zero year by year.

(d) Accepted-liver quality. A natural alternative explanation for the transplant outcome is that changes in access affected the quality of the livers ultimately transplanted. Table 8 examines three standard proxies for liver quality: donor age, an indicator for donation after circulatory death (DCD), and cold-ischemia time. Across these measures, the interaction terms are small and statistically insignificant. If anything, the point estimates move in the direction of higher quality, toward younger donors and shorter cold-ischemia time, although they are imprecisely estimated.

Taken together, these results provide no evidence that the rise in graft failure is driven by a deterioration in observable liver quality. Instead, they are more consistent with a shift in the composition of transplant recipients toward sicker candidates.

5. A Dynamic Model of Acceptance Decisions

The reduced-form evidence suggests that candidates adjust their acceptance decisions in a forward-looking manner when their access changes. However, these estimates alone cannot reveal the underlying forces driving this behavior or quantify the trade-offs candidates face when deciding whether to accept an offer or continue waiting. In particular, they do not identify how candidates value future transplant opportunities relative to the risks of remaining on the waitlist, nor do they allow us to predict how acceptance behavior would change under alternative allocation policies.

To address these limitations, I develop a dynamic model of offer acceptance. The model provides a framework for recovering the key decision parameters, and evaluating counterfactual policies that lie beyond the scope of the reduced-form analysis.

5.1. Environment

I model the waitlist as a continuous-time dynamic decision problem. Candidates face a sequence of opportunities to accept or decline liver offers that arrive over time. A candidate decides whether to accept or decline each offer, taking into account that future offers may arrive but are uncertain in both timing and quality.

Each candidate i is described by a state vector (x_i, t_i, d_i) that summarizes the key aspects of their situation on the waitlist:

- $x_i \in \mathcal{X}$: current health, measured by the MELD score. $\mathcal{X}$ is a finite set of MELD levels, including Status 1A. Health evolves stochastically over time, allowing for both improvement and deterioration.
- $t_i \in [0, T_i]$: time since listing. The horizon T_i is finite and varies across candidates, reflecting the fact that candidates do not remain on the waitlist indefinitely.
- $d_i \in \mathcal{D}$: the Donation Service Area (DSA) of the transplant center where the candidate is listed.

Each deceased-donor liver j is characterized by a scalar quality index $q_j \in [0, 1]$. Because SRTR data do not provide a single summary measure of liver quality, I construct q_j using a machine-learning model that predicts the probability that a recovered liver is transplanted rather than going unused, based on a rich set of donor characteristics (Appendix B). Higher values of q_j correspond to livers that are less likely to go unused and therefore more desirable.

For any candidate–donor pair (i, j), let $\ell_{ij} \geq 0$ denote the distance between the donor hospital and the candidate's transplant center. Distance matters because longer travel increases cold-ischemia time and can reduce effective organ quality at transplantation.

The allocation policy maps health and geographic information into a priority class $c_{ij} \in \mathcal{C}_r$, where r indexes the allocation rule. The precise definition of these classes varies across allocation systems. Table 9 summarizes the priority classes under both the DSA-based system and Acuity Circles rules.

5.1.1. Allocation Rule

Under a given allocation rule, each recovered liver is placed through the following sequence of steps:

Step 1 (Ordering). All feasible candidate–donor pairs (i, j) are assigned to a priority class c_{ij} based on the allocation rule. Priority depends primarily on the candidate's MELD score and geographic relationship to the donor. Candidates are ranked first by class and then by MELD score within class.

Step 2 (Offers). Because livers are time-sensitive, offers are not extended sequentially. Instead, they are sent simultaneously to a group of high-priority candidates.

Step 3 (Decision). Each candidate who receives an offer decides whether to accept or decline. Since multiple candidates may receive the same liver offer, more than one may choose to accept.

Step 4 (Assignment). The liver is assigned to the highest-priority candidate among those who accept. Candidates who accept but are not selected remain on the waitlist and continue to receive future offers.

Step 5 (Departures and transitions). Candidates leave the waitlist either through transplantation or other routes, including death, becoming too sick for transplant, receiving a living-donor liver, or transferring centers. While waiting, their MELD scores evolve over time and may increase or decrease.

5.1.2. Payoffs

Preferences are summarized by three primitives. First, $\Gamma(x_i, q_j, \ell_{ij})$ denotes the expected net present value of transplantation, which depends on the candidate's MELD score x_i, liver quality q_j, and distance ℓ_{ij}. Second, $D(x_i)$ is the expected net present value of departing the waitlist without a transplant. Third, $w(x_i, t_i)$ is the flow payoff from remaining on the waitlist.

Assumption 1.

(i) The value of transplantation is additively separable in an idiosyncratic shock:

$$\Gamma(x_i, q_j, \ell_{ij}) + \epsilon_{ij}.$$

(ii) The shocks ϵ_{ij} are independently and identically distributed across offer arrivals and follow the standard logistic distribution H.

(iii) The flow payoff $w(x_i, t_i)$ and the expected payoff from departing without a transplant $D(x_i)$ are exogenous and depend on the state only through (x_i, t_i) and x_i, respectively.

The additive separability in Assumption 1(i) and the independence assumption in (ii) are standard in the dynamic discrete choice literature (Rust, 1987; Hotz and Miller, 1993). Together, they allow the model to be estimated using the inversion approach of Hotz and Miller (1993), which avoids repeatedly solving the dynamic program during estimation.

The model abstracts from persistent unobserved heterogeneity on both the candidate and donor sides. This is motivated by two considerations. First, the liver quality index q_j is constructed from a rich set of observed donor characteristics. Second, the state vector $(x_i, t_i, q_j, \ell_{ij}, c_{ij}, d_i)$ captures the main determinants of priority and match value in this setting. Extending the framework to allow for a finite number of unobserved candidate types, as in Arcidiacono and Miller (2011), is a natural next step but is beyond the scope of the present analysis.

5.1.3. Arrivals, Departures and MELD Transitions

Time is modeled as continuous. Offer arrivals, departures and MELD transitions follow exogenous Poisson and jump processes, as in Arcidiacono et al. (2016) and the kidney waitlist application of Agarwal et al. (2021).

Assumption 2.

(i) Offer arrivals and pre-transplant departures are governed by independent, exogenous Poisson processes.

(ii) Offers in class $c \in \mathcal{C}_r$ arrive to candidate i at rate $\lambda_c(x_i; d_i)$.

(iii) The candidate departs the waitlist without transplantation at rate $\delta(x_i)$. In addition, each candidate has a finite terminal date $T_i < \infty$, at which departure occurs with probability one.

(iv) The MELD score evolves as a continuous-time Markov process: it jumps from x to $x' \neq x$ at exogenous rate $F_{xx'}$, where F is the transition matrix.

Following Agarwal et al. (2021), I assume that all candidates depart on or before their 100th birthday, so T_i corresponds to the time from listing until that date. Modeling time in continuous form removes the simultaneity that complicates discrete-time

dynamic models (Arcidiacono et al., 2016): at any instant at most one event occurs, so the likelihood factors cleanly across the rate processes. It also matches the institutional reality that offers arrive at irregular, stochastic times rather than on a fixed schedule.

5.2. A Single Candidate's Problem

Upon receiving an offer, a candidate chooses whether to accept it—and be transplanted if no higher-priority candidate also accepts—or to decline and wait for future offers. I treat this decision as the problem of a single agent who takes the primitive rates $\{\lambda_c, \delta, F\}$ and the assignment risk π as given.

5.2.1. Acceptance and Assignment

A distinctive feature of deceased-donor organ allocation is that accepting an offer does not guarantee transplantation. Because offers are sent to multiple candidates simultaneously, several candidates may accept the same organ. The liver is then assigned to the highest-priority acceptor.

I summarize this competition through an assignment probability

$$\pi_{ij} \;\equiv\; \pi(x_i, t_i, c_{ij}, q_j, d_i) \;=\; \Pr\big(i \text{ is assigned the liver } j \,\big|\, i \text{ accepts}\big).$$

Assignment probabilities depend on the candidate's priority, the characteristics of the offered liver, and the composition of competing candidates. At the time an offer arrives, candidate i observes the offer characteristics (c_{ij}, q_j, ℓ_{ij}) and the payoff shock ϵ_{ij}.

Assumption 3. For every state:

(i) Accepting an offer but not being assigned leaves the candidate's state unchanged and does not impose any cost.

(ii) The assignment probability satisfies

$$\pi(x_i, t_i, c_{ij}, q_j, d_i) > 0.$$

Let $V^r(x_i, t_i; d_i)$ denote the value of waiting under allocation rule r. Conditional on receiving an offer, accepting yields the value of transplantation, $\Gamma(x_i, q_j, \ell_{ij}) + \epsilon_{ij}$, if the candidate is assigned the liver (with probability π_{ij}) and the value of waiting, $V^r(x_i, t_i; d_i)$, otherwise. Declining yields the value of waiting with certainty.

The expected payoff from accepting is therefore

$$\pi_{ij}\left(\Gamma(x_i, q_j, \ell_{ij}) + \epsilon_{ij}\right) + (1 - \pi_{ij}) V^r(x_i, t_i; d_i),$$

while the payoff from declining is

$$V^r(x_i, t_i; d_i).$$

A candidate accepts whenever the expected payoff from accepting exceeds the value of declining:

$$\pi_{ij}\left(\Gamma(x_i, q_j, \ell_{ij}) + \epsilon_{ij}\right) + (1 - \pi_{ij})V^r(x_i, t_i; d_i) \geq V^r(x_i, t_i; d_i)$$

Rearranging gives the acceptance rule

$$\Gamma(x_i, q_j, \ell_{ij}) + \epsilon_{ij} \geq V^r(x_i, t_i; d_i)$$

This condition highlights an important feature of the model. Because candidates who accept but are not assigned remain on the waitlist without penalty, assignment risk does not affect the acceptance threshold itself. Candidates accept whenever the value of transplantation exceeds the value of waiting. Assignment risk instead affects the expected value of offers by changing the probability that an accepted offer results in transplantation.

In principle, assignment probabilities are equilibrium objects because they depend on the behavior of other candidates who receive the same offer. I do not model these strategic interactions explicitly. Instead, I treat assignment probability as a reduced-form object estimated from the data.

This approximation is well suited to the institutional setting. Offers are distributed in large batches that include candidates from many transplant centers. In the data, the 95th percentile of first-batch size is 296 candidates before the reform and 1,478 candidates under Acuity Circles. Although higher-priority candidates affect whether a particular candidate ultimately receives a liver, any individual candidate has a negligible effect on their own assignment probabilities. The model therefore adopts a competitive approximation in which candidates take assignment risk as given when making acceptance decisions, much as agents in large markets take prices as given.

5.2.2. Value Functions

Let $V^r(x, t; d)$ denote the value of remaining on the waitlist under allocation rule r for a candidate with health state x, waiting time t, and listing location d. The value function reflects four forces: the flow payoff from remaining active on the waitlist, changes in health over time, the risk of departure without transplantation, and future opportunities to receive offers. These forces determine the value of waiting through the Hamilton–Jacobi–Bellman equation

$$\rho V^r - \dot{V}^r = w(x,t) + \sum_{x' \neq x} F_{xx'}[V^r(x',t;d) - V^r(x,t;d)] + \delta(x)\left[D(x) - V^r(x,t;d)\right] + \sum_{c \in \mathcal{C}_r} \lambda_c(x,d) \iint \pi_{ij} \int \max\left\{\Gamma(x,q,\ell) + \epsilon - V^r(x,t;d), 0\right\} dH(\epsilon) dG_c(q,\ell), \tag{5.1}$$

where ρ is the discount rate and $\dot{V}^r \equiv \partial V^r / \partial t$. The final term represents the option value of future offers and is the primary channel through which allocation rules affect behavior. Changes in offer-arrival rates or offer composition alter the value of waiting and therefore shift the acceptance threshold.

5.2.3. Normalization

Choice data alone do not identify all payoff levels in dynamic discrete choice models, so the payoff from one action must be normalized at each state (Magnac and Thesmar, 2002). Following Agarwal et al. (2021), let $O(x,t)$ denote the value of always declining. Such a candidate receives the flow payoff from remaining on the waitlist and faces the same health transitions and departure risks as any other candidate, but never receives the payoff from transplantation. The value of this outside option satisfies

$$\rho O(x,t) - \dot{O}(x,t) = w(x,t) + \sum_{x' \neq x} F_{xx'}\left[O(x',t) - O(x,t)\right] + \delta(x)\left[D(x) - O(x,t)\right], \tag{5.2}$$

with terminal condition $O(x,T_i) = D(x)$. Under Assumptions 1(iii) and 2(i), $O(x,t)$ depends only on $\{w, D, F, \delta\}$ and is therefore exogenous.

The next proposition shows that subtracting $O(x,t)$ from both the value of waiting and the value of transplantation leaves behavior unchanged.

Proposition 1. *Let $V^r(x,t;d)$ solve Equation (5.1) with terminal condition $V^r(x,T_i;d) = O(x,T_i) = D(x)$, and let $O(x,t)$ solve Equation (5.2). Define*

$$\tilde{V}^r(x,t;d) \equiv V^r(x,t;d) - O(x,t), \qquad \tilde{\Gamma}(x,q,\ell,t) = \Gamma(x,q,\ell) - O(x,t).$$

Under Assumptions 1–3, $\tilde{V}^r$ satisfies

$$\rho \tilde{V}^r - \dot{\tilde{V}}^r = \sum_{x' \neq x} F_{xx'}\left[\tilde{V}^r(x',t;d) - \tilde{V}^r(x,t;d)\right] - \delta(x)\tilde{V}^r(x,t;d) + \sum_{c \in \mathcal{C}_r} \lambda_c(x,d) \iint \pi_{ij} \int \max\left\{\tilde{\Gamma}(x,q,\ell,t) + \epsilon - \tilde{V}^r(x,t;d), 0\right\} dH(\epsilon) dG_c(q,\ell), \tag{5.3}$$

with terminal condition $\tilde{V}^r(x, T_i; d) = 0$. *Moreover, the acceptance rule is invariant to* O*: for every offer,* $\Gamma + \epsilon \geq V^r$ *if and only if* $\tilde{\Gamma} + \epsilon \geq \tilde{V}^r$.

The proof is in Appendix C. Because the acceptance rule is invariant to $O(x,t)$, only value differences matter for behavior. I therefore normalize $O(x,t) \equiv 0$ for all $(x,t) \in \mathcal{X} \times [0, T_i]$ and interpret all subsequent value functions relative to the payoff from always declining offers. Under this normalization, $\tilde{\Gamma}(x, q, \ell, t) = \Gamma(x, q, \ell)$ and carries no time argument. With a slight abuse of notation, I drop tildes and write V^r and Γ for the normalized objects throughout the remainder of the paper.

5.3. Estimation

I estimate the model by expressing both the value of waiting, $V^r(x, t; d)$, and the value of transplantation, $\Gamma(x, q, \ell)$, in terms of conditional choice probabilities (CCPs). Estimation proceeds in two steps, following the CCP approach of Hotz and Miller (1993), Aguirregabiria and Mira (2002), and Arcidiacono and Miller (2011), adapted to continuous time as in Arcidiacono et al. (2016).

5.3.1. CCP Representation

Directly solving Equation (5.3) using a nested fixed-point algorithm (Rust, 1987) is computationally burdensome. The CCP approach avoids this by rewriting the value function in terms of observed choice probabilities, eliminating the need to solve the dynamic program inside the estimation loop.

Let P_{ijt} denote the conditional probability that candidate i, in waitlist state (x, t, d), declines an offer with characteristics (q, ℓ).

Proposition 2. *Suppose Assumptions 1–3 hold. Then,*

(i) The conditional choice probabilities satisfy

$$P_{ijt} = H(V^r(x_i, t; d_i) - \Gamma(x_i, q_j, \ell_{ij}))\,,$$

so that

$$V^r(x_i, t; d_i) - \Gamma(x_i, q_j, \ell_{ij}) = H^{-1}(P_{ijt}) = \log\left(\frac{P_{ijt}}{1 - P_{ijt}}\right). \tag{5.4}$$

(ii) There is a known function ψ *such that, for* $P_{ijt} \in (0, 1)$,

$$\psi(P_{ijt}) = \int \max\{\Gamma(x_i, q_j, \ell_{ij}) + \epsilon_{ij} - V^r(x_i, t; d_i),\, 0\}\; dH(\epsilon) = -\log P_{ijt}. \tag{5.5}$$

The proof is in Appendix C. Part (i) follows directly from the logistic assumption on ϵ_{ij}, and (ii) evaluates the option value in closed form. Substituting Equation (5.5) into Equation (5.3) yields

$$\rho\, V^r(x,t;d) - \dot{V}^r(x,t;d) = \sum_{x' \neq x} F_{xx'}[V^r(x',t;d) - V^r(x,t;d)] - \delta(x) V^r(x,t;d) + \sum_{c \in \mathcal{C}_r} \lambda_c(x,d) \iint \pi_{ij}\ (-\log P_{ijt})\ dG_c(q,\ell), \tag{5.6}$$

so that the offer term no longer depends on the value function inside the integral.

Given estimates of $\{\delta(\cdot), \lambda_c(\cdot), F, \pi\}$ and the CCPs P_{ijt}, the right-hand side of Equation (5.6) is known. The value function V^r can therefore be recovered by backward recursion. The transplant value is then obtained pointwise from Equation (5.4):

$$\Gamma(x_i, q_j, \ell_{ij}) = V^r(x_i, t; d_i) + \log\left(\frac{1 - P_{ijt}}{P_{ijt}}\right), \tag{5.7}$$

without solving the dynamic program during estimation.

5.3.2. Identification

Identification exploits the fact that the data record the full sequence of offers, acceptance decisions, health transitions, and waitlist outcomes for all candidates. These data allow the offer process, departure process, health dynamics, and choice behavior to be observed separately, which simplifies identification of the key primitives.

The rate primitives $\{\lambda_c(x;d), \delta(x), F\}$ are identified from observed event histories. In continuous time, the timing and type of events provide separate information about offer arrivals, departures, and MELD changes. Variation in the duration between successive events, together with the realized event type, identifies the corresponding hazard rates under Assumption 2.

Assignment risk $\pi(x,t,c,q,d)$ is identified from the distinction between acceptance and transplantation. Conditional on observed candidate states and offer characteristics, assignment risk is the probability that an accepted offer results in transplantation. Because the data separately record whether a candidate accepts an offer and whether that acceptance leads to a transplant, assignment probabilities can be recovered directly from observed outcomes.

Preferences are identified from variation in acceptance behavior across candidate states and offer characteristics. Under Assumption 1, the logit structure implies that observed conditional choice probabilities identify the value difference between accepting and waiting:

$$V^r(x,t;d) - \Gamma(x,q,\ell).$$

Intuitively, candidates facing identical states but different offers reveal how offer quality and distance affect the value of transplantation, while differences in acceptance behavior across candidate states reveal variation in the value of waiting.

Choice data identify only value differences rather than absolute payoff levels. The normalization addresses this issue by measuring all values relative to the payoff from always declining offers. Given $\{\lambda_c, \delta, F, \pi\}$ and observed CCPs, the normalized value function V^r is identified as the unique solution to Equation (5.6) subject to the terminal condition. The transplant value Γ is then recovered pointwise from Equation (5.7).

Finally, the discount rate ρ is not identified from observed choices alone in dynamic discrete choice models (Magnac and Thesmar, 2002). I therefore treat ρ as fixed throughout the analysis.

5.3.3. Estimation Procedures

I estimate the model in two steps, following the CCP approach of Hotz and Miller (1993) adapted to continuous time. I fix the discount rate at $\rho = 0.05$ per year.

First step: reduced-form objects. I begin by estimating the primitives governing offer arrivals, departures, health dynamics, assignment risk, and acceptance behavior.

Arrival, departure, and MELD transitions. Using matchrun data, I construct candidate-level event histories and estimate the offer-arrival rates $\lambda_c(x; d)$, the departure rate $\delta(x)$, and the MELD transition matrix F jointly by maximum likelihood. The likelihood combines information on event timing and event type (offer arrival, departure, MELD transition, or censoring). Arrival rates vary by priority class, depend on MELD score within class, and differ across DSAs. Departure and transition rates are allowed to vary flexibly with MELD.

