



MIT Open Access Articles

Dynamic Incentives in Wait List Mechanisms

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Agarwal, Nikhil, Ashlagi, Itai, Somaini, Paulo and Waldinger, Daniel. 2018. "Dynamic Incentives in Wait List Mechanisms."
As Published	10.1257/PANDP.20181079
Publisher	American Economic Association
Version	Final published version
Citable link	https://hdl.handle.net/1721.1/138035
Terms of Use	Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

Dynamic Incentives in Wait List Mechanisms[†]

By NIKHIL AGARWAL, ITAI ASHLAGI, PAULO SOMAINI, AND DANIEL WALDINGER*

Many scarce public resources are allocated through a wait list. A particularly salient example is the kidney wait list in the United States, which now has almost 100,000 patients waiting for a lifesaving transplant.

While the design of such systems has garnered significant research attention, most theoretical results yield answers that depend on the primitives of the market (compare Su and Zenios 2004; Leshno 2017; Bloch and Cantala 2017). Moreover, the state of the art empirical methods used to prospectively evaluate wait list designs do not incorporate the dynamic incentives that are central to the theoretical literature. Perhaps the most prominent example of an empirically guided redesign is the 2014 reform of the deceased donor kidney allocation system. The kidney committee used the Kidney-Pancreas Simulated Acceptance Module (KPSAM) to predict the transplants that would result from various organ allocation rules. KPSAM allowed the committee to experiment with the priority system, evaluate outcomes and make an informed decision. However, this decision tool simplifies patient acceptance behavior by assuming that it is invariant to priority rules and therefore ignores patients' dynamic incentives.

This article uses a combination of theoretical and empirical arguments to show that considering dynamic incentives is important

for evaluating wait list mechanisms. We present examples to illustrate the interaction between dynamic incentives, preferences, and wait list design (Section II) as well as evidence consistent with agent choices being influenced by dynamic incentives (Section III). These results motivate methodological and empirical work studying dynamic assignment systems more broadly.

I. Background

In the United States, kidneys from deceased donors are allocated through a wait list. Each organ is offered to patients according to an organ-specific priority rule. Patients may accept or decline an offer, with no penalty for refusing. Each organ is assigned to the highest priority biologically compatible agent that accepts the organ. These assignments must take place quickly because it is difficult to maintain a kidney's viability after the donor has deceased.

A new priority system for allocating kidneys was adopted on December 4, 2014. Prior to the reform, organs were offered first to patients that had a rare perfect immunological match, then to patients within the local area of the donor, next to patients in the broader geographical region, and finally to