Assignment risk. I estimate $\pi = \Pr(\text{transplant} \mid \text{accept})$ nonparametrically from accepted offers. Conditional on observed candidate states and offer characteristics, assignment risk is measured as the empirical frequency that an accepted offer results in transplantation.

Conditional choice probabilities. I estimate the CCPs using a logit specification that includes flexible functions of MELD (x), waiting time (t), liver quality (q), and distance (ℓ), together with DSA fixed effects.

Second step: value recovery. Given estimates of $\{\lambda_c, \delta, F, \pi\}$ and the CCPs, I recover the value function using Equation (5.6). Because the equation is linear in V^r, I solve for the value of waiting by backward recursion, using the terminal condition $V^r(x, T_i; d) = 0$.

For each offer, I recover Γ using the inversion formula in Equation (5.7). I then parameterize $\Gamma(x, q, \ell)$ as a flexible function of MELD, liver quality, and distance, and

estimate the corresponding parameters by regressing the recovered values on these covariates. Standard errors account for first-stage estimation uncertainty.

5.4. Results

5.4.1. Parameter Estimates

First-stage Estimates. Figure 8 reports estimated offer-arrival rates across MELD scores, comparing the model to the data before and after the reform. The reform shifts the entire curve upward: candidates at nearly all MELD levels receive offers more frequently, with especially large increases at higher MELD scores. This reflects two changes that operate together. First, broader geographic sharing lets candidates receive offers from donors far beyond their own DSA, so each candidate draws from a much larger donor pool. High-MELD candidates gain the most, because they now rank near the top for many more donors. Second, batches become larger: the average number of offers per donor nearly doubles after the reform (Table 1). The allocation rule itself does not fix the batch size. Instead, it is an operational choice of the OPO, which depends on the quality of the liver and plausibly responds to the higher decline rates that broader sharing creates.

Figure 9 reports the estimated hazard of departure from the waitlist without transplantation. This risk is low and relatively flat up to a MELD score of about 25, and then rises sharply between MELD 30 and 40. Figure 10 shows the estimated transition dynamics of MELD. Transition rates increase with medical urgency: both improvement and deterioration occur more frequently at higher MELD levels, but worsening becomes more likely at the top of the distribution.

Figure 11 reports assignment probabilities conditional on acceptance. For candidates below MELD 30, fewer than 5% of accepted offers result in transplantation. Assignment probabilities increase sharply above MELD 35, but remain roughly 25% even at the highest levels. Thus, for most candidates, accepting an offer confers a chance at transplantation rather than a guarantee. In the model, this affects the value of accepting—through the expected surplus—but not the acceptance threshold itself. The reform also alters assignment risk, and in different directions at different MELD levels. A candidate's assignment probability depends on the higher-ranked candidates in the batch. The reform changes both the number of these competitors and their behavior. Broader sharing places more high-MELD candidates in each batch, which lowers assignment probabilities. At the same time, these high-MELD candidates become more selective after the reform, which raises assignment probabilities for candidates ranked below them. For candidates in the MELD 30–35 range, the second force dominates. At the highest MELD levels, the first force dominates.

Table 10 reports estimates from a flexible logit model of the decline decision. The outcome equals one if the offer is declined, and continuous covariates are standardized, so each coefficient gives the change in the log-odds of declining for a one-standard-deviation change in the covariate. Column (1) reports pre-reform estimates and column (2) reports post-reform changes, so the post-reform effect is the sum of the two.

Liver quality is the strongest predictor of the decision. Evaluated at the mean of q, a one-standard-deviation increase in quality reduces the odds of declining by a factor of about four, since $e^{-1.373}$ is approximately 0.25. Starting from the pre-reform decline rate of 62% (Table 1), this corresponds to the probability of declining falling to roughly 30%. The negative coefficient on q^2 means the effect is stronger at higher quality levels. Sicker candidates respond less to quality. The positive interaction between MELD and q implies that, for a candidate one standard deviation above the mean MELD, the quality effect falls from -1.373 to -1.141, so the same improvement in quality reduces the odds of declining by a factor of about three rather than four. This weaker response is consistent with a lower option value of waiting when the risk of dying without a transplant is high.

Distance works in the opposite direction, and its importance depends strongly on urgency. For the least urgent candidates, with MELD between 6 and 14, a one-standard-deviation increase in log distance nearly doubles the odds of declining, since $e^{0.615}$ is approximately 1.85. The distance penalty declines monotonically with MELD and essentially disappears for candidates with MELD 35 and above, where the interaction term offsets the main effect almost exactly. Waiting time plays a limited role. The coefficients on the waiting time terms are small, and a one-standard-deviation increase in waiting time changes the odds of declining by only a few percent. This magnitude is far smaller than the effects of quality and distance. This is consistent with the design of the liver allocation rule, which, unlike kidney allocation, does not reward accumulated waiting time.[24] Because priority depends on current MELD and geographic proximity to the donor rather than time on the waitlist, time already spent on the list carries little payoff-relevant information.

Column (2) shows how these patterns change after the reform. The quality effect at the mean weakens, from -1.373 to -1.069, while the effect at very high levels of q becomes stronger. Candidates discriminate somewhat less across ordinary livers but more sharply among the best ones. The role of distance also changes with medical urgency in a specific way. The distance penalty falls for the healthiest candidates, from 0.615 to 0.445, but rises for mid-urgency candidates, from 0.464 to 0.616 for MELD 15–20 and from 0.401 to 0.501 for MELD 21–27, leaving these the most distance-averse group after the reform. For MELD 28–34, the penalty is essentially unchanged, and for MELD 35 and above the point estimate turns slightly negative: distance does not deter the sickest candidates

[24] Waiting time enters liver allocation only as a tie-breaker among candidates with the same score.

either before or after the reform.

Figures 12–15 translate the coefficients into choice probabilities. They plot the estimated CCPs across key state variables, averaged over the offers that candidates actually receive. Candidates decline most offers. Those with MELD below 15 decline more than 90% of offers. These candidates are relatively healthy, so the benefit of transplantation is small relative to the value of continued waiting. The decline probability falls sharply at MELD 15, the score at which transplantation is generally judged to provide a survival benefit.[25] It remains roughly flat through the 20s and reaches its lowest point around MELD 30. It then rises again toward the top of the distribution, where frequent offers raise the option value of waiting. Therefore, acceptance is highest in the middle of the MELD distribution, where urgency is high but offers remain scarce.

The figures also show that decline rates rise by roughly ten percentage points at moderate MELD levels after the reform. This increase is larger than the post-reform coefficient changes in Table 10 alone would imply, because the figures average over the actual post-reform offer mix. Under broader sharing, candidates receive offers from a wider geographic area, so the offers are more distant on average, and Table 10 shows that distant offers are more likely to be declined. The rise in decline rates therefore reflects both the change in the offers candidates receive and the change in how they respond to them.

The Value of Transplantation. Table 11 reports the second-stage estimates of $\Gamma(x, q, \ell)$ and Figure 16 plots the implied value of quality and cost of distance. The level of the transplant value has no natural interpretation so I describe the estimates in relative terms.

Liver quality is the dominant determinant of transplant value, and its effect is strongly convex. For a candidate with MELD 20 receiving an offer at a distance of 250 nautical miles, a marginal increase in quality is worth roughly seven times more for a high-quality liver ($q = 0.75$) than for a low-quality one ($q = 0.25$).[26] The difference between a good liver and an excellent liver matters far more than the difference between a poor liver and an adequate one. This pattern is consistent with the first-stage acceptance behavior, where candidates discriminate most sharply among the best livers.

These effects vary systematically with medical urgency. The negative MELD-by-quality interaction implies that the marginal value of quality declines steadily with urgency. At mid-range quality ($q = 0.5$), it falls by roughly a factor of five between MELD

[25] For candidates below MELD 15, the mortality risk of transplant surgery exceeds the mortality risk of remaining on the waitlist (Merion et al., 2005). Clinical practice guidelines accordingly use MELD 15 as the threshold for recommending transplant evaluation (Martin et al., 2014).

[26] Marginal effects are computed from the coefficients in Table 11, with distance evaluated at 250 nautical miles, the radius of the second Acuity Circles allocation boundary. The marginal value of quality is $\partial\Gamma/\partial q = 1.6476 + 18.7612q - 0.2300(\text{MELD}) - 0.0299\ln(1 + \ell)$, which, at MELD 20 and $\ell = 250$, equals 1.6 at $q = 0.25$ and 11.0 at $q = 0.75$.

10 and MELD 40.[27] The positive MELD-by-distance interaction points the same way. The penalty from a more distant liver falls by roughly half over the same range.[28] For the sickest candidates, the value of transplantation is much less sensitive to the liver's quality and distance, although it still rises with quality.

5.4.2. Welfare Comparison

To quantify the welfare effects of the reform, I follow Agarwal et al. (2021) and express welfare changes in arrival-equivalent units: the proportional increase in offer-arrival rates under the DSA system that would generate the same change in value as the policy itself.

For a candidate listed in state (x_i, d_i), let

$$dV_i = V^{AC}(x_i, 0; d_i) - V^{DSA}(x_i, 0; d_i)$$

denote the change in value at listing induced by the reform. If λ' denotes the arrival rate that would generate the same change in value under the DSA system, the arrival-equivalent gain satisfies

$$\frac{dV_i}{V^{DSA}(x_i, 0; d_i)} \approx \frac{\lambda' - \lambda}{\lambda},$$

where the approximation follows because the value of waiting is approximately linear in the arrival rate (Equation (5.1)). By construction, the measure is invariant to the scale of utility, so welfare changes can be compared across candidates. This metric is analogous to an equivalent-variation measure evaluated at listing and provides a common and interpretable unit for comparing different channels of welfare change. A useful feature is that the decomposition is additive, so each component can be interpreted in the same arrival-equivalent units.

To understand sources of welfare change, I decompose dV_i into four channels: offer-arrival rates (λ), assignment probabilities (π), offer composition (q, ℓ), and candidate behavior (P). The first three channels hold the acceptance threshold at its pre-reform optimum and change the environment candidates face. The final channel changes the acceptance rule itself. It measures the gain from replacing the pre-reform acceptance threshold with the optimal threshold under the new environment. The channels are constructed sequentially so each component is the incremental value of switching that feature given the earlier switches, and the components sum to dV_i by construction. Because arrival rates and assignment probabilities jointly determine access to transplantation, I group them into a single effective-arrival channel.

[27] At $q = 0.50$ and $\ell = 250$, $\partial\Gamma/\partial q$ equals 8.56 at MELD 10 and 1.66 at MELD 40.

[28] The marginal cost of distance is $\partial\Gamma/\partial \ln(1+\ell) = -0.3907 - 0.0100 \ln(1+\ell) + 0.0071(\text{MELD}) - 0.0299q$, which falls in magnitude from 0.39 at MELD 10 to 0.18 at MELD 40 at $q = 0.5$.

Aggregate effects. The welfare gains are large (Figure 17). On average, the reform increases welfare at listing by 1.06 arrival-equivalent units, corresponding to approximately a doubling of donor supply. Most of this gain operates through the effective-arrival channel, which contributes 1.22 units. This channel reflects two large but offsetting forces. Broader sharing dramatically increases offer opportunities, producing a gain equivalent to more than tripling donor supply (3.10). At the same time, increased competition reduces the probability that an accepted offer results in transplantation, generating an offsetting effect of −1.88. The net effect is strongly positive because the increase in offer opportunities dominates the decline in assignment probabilities.

In contrast, changes in offer composition contribute relatively little to welfare. Candidates receive livers that are, on average, slightly lower quality (−0.08) and recovered from greater distances (−0.09). These effects modestly reduce welfare but are small relative to the gains from increased access. The behavior channel contributes almost nothing (0.01). This implies that the decisions that change under the new threshold involve offers whose value is nearly the same as the value of waiting, so reversing them changes welfare little. The reform therefore affects welfare primarily through quantity rather than quality: candidates benefit because transplant opportunities become substantially more abundant.

Heterogeneity by MELD. Figure 18 shows that welfare gains differ across levels of medical urgency. Gains are positive throughout the MELD distribution but are not monotonic. They are approximately 1.4 arrival-equivalent units at MELD 6, decline to about 0.7 near MELD 18, increase to roughly 1.8 at MELD 30–32, and then fall at the highest MELD levels. Low-MELD candidates gain despite losing relative priority because the value at listing reflects future health states, in which the reformed environment delivers far more offers. The effective-arrival channel is positive throughout the distribution. The dip at intermediate MELD levels is driven primarily by the composition channel. Candidates with MELD scores between 14 and 32 are most exposed to longer travel distances and slightly lower quality livers. The behavior channel contributes little to welfare across the distribution.

Figure 19 further decomposes the effective-arrival channel into arrival-rate and assignment-probability components. For low-MELD candidates, both effects are large and largely offsetting. These candidates receive substantially more offers under broader sharing, but they also face lower assignment probabilities because they compete in a larger pool. As MELD score increases, both effects weaken. For candidates around MELD 27–32, the assignment channel becomes positive, reinforcing the increase in arrival rates and generating the largest welfare gains. At the highest MELD levels, gains are smaller, consistent with the fact that these candidates were already highly prioritized under both systems.

Who loses, and where. Despite the substantial redistribution of livers, the reform is close to a Pareto improvement. Only 1.6% of candidates experience a welfare loss at listing. These losses, however, are geographically concentrated. In five DSAs, nearly half of candidates (48.8%) are worse off, while the fraction is negligible elsewhere (Table 12).

These DSAs tend to be relatively organ-rich, with higher donor supply relative to demand and shorter waiting times. Under the DSA system, candidates in these areas benefited from favorable local supply conditions. Broader sharing reallocates toward higher-priority candidates in other parts of the country, reducing these local advantages. As a result, welfare gains are smaller in these areas, and a subset of candidates experiences net losses.

5.5. Counterfactuals

I use the estimated model to evaluate allocation rules other than the one observed. I consider two exercises. The first looks back at the 2020 reform and asks which design features produced the welfare gains in Section 5.4.2. The second looks forward to a rule that has not been implemented for livers. That rule replaces discrete geographic and urgency tiers with a single continuous priority score. In both exercises the estimated primitives are held fixed and only the allocation rule changes. For each counterfactual rule, I solve for a steady state in which candidate values, acceptance decisions, and assignment probabilities are mutually consistent.

I adopt one convention throughout this section. In the observed data, the number of offers a donor generates varies across donors. The entire distribution shifts to the right when sharing is broadened, as shown in Appendix Figure A5. The number of offers per donor, which I call batch size, is an operational choice of the OPO rather than a written feature of allocation policy. Because all offers in a batch are sent simultaneously, the batch is assembled before any candidate in it responds, so batch size is not determined by the acceptance decision I model. I therefore treat it as an exogenous feature of each allocation regime. Batch size is held fixed within each rule at 87 offers per donor under the pre-reform rules and 165 under Acuity Circles. These values match the observed averages in Table 1, Panel C. For Continuous Distribution, which has not been implemented for livers, I recover the corresponding number in Section 5.5.2. I recompute V^{AC} and V^{DSA} using the fixed batch size so that all rules are evaluated in a common environment.

Because the baselines are simulated in this common environment, the numbers reported for Acuity Circles in this section differ from those in Section 5.4.2. That section measures the change candidates experienced between the observed pre-reform and post-reform environments. This section re-runs allocation on the same donors and candidates under alternative rules. All comparisons below are made within this simulated environment. As before, I express welfare changes in arrival-equivalent units, now relative to the

simulated DSA baseline.

5.5.1. Components of Acuity Circles

Acuity Circles changed three parts of the allocation system at once. It replaced the administrative DSA boundary with circles of fixed radius around the donor hospital. It split the wide MELD priority bands into narrower ones. It also increased the number of candidates who receive a simultaneous offer for a given liver. Section 5.4.2 studies the channels through which welfare changed. This section asks which of the three design features generated that change. Each lever is set to either its pre-reform or post-reform value. This creates a $2 \times 2 \times 2$ design with eight allocation rules, of which the DSA system and Acuity Circles are two. Panels A and B of Table 13 report the results. Figure 20 plots the average welfare gain against the share of candidates made worse off.

What each lever does. Circles and bands move offers between candidates, but along different dimensions. Circles reassign offers geographically, while bands reassign them by medical urgency. Introducing circles, holding the other levers fixed, replaces the DSA boundary with a 150-nautical-mile radius around the donor hospital. A candidate just outside the DSA gains priority, and a candidate far from the donor but inside the DSA loses it. Therefore circles redistribute offers by distance. Narrower bands under the DSA boundary move a low-MELD candidate in the donor's DSA down the sequence, because higher-MELD candidates outside the DSA are now prioritized over them. Bands shift offers toward sicker candidates. The larger batch adds offers to new candidates without taking any away.

Circles are the largest single lever, for both welfare and redistribution. Replacing the DSA boundary with circles raises welfare by 0.332 arrival-equivalent units, equal to a 33% increase in donor supply, which is about half the gain of the full reform. Circles also make more candidates worse off than any other lever, 34.2%, the highest share among the eight rules. Because circles move livers geographically, average recipient distance falls by 12.5% relative to the DSA baseline, while recipient MELD score changes little (+0.7).

Narrower bands are close to a pure transfer. On their own, they raise welfare by 0.051 units, less than one-tenth of the full reform, and make 27.2% of candidates worse off. Bands redistribute offers by MELD score, yet recipient MELD is almost unchanged relative to the DSA rule (+0.2). The offer moves to a sicker candidate, but the transplant does not. Figure 12 shows why. The probability of declining an offer rises above MELD 30, so the sicker candidates moved to the front would decline more often, and the liver passes down the sequence to a candidate close to the one who would have received it under the DSA system. Because bands send more offers outside the DSA, average recipient distance rises by 4.9%.

The larger batch does not change any candidate's position in the queue. Every candidate who received an offer under the DSA system still receives one, and 78 additional candidates per donor receive one as well. Thus, no candidate is worse off, which is why the share of candidates made worse off is 0.0%. Welfare rises by 0.231 units, coming entirely from the candidates added at the new positions. Average recipient distance rises from 144 to 182 nautical miles because the additional offers reach candidates further down the list, who are on average farther from the donor hospital. One caveat applies to this lever. Making additional simultaneous offers consumes OPO, transplant center, and candidate resources but the model does not price this cost.

How the levers relate. Narrower bands and circles are complements. Added to the DSA system alone, bands raise welfare by only 0.051 units, and combined with the larger batch they only add 0.046. Combined with circles, however, they add approximately 0.11, with or without the larger batch. Bands therefore contribute roughly twice as much to welfare when circles are in place, and the combination also leaves fewer candidates worse off.

The larger batch does not interact with the other levers in welfare terms. With or without the other levers, it adds roughly 0.231 units. Its role is instead to recover some of the candidates whom bands or circles push out of the batch. Added to bands, it lowers the share of candidates made worse off from 27.2% to zero. Added to circles, it lowers the share from 34.2% to 25.1%. The additional offers are enough to cover the candidates displaced by narrower bands, but not all of those displaced by circles.

Acuity Circles. With all three levers on, welfare rises by an amount equivalent to a 65% increase in donor supply, and 23.1% of candidates are worse off. Each number traces to a different lever. The share of candidates made worse off comes mainly from circles, while the welfare gain is a composite to which circles contribute the largest part. These results show why the OPTN introduced all three design changes at once. Circles alone generate a large welfare gain but leave a substantial share of candidates worse off. Adding narrower bands and the larger batch raises the welfare gain above circles alone while shrinking the share of candidates made worse off.

5.5.2. Continuous Distribution

Although Acuity Circles broadened sharing, priority still changes discontinuously at geographic boundaries and MELD thresholds. A natural alternative is a continuous distribution system that replaces these discrete cutoffs with a single priority score that combines medical urgency and geographic proximity. This approach has already been adopted in lung allocation and is currently under consideration for liver allocation.

This section uses the estimated model to evaluate such rules. Continuous Distribution can increase access for high-urgency candidates regardless of their location. However, it may also increase transport distances, raise cold-ischemia time, and reduce the value of transplantation. The model quantifies this trade-off by solving for the steady state under alternative urgency-proximity weights.

Under a Continuous Distribution rule, each candidate–donor pair is assigned a scalar priority score:

$$S(x, \ell) \;=\; w_u\, x \;+\; w_d\, g(\ell), \qquad w_u, w_d \geq 0, \quad w_u + w_d = 1. \tag{5.8}$$

Here, $g(\ell)$ is a decreasing function of distance. The weights determine the relative importance of medical urgency and geographic proximity in allocation.[29] A liver is offered according to the national ranking induced by S. The case $w_u = 1$ corresponds to allocation based solely on medical urgency, while $w_d = 1$ corresponds to allocation based solely on proximity. I evaluate the grid $w_u \in \{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$ spanning these two extremes.

The counterfactual requires a batch size. I do not hold it at the Acuity Circles value, because larger batches are part of the practical response to broader sharing. Each time the sharing area has expanded, OPOs have increased the number of simultaneous offers, and I assume a Continuous Distribution rule also would lead to the same adjustment. Because no such policy exists for livers, I recover the batch size from lung allocation, the only organ operating under Continuous Distribution.[30] Lungs provide a credible benchmark because the two organs responded almost identically to the preceding reform: when distance-based circles replaced administrative boundaries, offers per donor rose by a factor of 1.88 for lungs and 1.91 for livers.[31] I assume the same proportional response under Continuous Distribution. Applying the lung response of 1.35 to the liver Acuity Circles value yields 224 offers per donor, which I hold fixed across all weight combinations.

Panel C of Table 13 reports outcomes for six weight combinations and Figure 21 plots the welfare gain against the share of candidates made worse off. Raising w_u shifts the allocation along two margins at once. It directs livers toward sicker candidates. Average recipient MELD rises from 16.3 under pure proximity ($w_u = 0$) to 26.9 under pure urgency ($w_u = 1$). It also moves livers farther. Average donor–recipient distance rises from 71 to 828 nautical miles.

[29] Medical urgency and proximity are normalized to a comparable scale before applying the weights.

[30] Implemented March 9, 2023.

[31] Author's calculation from SRTR data. Average batch size rose from 50.48 to 94.77 for lungs and from 86.56 to 165.24 for livers after moving from DSA boundaries to distance-based circles. Under Continuous Distribution, lung batch size increased to 128.25.

Extremes are not optimal. Neither end of the grid performs well, and each fails for a different reason. Pure proximity produces a large average welfare gain, but it also generates the largest share of candidates made worse off, 59.1%. Under any rule with $w_u > 0$, a candidate who declines an offer moves up the ranking as their MELD score rises, so waiting buys priority. Under pure proximity, the score depends only on distance, which does not change with health. Health deterioration no longer improves a candidate's position, and the dynamic incentive to wait disappears. Candidates therefore accept earlier, which pulls a large share of low-MELD candidates into the recipient pool: 46.2% of transplants go to candidates with MELD 6-14. These are candidates who would not have received offers, or would have waited years, under the DSA system, so their baseline value is low and any transplant is a large proportional gain. Therefore, the average welfare gain is high not because livers reach the candidates who value them most, but because the gainers start from so little. Every liver that goes to a low-MELD candidate is one a sicker candidate would otherwise have received, so pure proximity redistributes welfare toward the healthy.

Pure urgency fails on the other margin. It delivers the largest average welfare gain on the grid, but it also wastes the most livers. Under pure urgency, the ranking ignores location, so every liver is offered first to the sickest candidates in the country wherever they are. As a result, even accepted livers travel 828 nautical miles on average, more than five times the distance under the DSA baseline. Distance lowers the value of the liver on offer, and Figure 14 shows that at 828 nautical miles the estimated probability of declining is about 90%. The candidates the rule prioritizes turn down most of what it sends them, and if every candidate in the batch declines, the liver goes unused. The nonuse rate exceeds the DSA baseline by 8.5 percentage points, the largest increase on the grid. Those livers were transplanted under the DSA system, so their loss falls on the candidates who would have received them, and the share of candidates made worse off, which falls steadily from 59.1% under pure proximity to 11.9% at (0.8, 0.2), rises to 13.1%. The welfare gain under pure urgency is concentrated on the sickest candidates who do accept, and it is paid for in livers that are never used.

Interior rules trade welfare against utilization. Moving weight from proximity to urgency changes who rises to the top of the ranking and how far the liver must travel to reach them. The first effect raises welfare because a sicker candidate values the liver more. The second effect lowers acceptance, because a liver that travels farther is worth less and is more likely to be declined. Which effect dominates depends on how far away the sickest candidates actually are.

At low values of w_u, the two effects do not conflict because distance still carries enough weight that the sickest candidates in the ranking are also nearby ones. Raising w_u from 0 to 0.4 increases average recipient MELD from 16.3 to 22.1 while average donor–recipient

distance barely moves, from 71 to 84 nautical miles. Livers go to sicker candidates without traveling farther, so welfare rises and organ nonuse falls together. At (0.4, 0.6) the nonuse rate is 10.5 percentage points below the DSA baseline, the lowest on the grid. Beyond that point, distance binds. Putting more weight on urgency now reaches sicker candidates only by sending livers farther, and average distance rises to 139, then 325, then 828 nautical miles. Acceptance falls with it, and the nonuse rate turns up. Welfare still rises because the sicker recipients who do accept gain more than the extra travel costs, but each step now buys that gain with more livers left unused.

Therefore, the trade-off appears only once w_u exceeds 0.4. Moving from (0.6, 0.4) to (0.8, 0.2) raises the welfare gain and lowers the share of candidates made worse off, but it also raises the nonuse rate and lengthens average time on the list. Neither rule dominates the other, and the choice between them depends on how much weight the planner places on livers that go unused.

Continuous Distribution improves on Acuity Circles. At $(0.6, 0.4)$ Continuous Distribution improves on Acuity Circles on welfare, organ nonuse, the share of candidates made worse off, and time on the list. It does so at nearly the same average recipient MELD, and with livers traveling no farther. This implies that the improvement does not come from transplanting more urgent candidates or from changing how far livers travel. It comes from replacing a hard geographic boundary with a continuous national ranking. A boundary rule is lexicographic, so a candidate inside the circle is ranked ahead of a candidate outside it however much sicker the outside candidate is. Under the continuous ranking, a liver is offered to the highest-priority candidates in its vicinity rather than to the highest-priority candidates inside a hard boundary. Taken together, the results suggest that a moderate Continuous Distribution rule improves on current policy, and that weighting urgency too heavily imposes its cost in the form of wasted livers.

6. Conclusion

When allocation rules for scarce goods are redesigned, the discussion usually centers on who should receive them. What this overlooks is that people are forward-looking. Even in a setting with extremely high stakes, transplant candidates decline what they are offered to wait for a better one. Yet the simulation model the OPTN uses to evaluate proposed policy designs holds candidates' behavior fixed at its historical pattern (Weaver and Pyke, 2023). This paper measures how the accept-or-decline decision responds to the allocation rule and evaluates liver allocation policy with that response built in.

I find empirical evidence that acceptance is not fixed. Candidates adjust their decisions to the allocation rule: a one-standard-deviation increase in expected future offer arrivals leads candidates to decline roughly one in eight offers they would otherwise have

accepted. Once this behavioral margin is taken seriously, the effect of a reform cannot be inferred from the rule alone, because the candidates a rule intends to help are not necessarily the ones who end up transplanted. I construct a model that incorporates these endogenous decisions, and it gives a clear answer. Acuity Circles was introduced to reduce geographic disparities and it did: welfare at listing rose by an amount equivalent to roughly doubling the supply of livers, and only 1.6% of candidates were worse off. Those candidates were the ones who had benefited from priority under the old local boundaries, so the cost of the reform fell exactly where a reform aimed at geographic disparity would place it.

For the policy now under consideration, a score weighting medical urgency against distance, the choice of weights is critical. A score based only on distance makes 59.1% of candidates worse off than under the old boundary system. A score based only on urgency raises average welfare the most but increases the share of livers that go unused by 8.5 percentage points. Given that 12,274 livers were recovered for transplant in 2024, this implies more than a thousand additional livers wasted each year. Neither extreme is optimal. I identify a range of intermediate weights that improves on Acuity Circles.

This trade-off exists because a liver deteriorates in transit. But that link is beginning to weaken. Preservation technology is changing how quickly a liver loses value as it travels, and machine perfusion is beginning to do exactly that: by 2023, 12.1% of livers recovered for transplant were machine-perfused (Wisel et al., 2024).[32] The allocation rule affects acceptance decisions by changing candidates' expectations about future offers; technology could shift their preferences themselves. If candidates become more willing to accept livers from farther away, the cost of weighting medical urgency could fall, and allocation could move toward the sickest candidates first without paying for it in wasted livers.

[32]The two liver perfusion devices received Food and Drug Administration (FDA) approval in late 2021.

References

Abadie, A. and S. Gay (2006). The impact of presumed consent legislation on cadaveric organ donation: a cross-country study. *Journal of health economics 25*(4), 599–620.

Abdulkadiroğlu, A., N. Agarwal, and P. A. Pathak (2017). The welfare effects of coordinated assignment: Evidence from the new york city high school match. *American Economic Review 107*(12), 3635–3689.

Agarwal, N., I. Ashlagi, M. A. Rees, P. Somaini, and D. Waldinger (2021). Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica 89*(1), 37–76.

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

Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica 70*(4), 1519–1543.

Arcidiacono, P., P. Bayer, J. R. Blevins, and P. B. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *The Review of Economic Studies 83*(3), 889–931.

Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*(6), 1823–1867.

Ashlagi, I. and A. E. Roth (2026). Organ allocation and transplantation. *Annual Review of Economics 18*.

Axelrod, D. A., P. A. Vagefi, and J. P. Roberts (2015). The evolution of organ allocation for liver transplantation: tackling geographic disparity through broader sharing. *Annals of surgery 262*(2), 224–227.

Bae, H. (2024). Can redrawing boundaries save lives? evidence from a reform of the kidney allocation system. Working paper.

Brown, S., Y. Savva, R. Barth, J. LaMattina, and P. J. Thuluvath (2020). Multiple regional listing increases liver transplant rates for those with model for end-stage liver disease score¡ 15. *Transplantation 104*(4), 762–769.

Callison, K. and A. Levin (2016). Donor registries, first-person consent legislation, and the supply of deceased organ donors. *Journal of health economics 49*, 70–75.

Choi, Y. (2020). The effect of broader liver sharing on patient outcomes and offer acceptance. *SSRN*.

Dickert-Conlin, S., T. Elder, and B. Moore (2011). Donorcycles: Motorcycle helmet laws and the supply of organ donors. *The Journal of Law and Economics 54*(4), 907–935.

Dickert-Conlin, S., T. Elder, and K. Teltser (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.

Ershoff, B. D., C. K. Lee, C. L. Wray, V. G. Agopian, G. Urban, P. Baldi, and M. Cannesson (2020). Training and validation of deep neural networks for the prediction of 90-day post-liver transplant mortality using unos registry data. In *Transplantation proceedings*, Volume 52, pp. 246–258. Elsevier.

Feng, S., N. P. Goodrich, J. L. Bragg-Gresham, D. M. Dykstra, J. D. Punch, M. DebRoy, S. M. Greenstein, and R. M. Merion (2006). Characteristics associated with liver graft failure: the concept of a donor risk index. *American journal of transplantation 6*(4), 783–790.

Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies 60*(3), 497–529.

Howard, D. H. (2002). Why do transplant surgeons turn down organs?: A model of the accept/reject decision. *Journal of health economics 21*(6), 957–969.

Kennes, J., D. Monte, and N. Tumennasan (2014). The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics 6*(4), 362–406.

Kessler, J. B. and A. E. Roth (2012). Organ allocation policy and the decision to donate. *American Economic Review 102*(5), 2018–2047.

Kwong, A. J., J. R. Lake, D. P. Schladt, A. Martinez, S. Weiss, D. Handarova, B. Schumacher, J. J. Snyder, A. Hart, and W. R. Kim (2026). Optn/srtr 2024 annual data report: Liver. *American Journal of Transplantation 26*(8), S214–S320.

Lai, J. C., S. Feng, and J. P. Roberts (2012). An examination of liver offers to candidates on the liver transplant wait-list. *Gastroenterology 143*(5), 1261–1265.

Lozanovski, V. J., B. Döhler, K. H. Weiss, A. Mehrabi, and C. Süsal (2020). The differential influence of cold ischemia time on outcome after liver transplantation for different indications—who is at risk? a collaborative transplant study report. *Frontiers in immunology 11*, 892.

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica 70*(2), 801–816.

Martin, P., A. DiMartini, S. Feng, R. Brown Jr, and M. Fallon (2014). Evaluation for liver transplantation in adults: 2013 practice guideline by the american association for the study of liver diseases and the american society of transplantation. *Hepatology 59*(3), 1144–1165.

Massie, A. B., B. Caffo, S. E. Gentry, E. C. Hall, D. Axelrod, K. L. Lentine, M. A. Schnitzler, A. Gheorghian, P. R. Salvalaggio, and D. L. Segev (2011). Meld exceptions and rates of waiting list outcomes. *American Journal of Transplantation 11*(11), 2362–2371.

Merion, R. M., D. E. Schaubel, D. M. Dykstra, R. B. Freeman, F. K. Port, and R. A.

Wolfe (2005). The survival benefit of liver transplantation. *American Journal of Transplantation 5*(2), 307–313.

Molinari, M., S. Ayloo, A. Tsung, D. Jorgensen, A. Tevar, S. H. Rahman, and N. Jonassaint (2019). Prediction of perioperative mortality of cadaveric liver transplant recipients during their evaluations. *Transplantation 103*(10), e297–e307.

Pan, E. T., D. Yoeli, N. T. N. Galvan, M. L. Kueht, R. T. Cotton, C. A. O'Mahony, J. A. Goss, and A. Rana (2018). Cold ischemia time is an important risk factor for post–liver transplant prolonged length of stay. *Liver Transplantation 24*(6), 762–768.

Rana, A., M. Hardy, K. Halazun, D. Woodland, L. Ratner, B. Samstein, J. Guarrera, R. Brown Jr, and J. Emond (2008). Survival outcomes following liver transplantation (soft) score: a novel method to predict patient survival following liver transplantation. *American Journal of Transplantation 8*(12), 2537–2546.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, 999–1033.

Singal, A. K. and P. S. Kamath (2013). Model for end-stage liver disease. *Journal of clinical and experimental hepatology 3*(1), 50–60.

Sweat, K. (2026). Endogenous priority in centralized matching markets: The design of the heart transplant waitlist. Working paper.

Waldinger, D. (2021). Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. *American Economic Review 111*(8), 2660–2696.

Weaver, T. and J. Pyke (2023, November). Background and methodology in simulation analysis: OASim. Technical report, Scientific Registry of Transplant Recipients.

Wey, A., J. Pyke, D. P. Schladt, S. E. Gentry, T. Weaver, N. Salkowski, B. L. Kasiske, A. K. Israni, and J. J. Snyder (2018). Offer acceptance practices and geographic variability in allocation model for end-stage liver disease at transplant. *Liver Transplantation 24*(4), 478–487.

Wisel, S. A., J. A. Steggerda, and I. K. Kim (2024). Use of machine perfusion in the united states increases organ utilization and improves dcd graft survival in liver transplantation. *Transplantation direct 10*(12), e1726.

Yeh, H., E. Smoot, D. A. Schoenfeld, and J. F. Markmann (2011). Geographic inequity in access to livers for transplantation. *Transplantation 91*(4), 479–486.

Zhang, J. (2010). The sound of silence: Observational learning in the us kidney market. *Marketing Science 29*(2), 315–335.

Table 1: Descriptive Statistics

	Mean		
	Overall	Before AC	After AC
Panel A: Offers (N=8,610,991)			
Medical urgency at offer			
a) Status 1A	0.66%	0.98%	0.52%
b) MELD 37-40	2.41%	3.51%	1.93%
c) MELD 33-36	2.94%	5.43%	1.85%
d) MELD 29-32	7.38%	12.34%	5.22%
e) MELD 15-28	63.69%	65.44%	62.93%
f) MELD 6-14	22.92%	12.30%	27.54%
Acceptance rate	36.35%	37.56%	35.82%
a) Status 1A	15.52%	20.57%	12.30%
b) MELD 37-40	20.97%	22.31%	20.13%
c) MELD 33-36	29.69%	30.15%	29.08%
d) MELD 29-32	37.90%	38.44%	37.34%
e) MELD 15-28	42.73%	44.82%	41.77%
f) MELD 6-14	24.42%	17.40%	25.91%
Panel B: Transplant Candidates (N=112,444)			
Wait time before first offer (in days)	11.23	14.86	6.60
a) Status 1A	1.70	1.97	1.41
b) MELD 37-40	2.01	2.81	1.40
c) MELD 33-36	2.95	4.42	1.82
d) MELD 29-32	8.62	10.53	6.85
e) MELD 15-28	8.32	10.38	6.18
f) MELD 6-14	21.59	28.88	14.16
Panel C: Donors (N=66,472)			
Number of offers per donor	129.54	86.56	165.24

Notes: The sample is drawn from the restricted Scientific Registry of Transplant Recipients (SRTR) datasets. Panel A reports offer-level outcomes. The sample includes responses of "yes", "no", or "provisional yes", and excludes bypassed offers, segmented transplants, and organs not recovered for transplantation. Panel B reports candidate-level outcomes, the sample includes candidates who were listed during the analysis period. Wait time is defined as the number of days between listing and the first offer received. Panel C reports donor-level outcomes. Each observation is a deceased donor whose liver was recovered for transplantation.

Table 2: Effect on Candidate Behavior: Offer Acceptance

	(1) Overall	(2)	(3)	(4)	(5)	(6)	(7)
		MELD tier					
		[6-14]	[15-28]	[29-32]	[33-36]	[37-40]	[1A]
Post × ΔAccess	−2.193***	−0.126	0.242	−6.024**	0.951	−0.326	–
	(0.563)	(0.818)	(1.224)	(2.691)	(1.406)	(1.025)	–
Mean of dep. var. (Pre, %)	37.56	17.40	44.82	38.44	30.15	22.31	20.57
St. dev. ΔAccess (pp)	2.16	2.16	2.16	2.16	2.16	2.16	–
Effect of 1 SD ΔAccess (pp)	-4.75	-0.27	0.52	-13.03	2.06	-0.71	–
Center FE	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
MELD-tier FE	✓						
Observations	8,610,991	1,973,609	5,484,274	635,632	252,799	207,627	57,050

Notes: The unit of observation is a deceased-donor liver offer; the dependent variable is an indicator equal to one if the offer is accepted and zero otherwise. Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2); coefficients are in probability units. All specifications also include ΔAccess in levels. Column (1) is the pooled specification; Columns (2)–(7) restrict the sample to the indicated MELD tier. For Status 1A the interaction is not identified because ΔAccess has insufficient within-cell variation. "Mean of dep. var. (Pre)" is the pre-policy acceptance rate in percent (April 2016–January 2020). "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value, in percentage points. Standard errors (in parentheses) are clustered at the MELD-tier × DSA cell level.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Effect on Candidate Behavior: Monthly New Registrations

	(1) Overall	(2)	(3)	(4)	(5)	(6)	(7)
		MELD tier					
		[6-14]	[15-28]	[29-32]	[33-36]	[37-40]	[1A]
Post × ΔAccess	7.287***	3.818	3.703	−0.900	5.056*	0.161	–
	(1.907)	(4.157)	(3.230)	(1.755)	(2.930)	(0.936)	–
Mean of dep. var. (Pre)	1.195	2.376	3.216	0.550	0.400	0.560	0.246
St. dev. ΔAccess (pp)	2.16	2.16	2.16	2.16	2.16	2.16	–
Effect of 1 SD ΔAccess	0.158	0.083	0.080	-0.019	0.109	0.003	–
Center FE	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
Month-of-year FE	✓	✓	✓	✓	✓	✓	✓
MELD-tier FE	✓						
Observations	84,258	14,043	14,043	14,043	14,043	14,043	14,043

Notes: The unit of observation is a transplant center–MELD group–month; the dependent variable is the monthly count of new waitlist registrations. Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2). All specifications also include ΔAccess in levels. Column (1) is the pooled specification; Columns (2)–(7) restrict the sample to the indicated MELD tier. In the Status 1A subsample, the interaction is omitted because ΔAccess has insufficient within-group variation. "Mean of dep. var. (Pre)" is the average monthly count over the pre-policy months (April 2016–January 2020). "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value and is in registrations per center–MELD group per month. Standard errors (in parentheses) are clustered at the MELD-tier × DSA cell level.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Effect on Candidate Behavior: Multi-listing

	(1) Overall	(2)	(3)	(4)	(5)	(6)	(7)
		MELD tier					
		[6-14]	[15-28]	[29-32]	[33-36]	[37-40]	[1A]
Post × ΔAccess	1.828	6.782	−0.706	−3.170	−2.645	0.901	–
	(1.924)	(5.706)	(2.408)	(2.540)	(2.960)	(0.685)	–
Mean of dep. var. (Pre)	0.858	1.680	2.713	0.345	0.148	0.164	0.100
St. dev. ΔAccess (pp)	2.16	2.16	2.16	2.16	2.16	2.16	–
Effect of 1 SD ΔAccess	0.040	0.147	-0.015	-0.069	-0.057	0.019	–
Center FE	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
MELD-tier FE	✓						
Observations	6,816	1,136	1,136	1,136	1,136	1,136	1,136

Notes: The unit of observation is a transplant center–MELD group–year; the dependent variable is the annual count of new multi-listing registrations, i.e., registrations by candidates already active at another center. Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2). All specifications also include ΔAccess in levels. Column (1) is the pooled specification; Columns (2)–(7) restrict the sample to the indicated MELD tier. In the Status 1A subsample, the interaction is omitted because ΔAccess has insufficient within-group variation. "Mean of dep. var. (Pre)" is the average annual count over the pre-policy period (April 2016–January 2020). "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value and is in registrations per center–MELD group per year. Standard errors (in parentheses) are clustered at the MELD-tier × DSA cell level.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Effect on Overall Mortality from Listing

	Overall death within horizon of listing			
	(1) 3 months	(2) 6 months	(3) 9 months	(4) 12 months
Post × ΔAccess	−0.420***	−0.455***	−0.476***	−0.540***
	(0.082)	(0.099)	(0.129)	(0.138)
Mean of dep. var. (Pre, %)	7.63	11.15	14.25	16.70
St. dev. ΔAccess (pp)	2.16	2.16	2.16	2.16
Effect of 1 SD ΔAccess (pp)	-0.909	-0.984	-1.031	-1.167
Center FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
MELD-tier FE	✓	✓	✓	✓
Observations	98,134	91,063	84,077	76,795

Notes. The unit of observation is a waitlist registration; each column is a linear probability model for the probability of death by the stated horizon. The dependent variable equals one if the candidate dies within the column horizon of listing through any pathway—while on the waitlist or after transplantation—and zero otherwise. Each column uses a balanced risk set: only registrations whose full horizon is observed within the sample period are retained, and registrations whose horizon window straddles the policy implementation date (February 4, 2020) are excluded. Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2); coefficients are in probability units. All specifications also include ΔAccess in levels and year, center, and MELD-tier fixed effects. "Mean of dep. var. (Pre)" is the average outcome over the pre-policy period (April 2016–January 2020), in percent. "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value, in percentage points. Standard errors, clustered at the MELD-tier × DSA cell, are in parentheses.
$^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table 6: Effect on Waitlist Mortality

	Waitlist death within horizon of listing			
	(1) 3 months	(2) 6 months	(3) 9 months	(4) 12 months
Post × ΔAccess	−0.391***	−0.428***	−0.420***	−0.492***
	(0.074)	(0.089)	(0.120)	(0.128)
Mean of dep. var. (Pre, %)	6.22	8.83	10.98	12.61
St. dev. ΔAccess (pp)	2.16	2.16	2.16	2.16
Effect of 1 SD ΔAccess (pp)	-0.846	-0.927	-0.909	-1.064
Center FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
MELD-tier FE	✓	✓	✓	✓
Observations	98,179	91,105	84,116	76,829

Notes. The unit of observation is a waitlist registration; each column is a linear probability model for the probability that waitlist death occurs within the stated horizon of listing. The dependent variable equals one if the candidate is removed from the waitlist owing to death or to becoming too sick to transplant within the column horizon of listing, and zero otherwise; transplantation and all other removals are treated as survival on the list. Each column uses a balanced risk set: only registrations whose full horizon is observable within the sample period are retained, and registrations whose horizon window straddles the policy implementation date (February 4, 2020) are excluded. Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2); coefficients are in probability units. All specifications also include ΔAccess in levels and year, center, and MELD-tier fixed effects. "Mean of dep. var. (Pre)" is the average outcome over the pre-policy period (April 2016–January 2020), in percent. "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value, in percentage points. Standard errors, clustered at the MELD-tier × DSA cell, are in parentheses.
$^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table 7: Effect on Post-Transplant Graft Failure

	Graft failure within horizon of transplant			
	(1) 3 months	(2) 6 months	(3) 9 months	(4) 12 months
Post × ΔAccess	0.224**	0.280**	0.310*	0.338**
	(0.109)	(0.141)	(0.160)	(0.169)
Mean of dep. var. (Pre, %)	5.27	6.72	7.93	8.87
St. dev. ΔAccess (pp)	2.16	2.16	2.16	2.16
Effect of 1 SD ΔAccess (pp)	0.484	0.605	0.670	0.731
Center FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
MELD-tier FE	✓	✓	✓	✓
Observations	60,733	58,960	55,834	51,966

Notes. The unit of observation is a deceased-donor liver transplant recipient; each column is a linear probability model for the probability that graft failure occurs within the stated horizon of transplant. The dependent variable equals one if the graft fails—through graft loss, re-transplantation, or patient death—within the column horizon of transplantation, and zero if the graft is known to be functioning at that horizon. Each column uses a recipient only if failure is observed within the horizon or follow-up extends through it; transplants censored before the horizon are excluded. Outcomes are measured from the date of transplant, not the date of listing. Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2); coefficients are in probability units. All specifications also include ΔAccess in levels and year, center, and MELD-tier fixed effects, where the MELD tier is measured at transplant. "Mean of dep. var. (Pre)" is the average outcome over the pre-policy period (April 2016–January 2020), in percent. "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value, in percentage points. Standard errors, clustered at the MELD-tier × DSA cell, are in parentheses.

$^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

Table 8: Effect on Accepted-Liver Quality

	(1) Donor age (years)	(2) DCD	(3) Cold ischemia (hours)
Post × ΔAccess	−18.368	−0.050	−5.238
	(14.853)	(0.275)	(5.676)
Mean of dep. var. (Pre)	39.30	7.10	5.99
St. dev. ΔAccess (pp)	2.16	2.16	2.16
Effect of 1 SD ΔAccess	-0.397	-0.108	-0.113
Center FE	✓	✓	✓
Year FE	✓	✓	✓
MELD-tier FE	✓	✓	✓
Observations	64,954	64,954	64,954

Notes. The unit of observation is a deceased-donor liver transplant. The dependent variables are characteristics of the accepted liver: donor age in years (Column 1); an indicator for donation after circulatory death, DCD (Column 2); and cold-ischemia time in hours (Column 3). Each column reports the coefficient on $\text{Post}_t \times \Delta\text{Access}_{gc(i)}$ from Equation (4.2); the DCD coefficient is in probability units. All specifications also include ΔAccess in levels and year, center, and MELD-tier fixed effects, where the MELD tier is measured at transplant. "Mean of dep. var. (Pre)" is the average outcome over the pre-policy period (April 2016–January 2020): years in Column 1, percent in Column 2, and hours in Column 3. "St. dev. ΔAccess" is the standard deviation of ΔAccess across candidates listed before the reform, used to scale all columns; "Effect of 1 SD" multiplies the coefficient by this value and is in the outcome's own units: years of donor age, percentage points of DCD, and hours of cold ischemia. Standard errors, clustered at the MELD-tier × DSA cell, are in parentheses.
$^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

Table 9: Priority classes under the two allocation rules

Class	Health status	Geographic tier
Panel A. Pre-reform (DSA-based) rule		
1	Status 1A	Donor's region
2	MELD 35–40	Donor's region
3	MELD 15–34	Donor's DSA
4	MELD 15–34	Donor's region, outside the DSA
5	Status 1A	Outside the region
6	MELD 15–40	Outside the region
7	MELD 6–14	Donor's DSA
8	MELD 6–14	Donor's region, outside the DSA
9	MELD 6–14	Outside the region
Panel B. Post-reform (Acuity Circles) rule		
Class	Health status	Geographic tier
10	Status 1A	Within 500
11–13	MELD 37–40	Within 150 / 150–250 / 250–500
14–16	MELD 33–36	Within 150 / 150–250 / 250–500
17–19	MELD 29–32	Within 150 / 150–250 / 250–500
20–22	MELD 15–28	Within 150 / 150–250 / 250–500
23	Status 1A	Beyond 500
24	MELD 30–40	Beyond 500
25	MELD 15–29	Beyond 500
26–28	MELD 6–14	Within 150 / 150–250 / 250–500
29	MELD 6–14	Beyond 500

Notes: A priority class is the intersection of a health-status tier and a geographic tier. Under the pre-reform rule, geography is defined by administrative boundaries (the donor's DSA and OPTN region); under the post-reform rule, by concentric circles of 150, 250, and 500 nautical miles around the donor hospital.

Table 10: Conditional Choice Probability of Declining an Offer

	Pre-policy (DSA) (1)	Post-policy change (2)
Liver quality q	-1.373 (0.008)	0.304 (0.009)
Liver quality q^2	-0.055 (0.005)	-0.088 (0.006)
Log distance $\ln(1+\ell)$	0.615 (0.015)	-0.170 (0.016)
$\times$ MELD 15–20	-0.151 (0.016)	0.322 (0.018)
$\times$ MELD 21–27	-0.214 (0.017)	0.270 (0.018)
$\times$ MELD 28–34	-0.303 (0.017)	0.147 (0.020)
$\times$ MELD 35+	-0.615 (0.024)	0.015 (0.033)
Log waiting time $\ln(1+t)$	-0.008 (0.009)	0.019 (0.011)
$\ln(1+t)^2$	-0.014 (0.009)	0.051 (0.010)
$\ln(1+t)^3$	0.007 (0.003)	0.008 (0.004)
MELD $\times$ q	0.232 (0.006)	-0.030 (0.007)
MELD $\times$ t	0.020 (0.005)	-0.020 (0.006)
Post-policy indicator		-0.249 (0.069)
MELD fixed effects (35)	Yes	
DSA fixed effects (52)	Yes	
$t \times q$ interactions	Yes	
Number of offers	1,737,877	

Notes: Logit estimates; the outcome is 1 if the offer is declined, so positive coefficients indicate a higher decline probability. Continuous variables t, ℓ, q are standardized. MELD enters interactions as a standardized score (MELD z) and as level fixed effects; Column (2) reports the coefficient on the covariate interacted with the post-reform indicator, so the post-period effect is the sum of columns (1) and (2). Maximum-likelihood standard errors in parentheses. All covariates are fully interacted with the post-period indicator; one collinear DSA dummy is dropped.

Table 11: Second-Stage Estimates of the Transplant Value Γ

	$\hat{\eta}$ (SE)
Constant	-6.6529 (0.1430)
MELD score	0.3050 (0.0039)
(MELD score)^2	-0.0037 (0.00005)
Liver quality	1.6476 (0.3522)
(Liver quality)^2	9.3806 (0.2344)
$\ln(1+\text{Distance})$	-0.3907 (0.0074)
$(\ln(1+\text{Distance}))^2$	-0.0050 (0.0004)
(MELD score) $\times$ (Liver quality)	-0.2300 (0.0038)
(MELD score) $\times$ $\ln(1+\text{Distance})$	0.0071 (0.0002)
(Liver quality) $\times$ $\ln(1+\text{Distance})$	-0.0299 (0.0069)
Number of offers	1,737,877

Notes: OLS estimates of the $\Gamma(x, q, \ell)$ on the values recovered from Equation (5.7). Bootstrap standard errors in parentheses (200 replications).

Table 12: Characteristics of losing versus gaining DSAs

	Loser DSAs (5)	All Gainers (42)	Neighbor Gainers (11)
Local donors	628.0	606.4	639.2
Mean MELD (waitlist)	20.17	19.39	18.85
Waitlist mortality	0.183	0.172	0.161
Supply/demand ratio	0.948	0.730	0.743
Median wait (days)	93.5	103.2	121.8
Share of candidates losing	48.8%	0.4%	1.4%

Table 13: Counterfactuals Outcomes

	Details	Batch size	Waitlist				Tx recipients		
			Welfare gain[a] (arrival-equiv.)	Δ Nonuse rate	Share of Losers	Years on list[b]	MELD	Distance (NM)	MELD 6–14 (%)
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Observed rules									
DSA	—	87	—	—	—	0.94	23.7	144	0.0
Acuity Circles	—	165	0.649	-6.3	23.1	0.92	23.8	152	0.0
Panel B: Acuity Circles levers									
Circles		87	0.332	0.0	34.2	0.97	24.4	126	0.0
Narrow bands		87	0.051	0.2	27.2	0.95	23.9	151	0.0
Large batches		165	0.231	-6.3	0.0	0.90	23.3	182	0.0
Circles + Narrow bands		87	0.448	0.8	29.5	0.98	24.6	139	0.0
Narrow bands + Large batches		165	0.277	-6.0	0.0	0.90	23.4	186	0.0
Circles + Large batches		165	0.539	-6.6	25.1	0.92	23.6	141	0.0
Panel C: Continuous distribution									
(w_u, w_d)	(0.0, 1.0)	224	0.521	-9.3	59.1	0.42	16.3	71	46.2
	(0.2, 0.8)	224	0.177	-10.0	48.7	0.73	21.1	69	19.2
	(0.4, 0.6)	224	0.310	-10.5	36.2	0.79	22.1	84	13.4
	(0.6, 0.4)	224	0.729	-9.9	19.0	0.87	23.3	139	6.6
	(0.8, 0.2)	224	1.570	-4.8	11.9	0.98	24.8	325	0.4
	(1.0, 0.0)	224	1.855	8.5	13.1	1.10	26.9	828	0.0

[a] Welfare gain at listing, expressed in arrival-equivalent units, $\frac{V^{cf} - V^{DSA}}{V^{DSA}}$

[b] Average time from listing to exit, over all exit reasons (transplant, and departure without transplantation).

Panel C rows are indexed by the score weights (w_u, w_d) in equation (5.8). Average MELD and average distance are computed over transplant recipients.

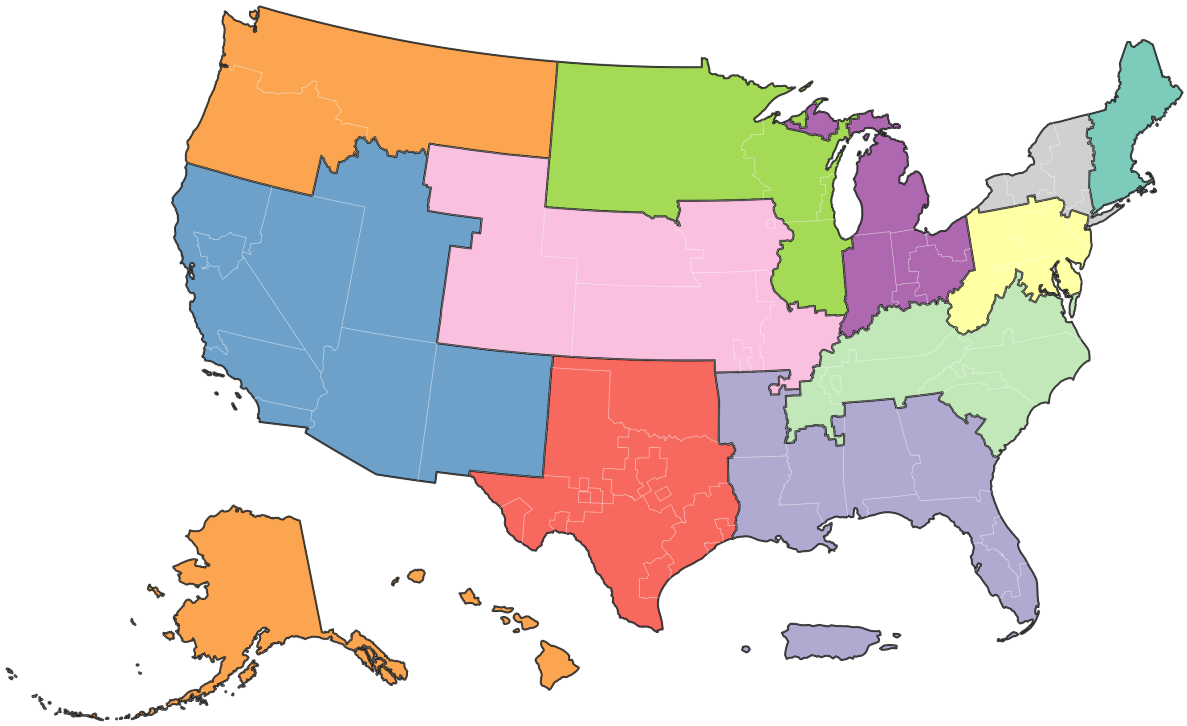


Figure 1: Map of OPTN Regions

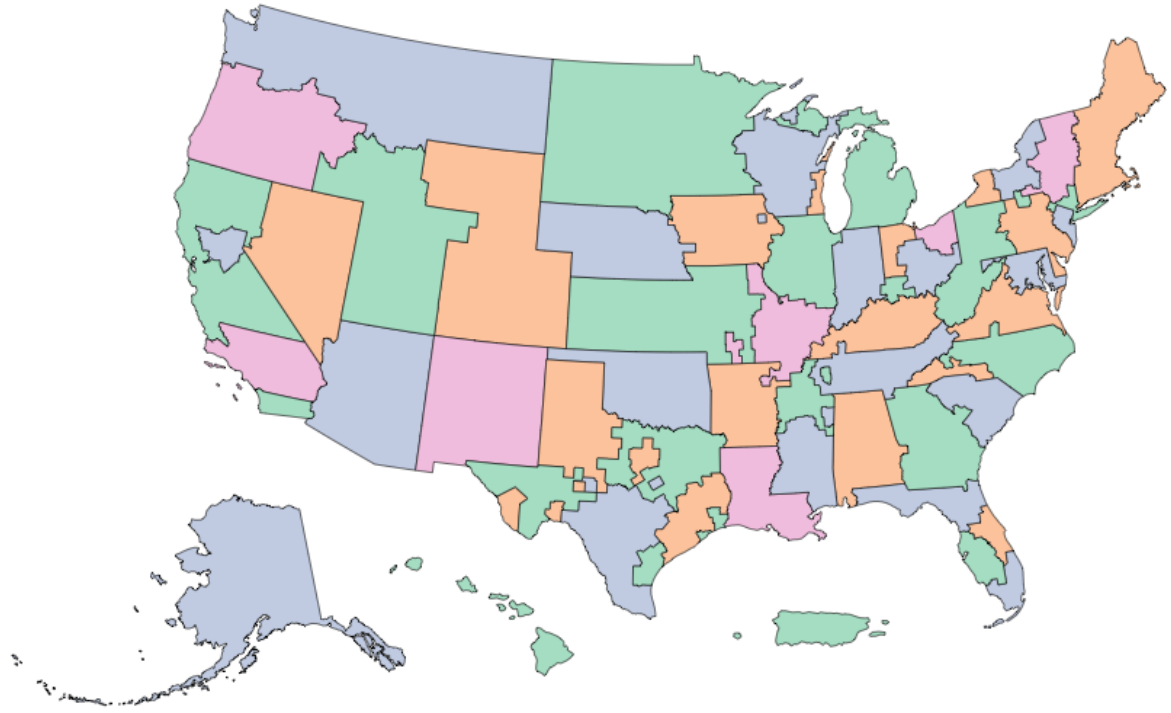
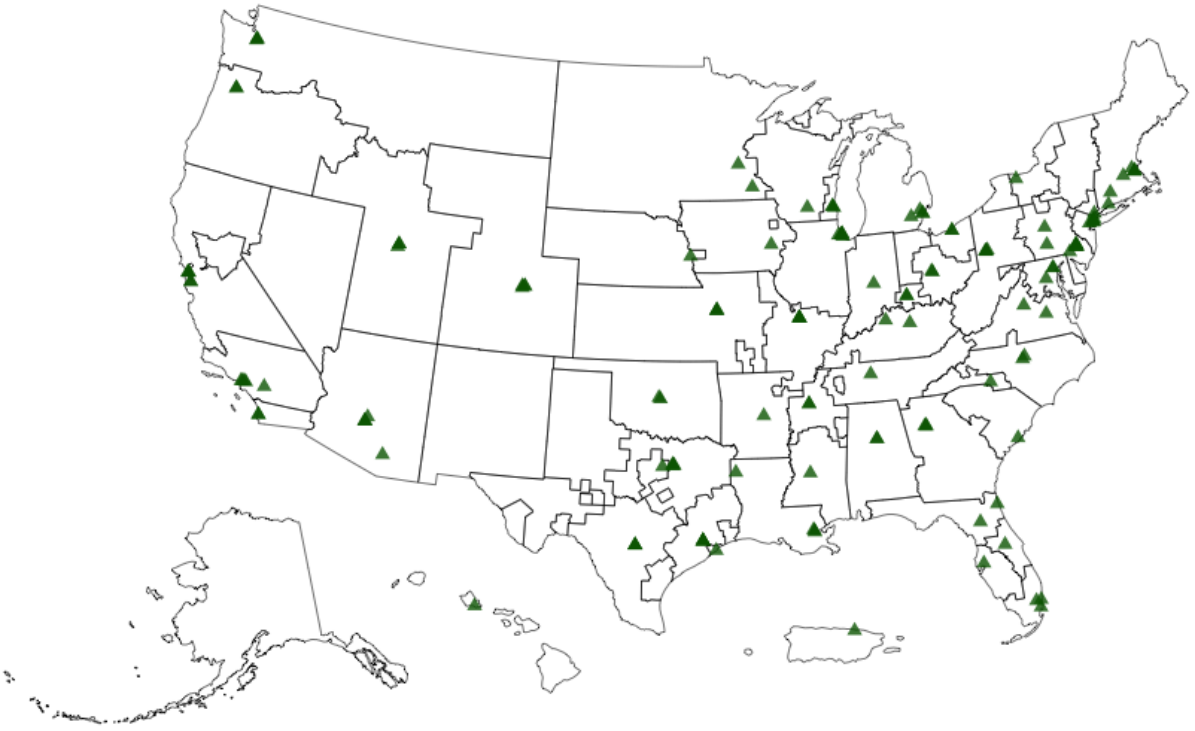
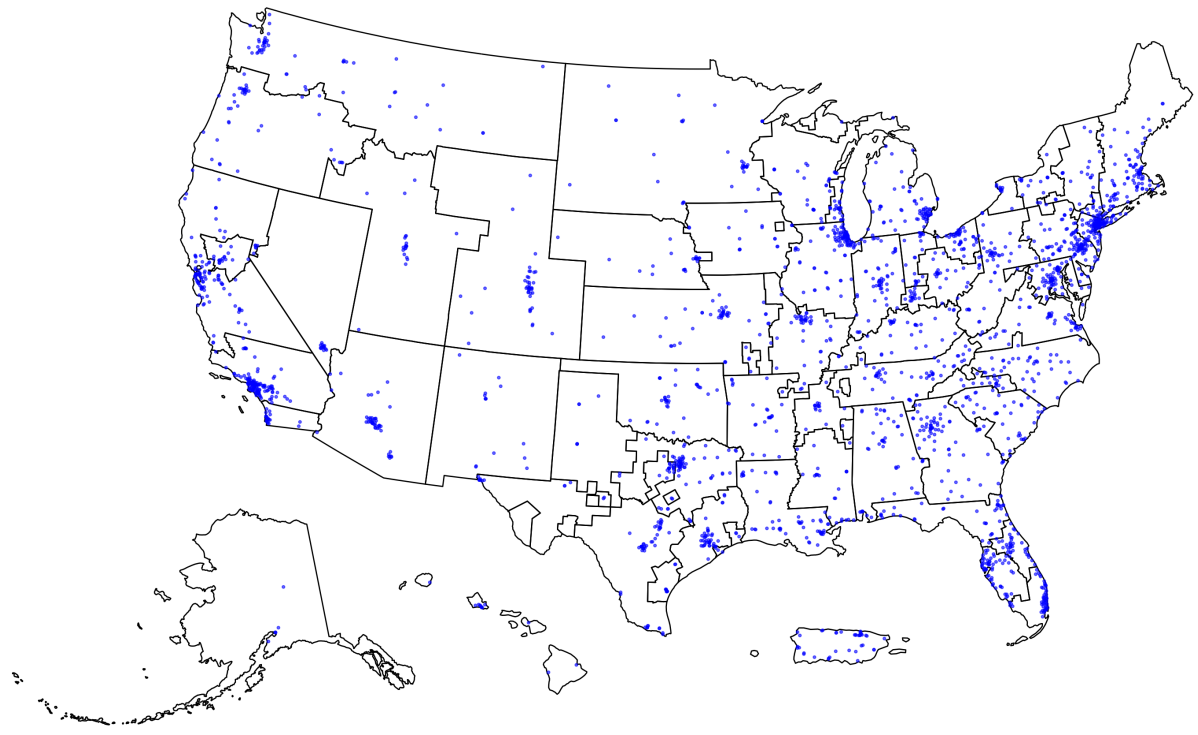


Figure 2: Map of Donation Service Areas

Figure 3: Map of Transplant Centers and Donor Hospitals

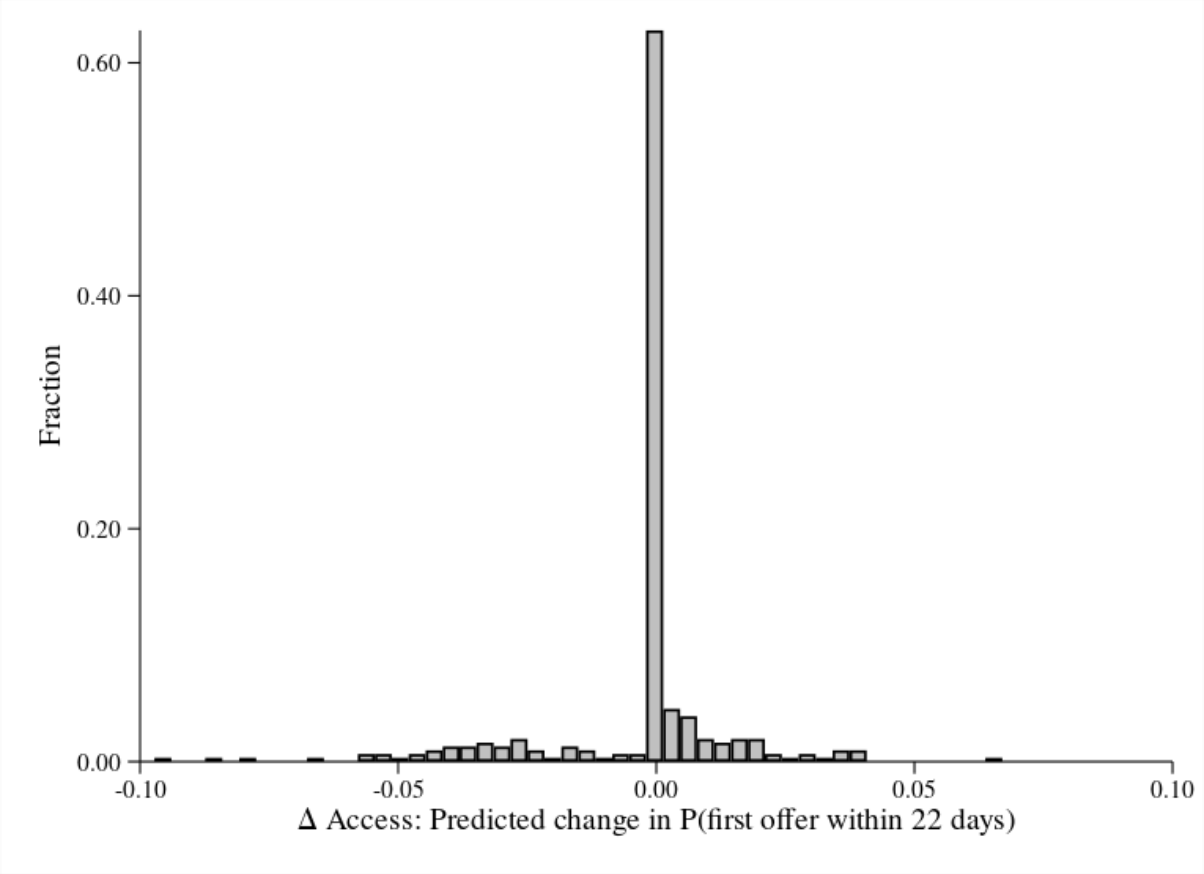


(a) Liver Transplant Centers



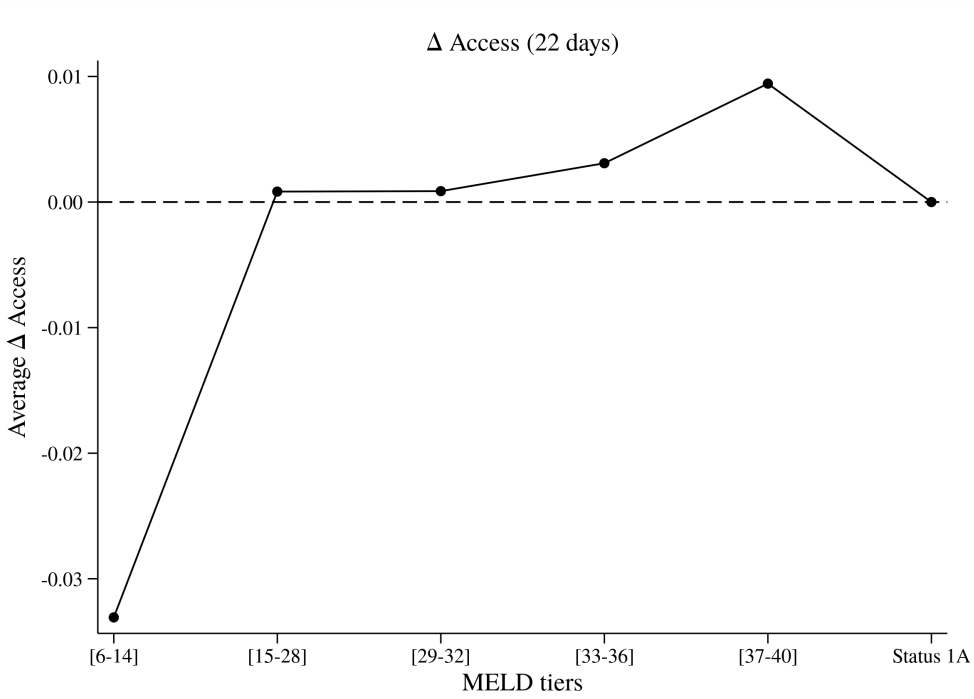
(b) Donor Hospitals

Figure 4: Distribution of Predicted Changes in Access to Transplant



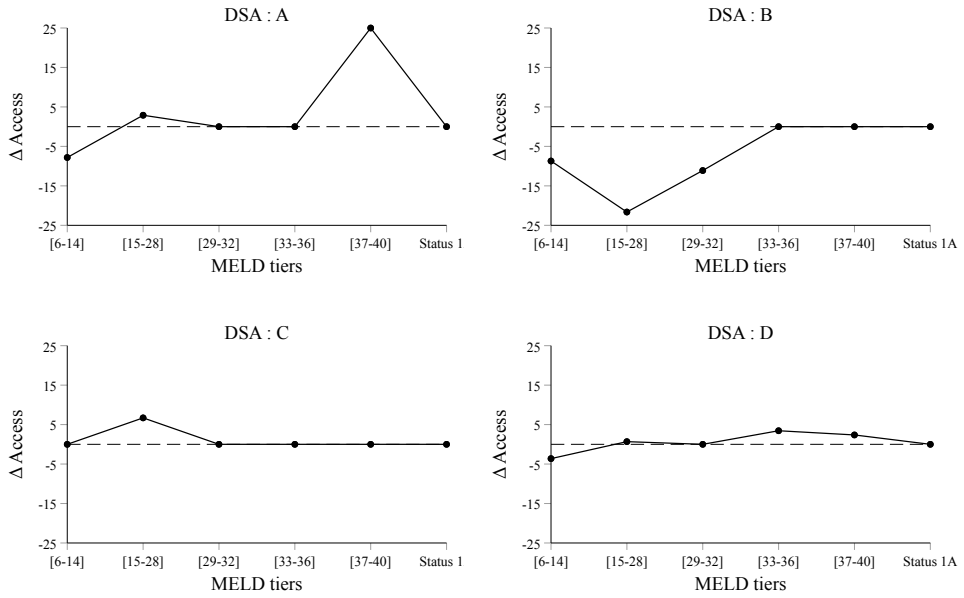
Notes: This figure displays a histogram of ΔAccess, defined in Equation (4.1), across the MELD-tier-by-DSA cells. The display is limited to $[-0.10, 0.10]$; the three cells (1.0%) with values outside this range are excluded. Probabilities are computed by simulating matchruns under both policies for each donor. The histogram uses 50 bins.
Source: Author calculations using Scientific Registry of Transplant Recipients (SRTR) datasets

Figure 5: Predicted Change in Access to Transplant



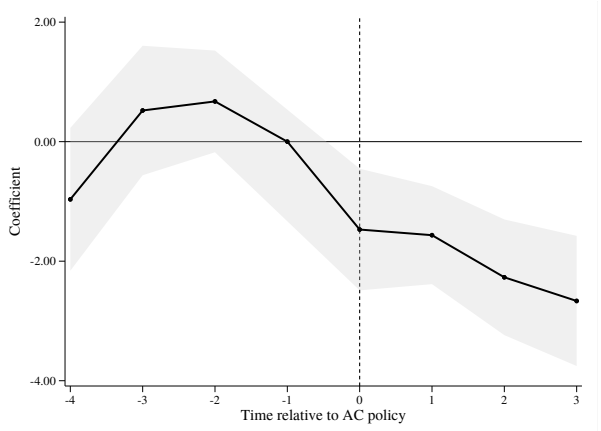
Notes: Each point shows the mean of ΔAccess–the predicted change in the probability of receiving a first offer within 22 days–within the indicated MELD tier.

(a) Average Predicted Changes in Access (22 days) by MELD Tier

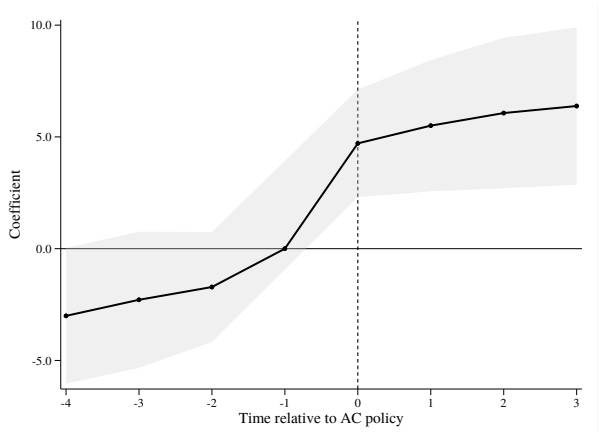


(b) Predicted Changes in Access (22 days) by MELD Tier–Selected DSAs

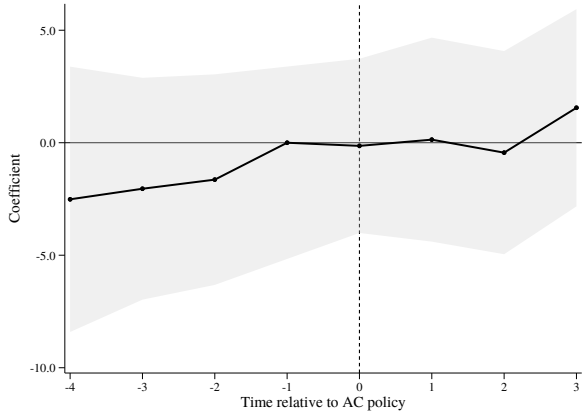
Figure 6: Event Study Figures on Candidate Behavior Outcomes



(a) Acceptance Rates



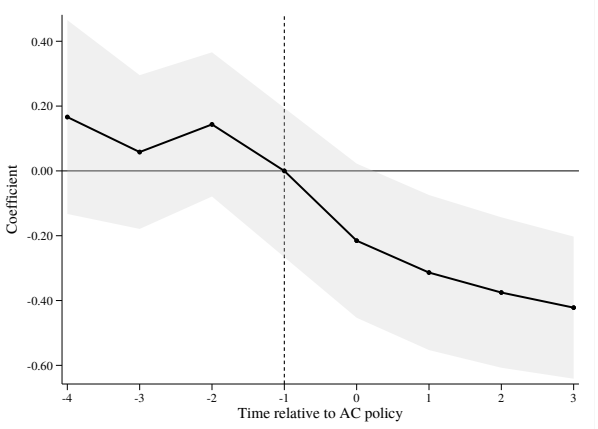
(b) Monthly New Registrations



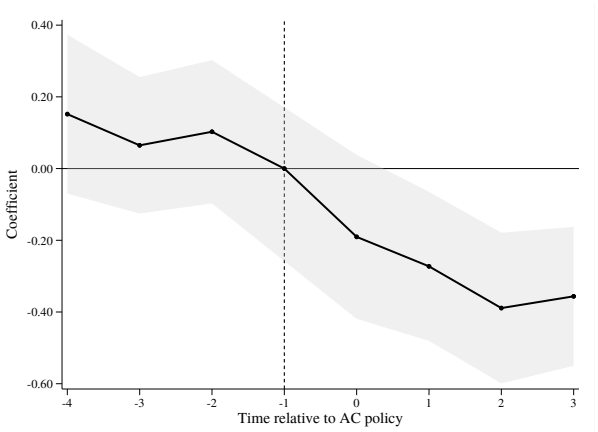
(c) Multi-listing

Notes: Figure 6.(a)-(c) display the coefficients and 95% confidence intervals from regression estimation in Equation (4.3)

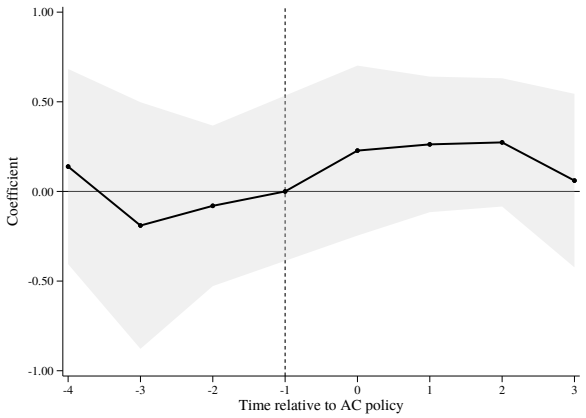
Figure 7: Event Study Figures on Mortality and Post-Transplant Outcomes



(a) Overall Mortality within 3 Months of Listing



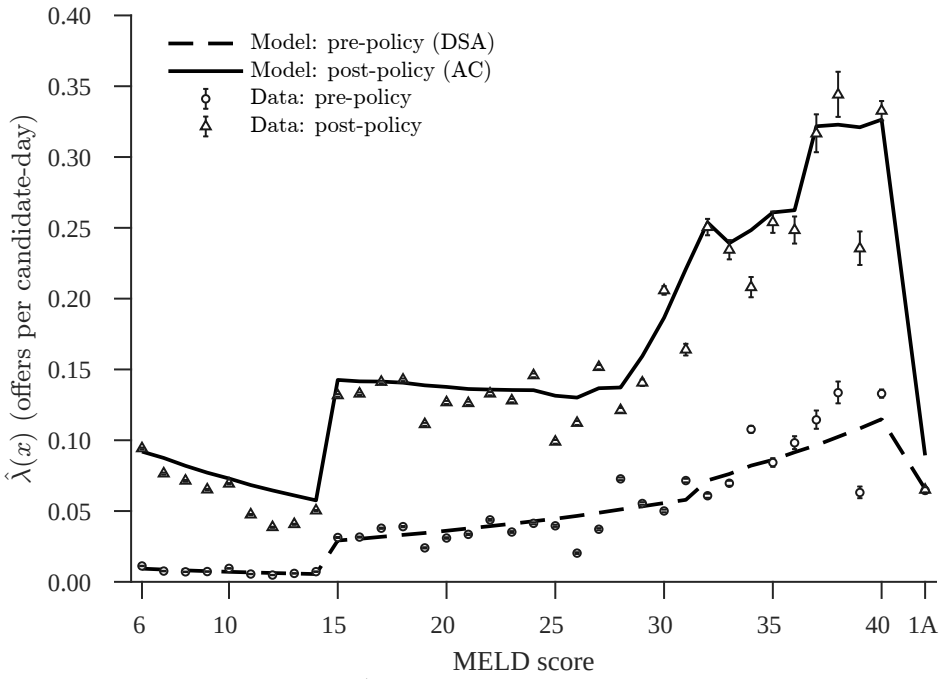
(b) Waitlist Mortality within 3 Months of Listing



(c) Graft Failure within 3 Months of Transplant

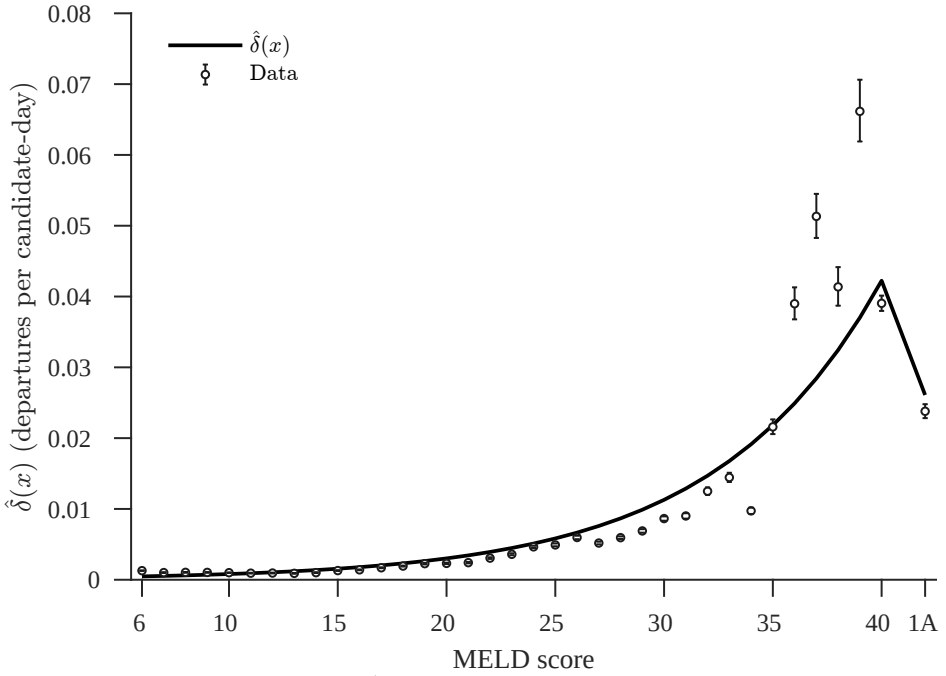
Notes: Figure 6.(a)-(c) display the coefficients and 95% confidence intervals from regression estimation in Equation (4.3)

Figure 8: Offer-arrival rates by MELD score



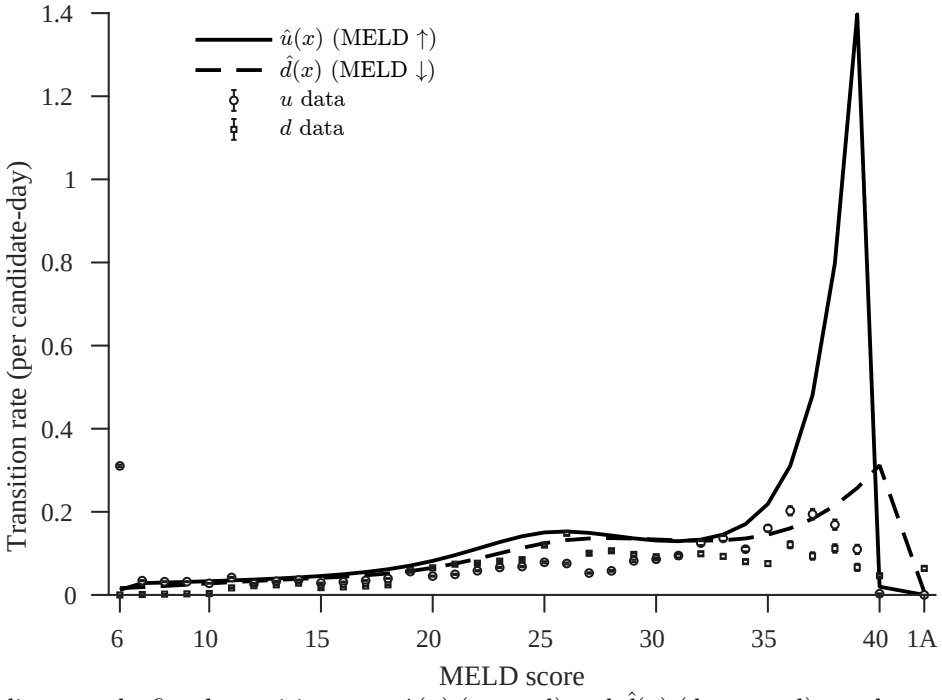
Notes: Lines are fitted arrival rates $\sum_c \hat{\lambda}_c(x)$ aggregated across priority classes and DSAs; markers are observed offer counts at each MELD score divided by the number of candidate-days spent at that score with 95% Poisson confidence intervals. Because MELD is observed only at offer times, the time a candidate spends at each MELD score is latent, so I compute expected candidate-days from the estimated process. "1A" denotes Status 1A.

Figure 9: Departure hazard $\hat{\delta}(x)$ by MELD score



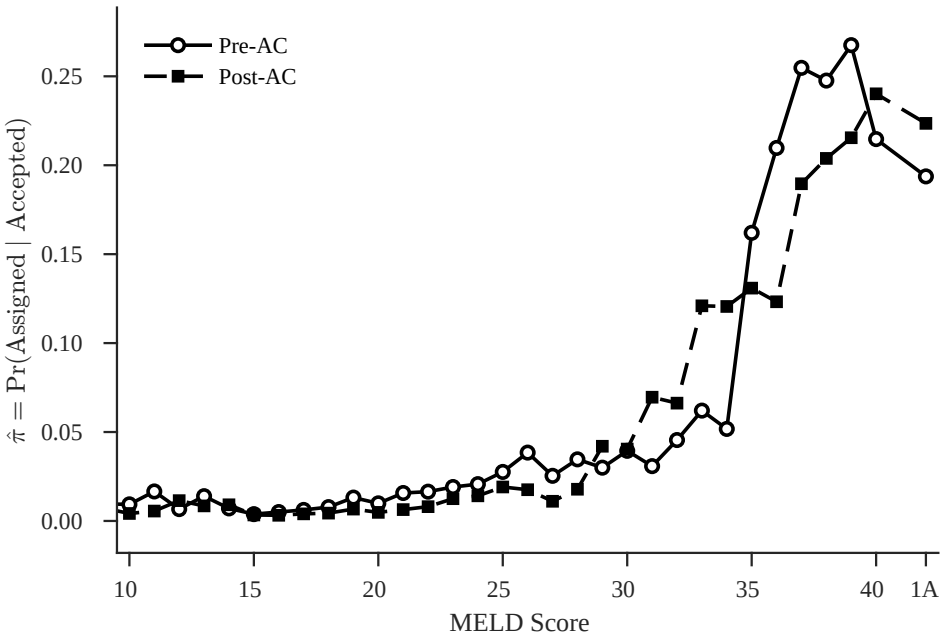
Notes: The line is the fitted departure rate $\hat{\delta}(x)$; markers are empirical departure rates per candidate-day with 95% confidence intervals. The MELD score is recorded at every offer but not at departure, so the empirical rates are binned by the score as of the most recent preceding offer.

Figure 10: MELD Transition rate $\hat{F}_{xx'}$



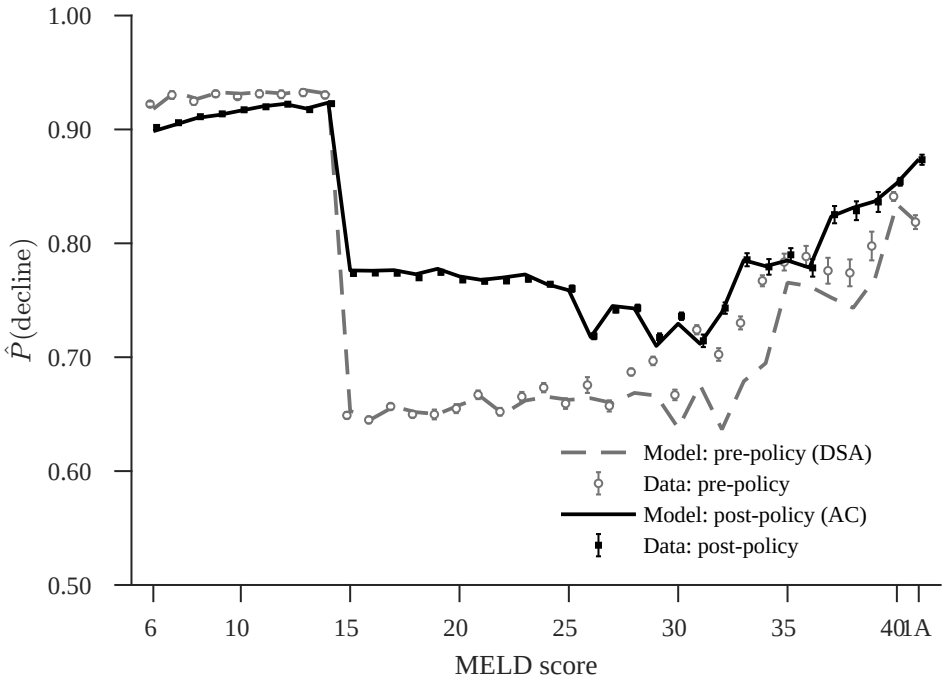
Notes: The lines are the fitted transition rates $\hat{u}(x)$ (upward) and $\hat{d}(x)$ (downward); markers are empirical transition rates per candidate-day with 95% confidence intervals, computed from consecutive offers no more than seven days apart. Because MELD changes between offers are not observed, empirical rates are binned by the MELD score at the most recent preceding offer; the seven-day window keeps that score close to the one held during the interval.

Figure 11: Assignment risk $\hat{\pi}$ by MELD score



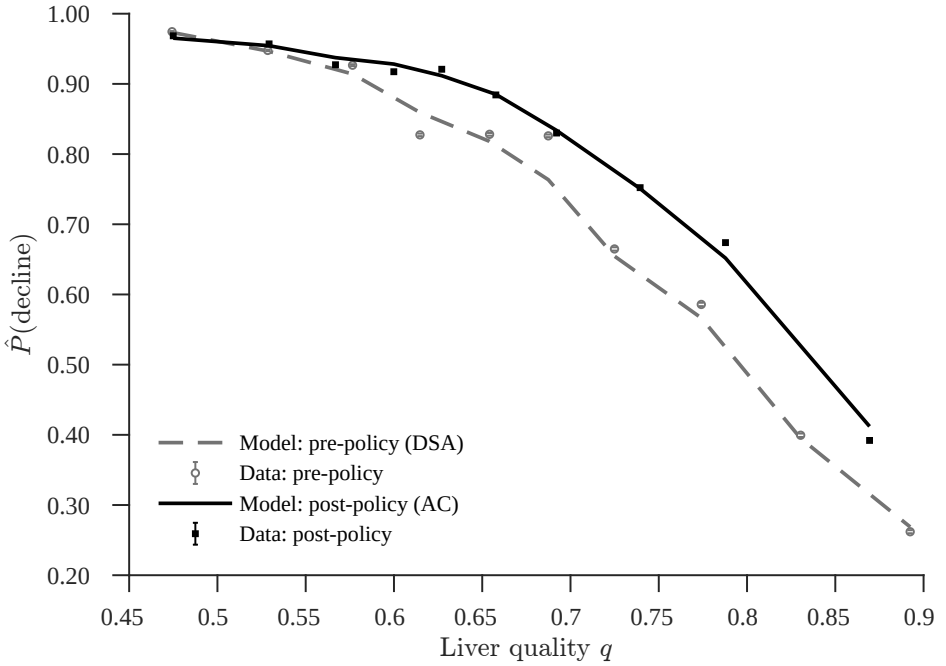
Notes: Assignment risk by MELD score and regime, aggregated over priority class, waiting time, quality, and DSA cells weighted by accepted offers.

Figure 12: CCP by MELD score



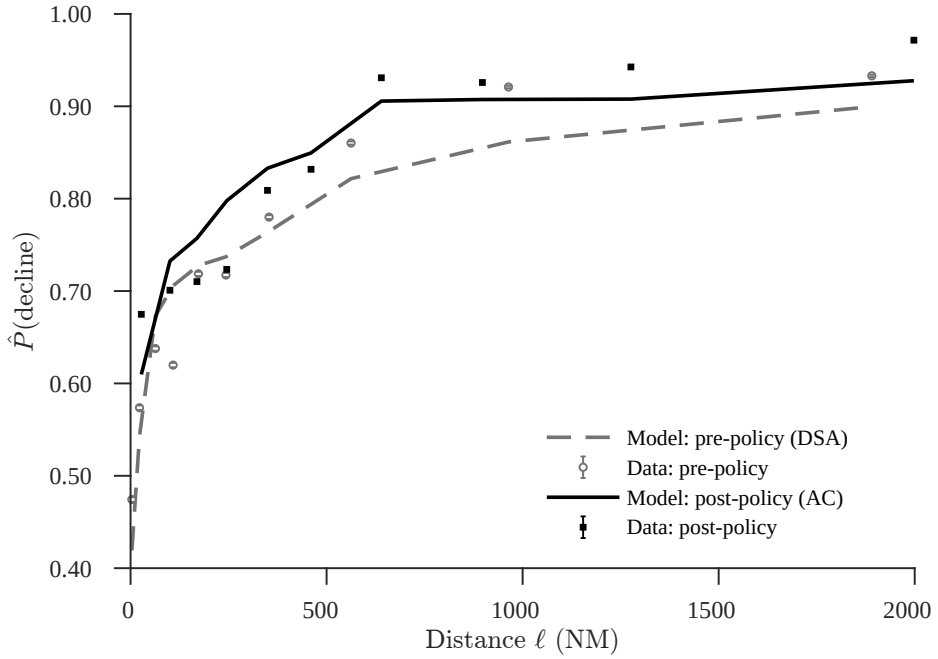
Notes: Lines are average CCPs (predicted probabilities averaged over the actual offers at each MELD level); markers are empirical decline rates with ±1 binomial standard error.

Figure 13: CCP by liver quality



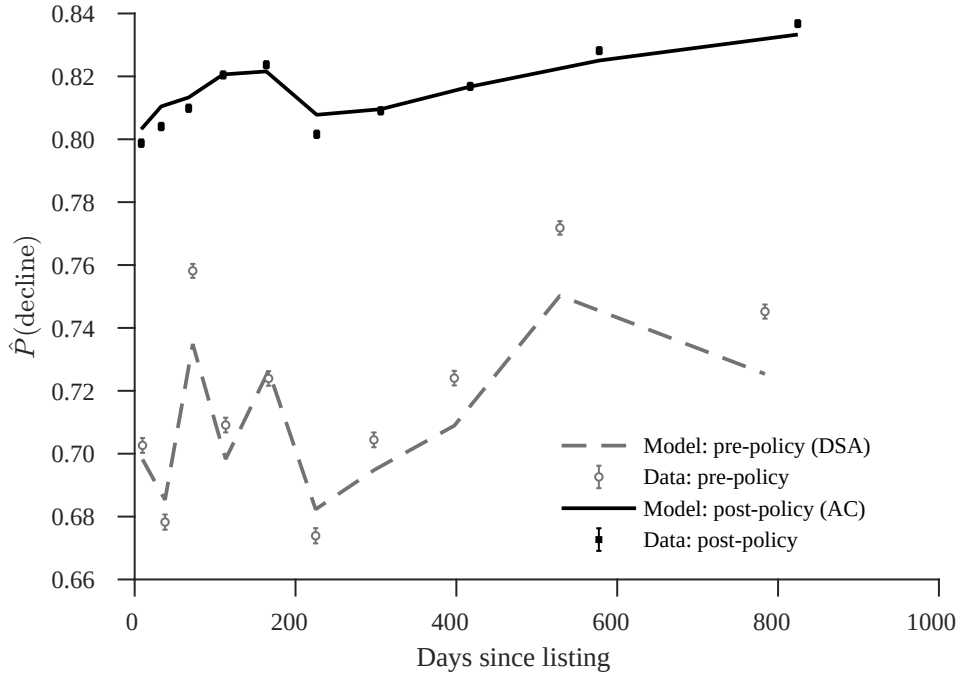
Notes: Lines are averages of estimated CCPs; markers are empirical decline rates with ±1 binomial standard error.

Figure 14: CCP by distance



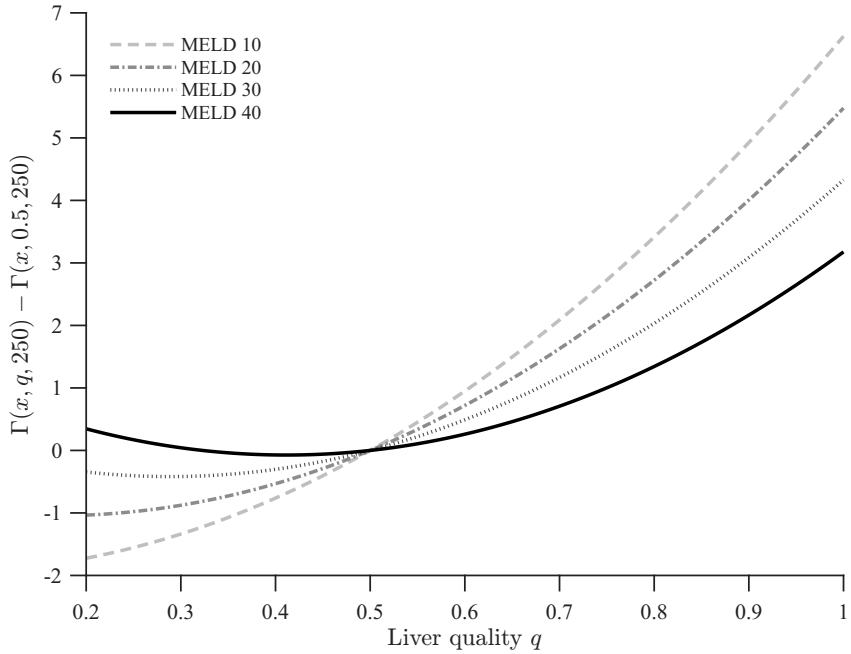
Notes: CCPs by distance between the donor hospital and the candidate's transplant center (nautical miles). Lines are estimated CCPs; markers are empirical decline rates with ±1 binomial standard error.

Figure 15: CCP by days since listing

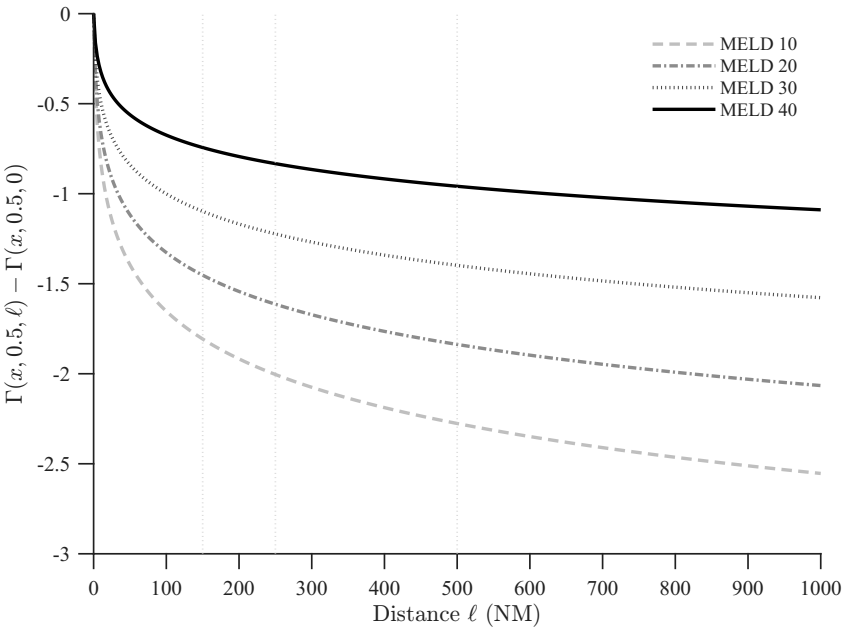


Notes: Lines are estimated CCPs; markers are empirical decline rates with ±1 binomial standard error.

Figure 16: Estimated Transplant Value $\Gamma(x, q, \ell)$



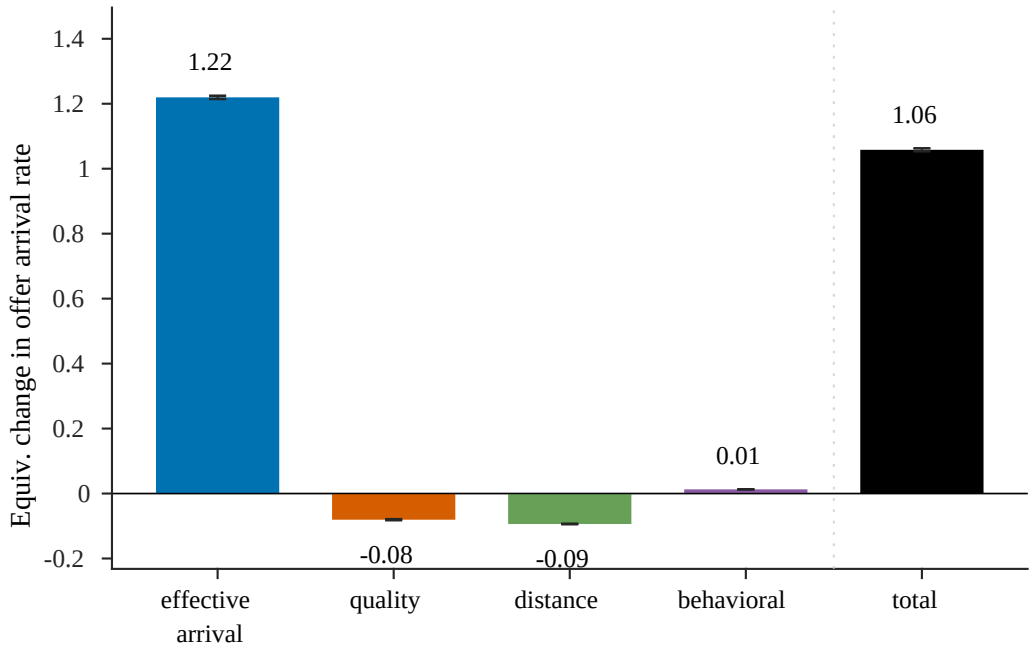
(a) Value of liver quality



(b) Penalty from distance

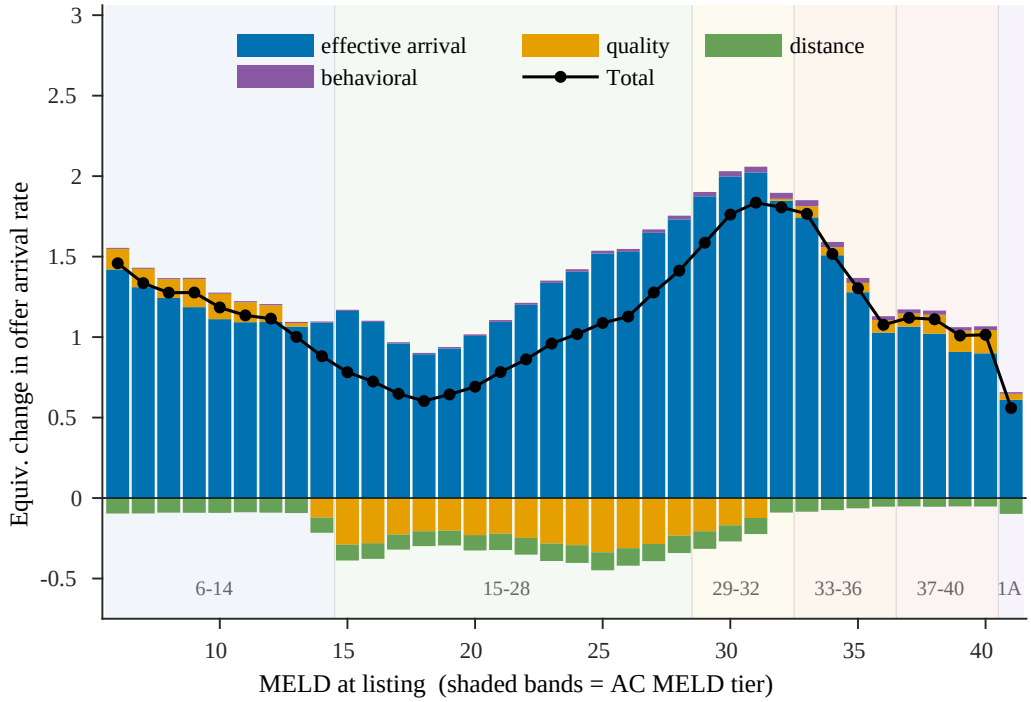
Notes: Computed from the estimates in Table 11. Because the level of the transplant value has no natural interpretation, each panel plots differences: Panel (a) shows the value of quality relative to $q = 0.5$, holding distance at 250 nautical miles; Panel (b) shows the penalty from distance relative to a zero-distance offer, holding quality at $q = 0.5$. The quality axis is restricted to the observed support of q among offers.

Figure 17: Decomposition of the Welfare Gain



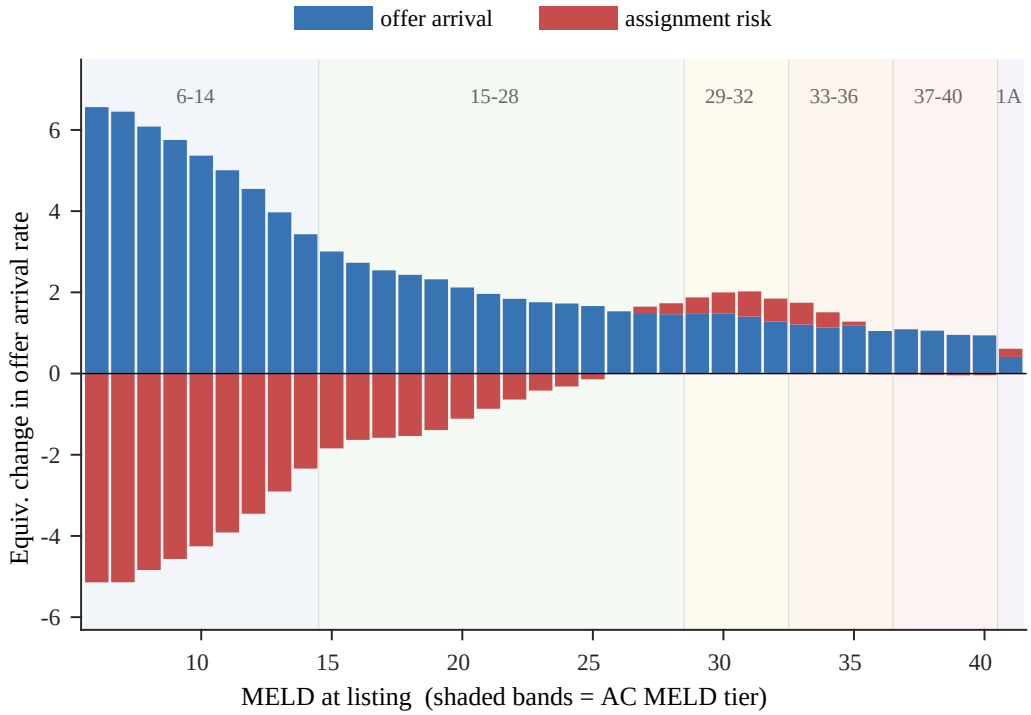
Notes: Each bar reports the contribution of one channel to the average welfare gain at listing, in arrival-equivalent units. The effective-arrival channel combines the change in offer-arrival rates and the change in the probability that an accepted offer results in transplantation. The channels sum to the total by construction. Error bars show 95% intervals based on variation across candidates.

Figure 18: Welfare Decomposition by MELD



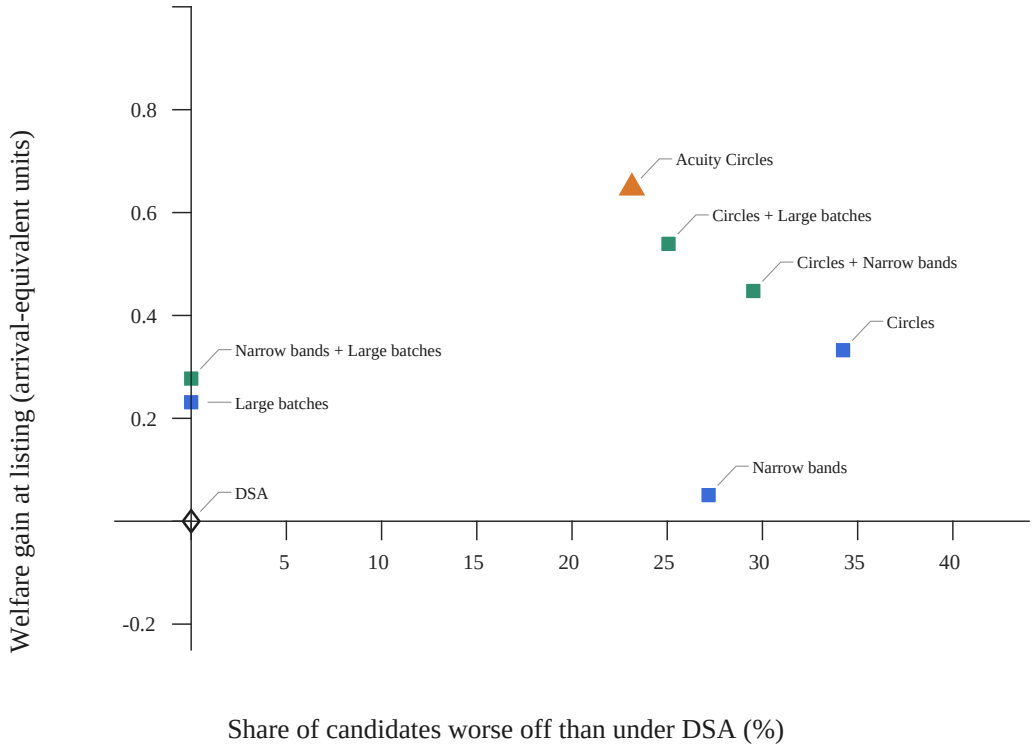
Notes: Bars report the average contribution of each channel to the welfare gain at listing, by MELD score at listing, in arrival-equivalent units. The black line is the total welfare gain, which equals the sum of the channels at each MELD score. Shaded bands mark the MELD tiers used under the Acuity Circles policy.

Figure 19: Effective-Arrival Channel Decomposition by MELD



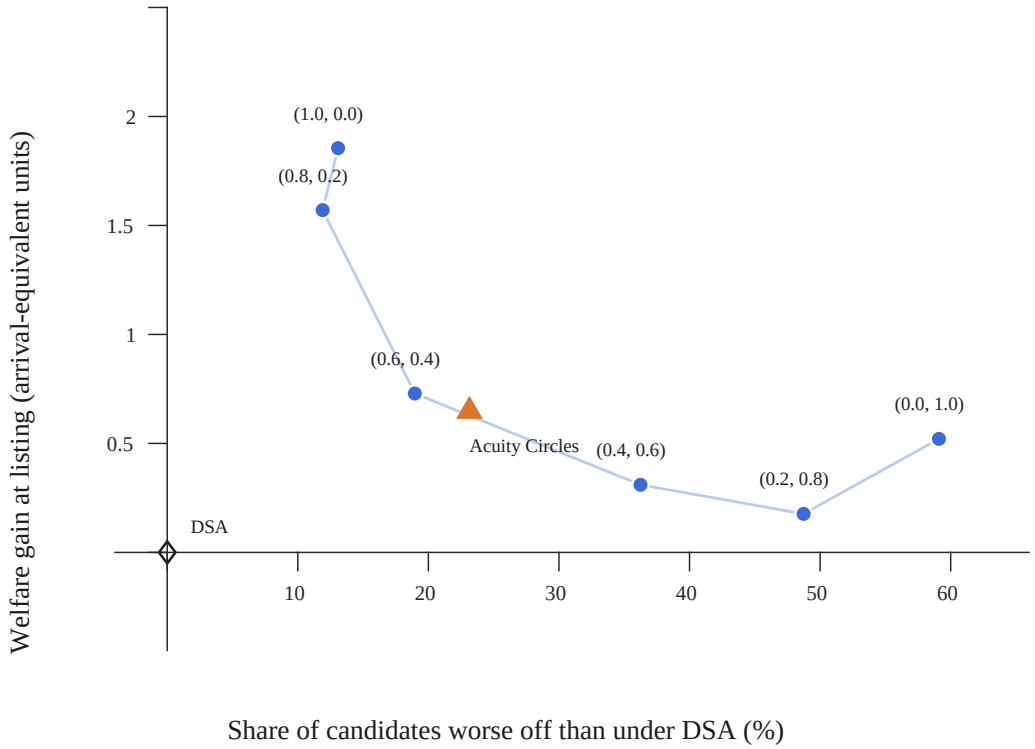
Notes: Bars decompose the effective-arrival channel into its offer-arrival and assignment components, by MELD score at listing, in arrival-equivalent units. The two components sum to the effective-arrival channel in Figure 18. Shaded bands mark the MELD tiers used under the Acuity Circles policy.

Figure 20: Welfare and Share of Losers under Partial Acuity Circles Reforms



Notes: Each square is a counterfactual in which a subset of the three Acuity Circles levers is turned on and the DSA system is retained otherwise. Blue squares turn on a single lever and green squares switch on two.

Figure 21: Welfare and Share of Losers under Continuous Distribution

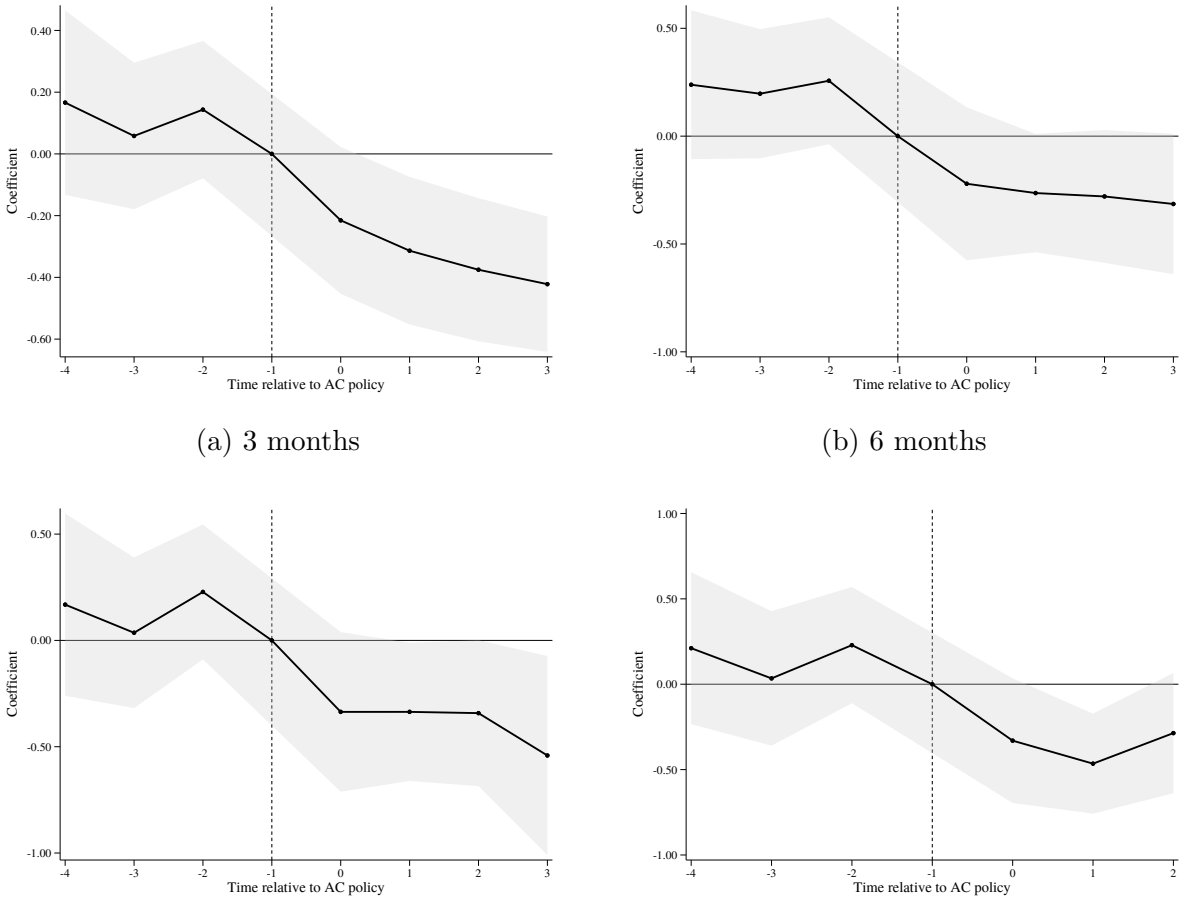


Notes: Each blue marker is a weight pair (w_u, w_d).

A. Additional Figures

This section reports event study estimates of Equation (4.3) for the mortality and post-transplant outcomes of Section 4.2.2 at all four horizons, complementing the three-month estimates in Figure 7a, as well as event studies for the accepted liver quality margins in Table 8.

Figure A1: Event Studies: Overall Mortality from Listing, by Horizon



(a) 3 months

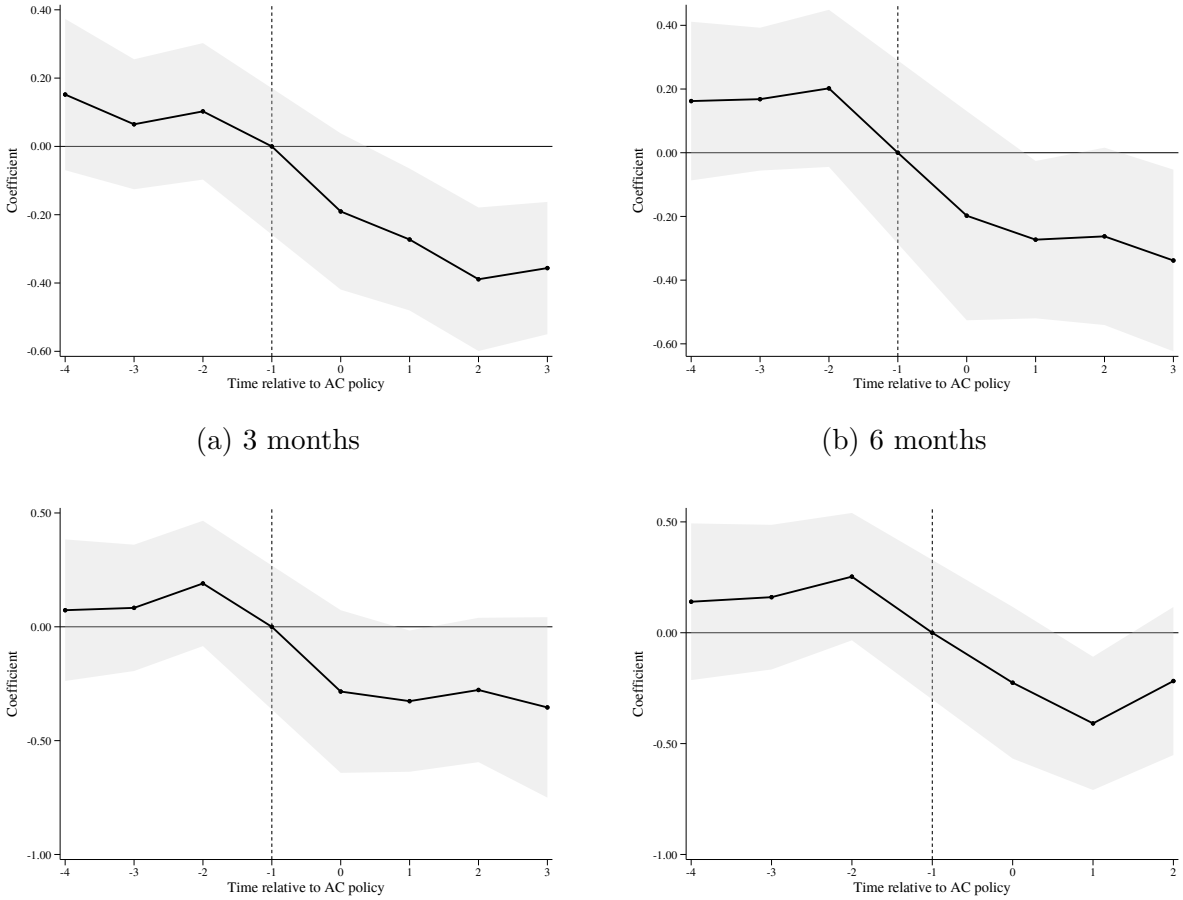
(b) 6 months

(c) 9 months

(d) 12 months

Notes: Each panel displays coefficients and 95% confidence intervals from Equation (4.3), where the dependent variable equals one if the candidate dies through any pathway within the stated horizon of listing. In Panel (d), the estimate for event year +3 is not identified because 2023 listings lack full twelve-month follow-up, so the window ends at +2. Standard errors are clustered at the MELD-tier × DSA cell level.

Figure A2: Event Studies: Waitlist Mortality, by Horizon



(a) 3 months

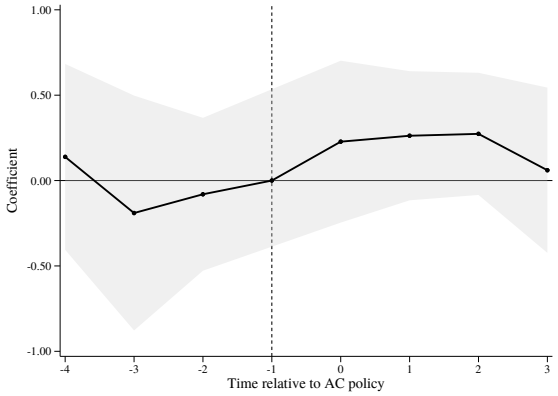
(b) 6 months

(c) 9 months

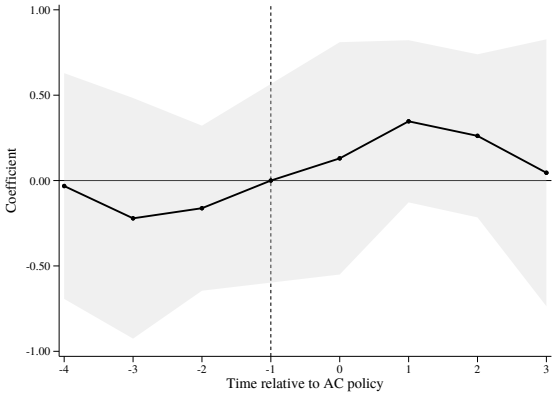
(d) 12 months

Notes: Each panel displays coefficients and 95% confidence intervals from Equation (4.3), where the dependent variable equals one if the candidate is removed from the waitlist owing to death or deteriorating health within the stated horizon of listing. In Panel (d), the window ends at event year +2 because 2023 listings lack full twelve-month follow-up. Standard errors are clustered at the MELD-tier × DSA cell level.

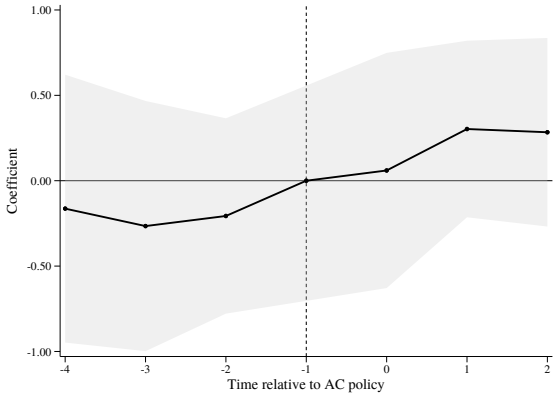
Figure A3: Event Studies: Post-Transplant Graft Failure, by Horizon



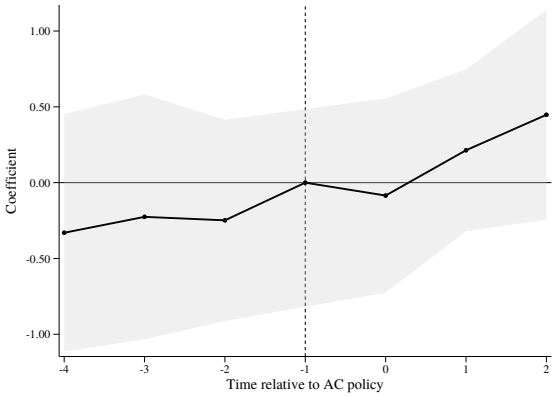
(a) 3 months



(b) 6 months



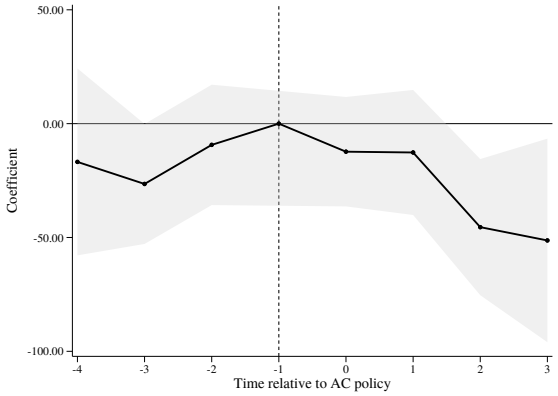
(c) 9 months



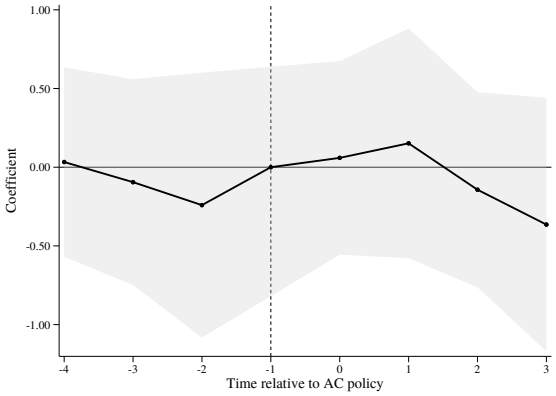
(d) 12 months

Notes: Each panel displays coefficients and 95% confidence intervals from Equation (4.3), where the dependent variable equals one if the graft fails—through graft loss, re-transplantation, or patient death—within the stated horizon of transplant. In Panels (c) and (d), the window ends at event year +2 because the most recent transplant cohorts lack full follow-up at these horizons. Standard errors are clustered at the MELD-tier × DSA cell level.

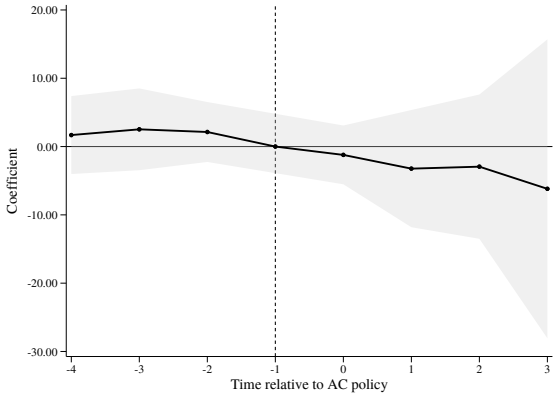
Figure A4: Event Studies: Accepted-Liver Quality



(a) Donor age (years)



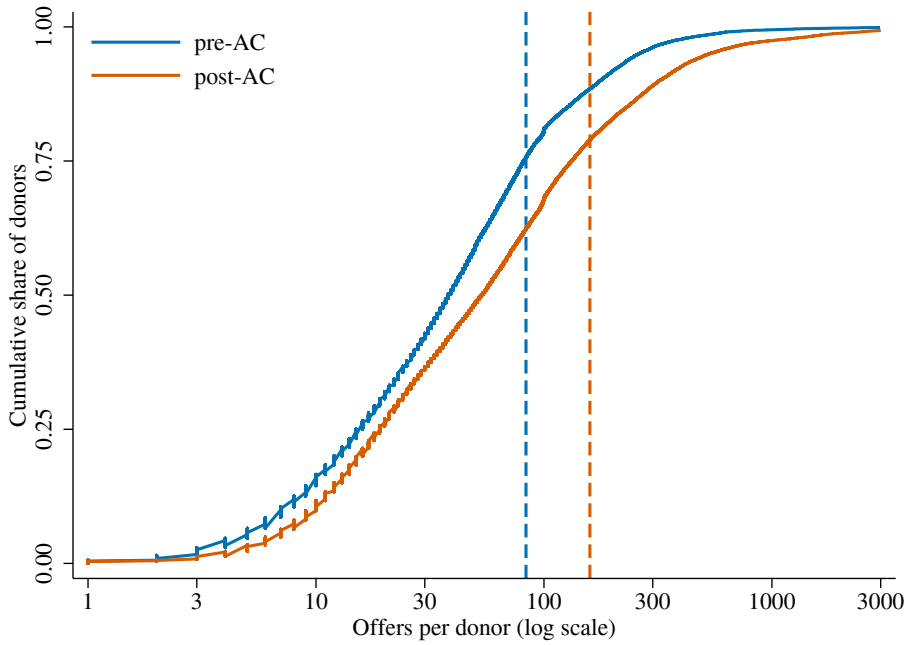
(b) DCD



(c) Cold-ischemia time (hours)

Notes: Each panel displays coefficients and 95% confidence intervals from Equation (4.3), where the dependent variable is the indicated characteristic of the accepted liver; the sample is deceased-donor transplants, as in Table 8. Standard errors are clustered at the MELD-tier × DSA cell level.

Figure A5: Cumulative Distribution of Batch Size



B. Constructing a Proxy for Liver Quality

The SRTR data do not provide a single summary measure of liver quality. Instead, candidates observe a rich set of donor characteristics, including demographics, cause of death, medical history, and laboratory values. Candidates combine this information with clinical priors to form an implicit assessment of underlying organ quality, which guides their decision to accept or decline an offer.

A large literature seeks to summarize liver quality using pre-transplant characteristics. The Donor Risk Index (DRI) aggregates donor characteristics into a single index (Feng et al., 2006). The Survival Outcomes Following Liver Transplantation (SOFT) score incorporates both donor and recipient attributes to predict 90-day post-transplant mortality (Rana et al., 2008). These approaches rely on linear combinations of covariates and may fail to capture the complex, nonlinear relationships that affect transplant outcomes.

Recent work applies machine learning (ML) methods to improve prediction. Molinari et al. (2019) use neural networks and classification trees to predict 90-day post-transplant survival, and Ershoff et al. (2020) applied deep neural networks for similar tasks. However, survival-based measures reflect both organ quality and recipient health. They are therefore not well suited for isolating the quality of the liver itself.

I construct a data-driven proxy for perceived liver quality by predicting organ disposition using only donor characteristics. The maintained assumption is that the probability that a liver is accepted reflects the transplant community's aggregate assessment of its

quality.

The outcome variable is an indicator equal to one if a recovered liver is not transplanted (nonuse) and zero if it is transplanted. I estimate four models: an OLS regression baseline and three flexible classifiers, Random Forest, XGBoost, and a multilayer perceptron (MLP). All models are trained using pre-policy data to avoid contamination from the Acuity Circles reform.

The SRTR data contain more than 100 donor covariates. After excluding variables with substantial missingness, I retain 59 donor characteristics. I randomly reserve 15% of the data as a hold-out test set. Hyperparameters are tuned using a validation split within the training data. Final model performance is evaluated on the test set.

All four models achieve similar predictive performance. Test-set AUC ranges from 0.783 to 0.794, with XGBoost performing best. The small performance differences indicate that the models produce similar rankings of livers by nonuse risk.

Table A1. Model performance (test-set AUC)

	OLS	Random Forest	XGBoost	MLP
AUC (test set)	0.7829	0.7897	0.7943	0.7847

Given similar predictive performance, model selection is guided by the intended use of the quality proxy. I use predicted probabilities rather than binary classification, so I require a smooth and finely graded measure.

Tree-based methods, both Random Forest and XGBoost, produce piecewise-constant predictions that generate discrete jumps in the probability distribution. In contrast, the MLP produces smooth predictions and captures higher-order interactions among donor attributes. This flexibility allows it to approximate the continuous quality assessment that underlies clinical decisions.

For each liver, I define the quality measure as:

$$\text{Quality} = 1 - P(\text{Nonuse})$$

where the probability is obtained from the MLP model. Higher values correspond to livers that are more likely to be desirable.

C. Proofs

C.1. Proof of Proposition 1

Let V^r solve Equation (5.1) with terminal condition $V^r(x, T_i) = D(x)$, and let $O(x,t)$ solve Equation (5.2) with the same terminal condition $O(x, T_i) = D(x)$. Define

$$V^r = \tilde{V}^r + O, \qquad \Gamma = \tilde{\Gamma} + O$$

Substituting into Equation (5.1) yields:

$$\begin{aligned}\rho(\tilde{V}^r + O) - (\dot{\tilde{V}}^r + \dot{O}) = w(x,t) &+ \sum_{x' \neq x} F_{xx'}\Big[\tilde{V}^r(x',t;d) - \tilde{V}^r(x,t;d) + O(x',t) - O(x,t)\Big] \\ &+ \delta(x)\big[D(x) - \tilde{V}^r(x,t;d) - O(x,t)\big] \\ &+ \sum_{c \in \mathcal{C}_r} \lambda_c(x,d) \iint \pi_{ij} \int \max\Big\{\tilde{\Gamma} + \epsilon - \tilde{V}^r(x,t;d),\, 0\Big\}\, dH(\epsilon)\, dG_c(q,\ell).\end{aligned}$$

Rearranging terms,

$$\begin{aligned}\rho\tilde{V}^r - \dot{\tilde{V}}^r + \big(\rho O - \dot{O}\big) = w(x,t) &+ \sum_{x' \neq x} F_{xx'}[O(x',t) - O(x,t)] + \delta(x)\big[D(x) - O(x,t)\big] \\ &+ \sum_{x' \neq x} F_{xx'}\Big[\tilde{V}^r(x',t;d) - \tilde{V}^r(x,t;d)\Big] - \delta(x)\tilde{V}^r(x,t;d) \\ &+ \sum_{c \in \mathcal{C}_r} \lambda_c(x,d) \iint \pi_{ij} \int \max\Big\{\tilde{\Gamma} + \epsilon - \tilde{V}^r(x,t;d),\, 0\Big\}\, dH(\epsilon)\, dG_c(q,\ell).\end{aligned}$$

By Equation (5.2), the bracketed term satisfies

$$\rho O - \dot{O} = w(x,t) + \sum_{x' \neq x} F_{xx'}[O(x',t) - O(x,t)] + \delta(x)\big[D(x) - O(x,t)\big].$$

Substituting and canceling terms gives

$$\begin{aligned}\rho\tilde{V}^r - \dot{\tilde{V}}^r = &\sum_{x' \neq x} F_{xx'}\Big[\tilde{V}^r(x',t;d) - \tilde{V}^r(x,t;d)\Big] - \delta(x)\tilde{V}^r(x,t;d) \\ &+ \sum_{c \in \mathcal{C}_r} \lambda_c(x,d) \iint \pi_{ij} \int \max\Big\{\tilde{\Gamma}(x,q,\ell,t) + \epsilon - \tilde{V}^r(x,t;d), 0\Big\}\, dH(\epsilon) dG_c(q,\ell)\end{aligned}$$

which coincides with Equation (5.3).

Finally, since $O(x,t)$ cancels from both Γ and V^r, the acceptance rule

$$\Gamma + \epsilon \geq V^r \quad \Longleftrightarrow \quad \tilde{\Gamma} + \epsilon \geq \tilde{V}^r$$

is invariant to the normalization. □

C.2. Proof of Proposition 2

The candidate declines an offer if and only if

$$\Gamma + \epsilon_{ij} < V^r$$

or equivalently $\epsilon_{ij} < V^r - \Gamma$. Under Assumption 1(ii), ϵ_{ij} has the standard-logistic distribution with CDF H, so

$$P_{ijt} = \Pr(\epsilon_{ij} < V^r(x_i, t; d_i) - \Gamma(x_i, q_j, \ell_{ij})) = H(V^r(x_i, t; d_i) - \Gamma(x_i, q_j, \ell_{ij}))\,.$$

Since $H(z) = 1/(1 + \exp(-z))$ is strictly increasing, inversion yields

$$V^r(x_i, t; d_i) - \Gamma(x_i, q_j, \ell_{ij}) = \log\left(\frac{P_{ijt}}{1 - P_{ijt}}\right),$$

establishing part (i).

For part (ii), define the option value

$$\psi \equiv \int \max\{\Gamma + \epsilon - V^r, 0\}\, dH(\epsilon) = \mathbb{E}\big[\max\{\Gamma + \epsilon - V^r, 0\}\big].$$

Using the identity $\mathbb{E}[\max\{Y, 0\}] = \int_0^\infty \Pr(Y > s)ds$, I obtain

$$\begin{aligned}
\psi &= \int_0^\infty \Pr(\Gamma + \epsilon - V^r > s)\, ds \\
&= \int_0^\infty \Pr(\epsilon > V^r - \Gamma + s)\; ds \\
&= \int_0^\infty [1 - H(V^r - \Gamma + s)]\, ds.
\end{aligned}$$

Substituting the logistic CDF,

$$1 - H(z) = \frac{1}{1 + \exp(z)},$$

gives

$$\psi = \int_0^\infty \frac{1}{1 + \exp(V^r - \Gamma + s)}\, ds = \log(1 + \exp(\Gamma - V^r))\,.$$

Using part (i),

$$\exp(\Gamma - V^r) = \frac{1 - P_{ijt}}{P_{ijt}},$$

so

$$\psi(P_{ijt}) = \log\left(1 + \frac{1 - P_{ijt}}{P_{ijt}}\right) = -\log P_{ijt},$$

which completes the proof. □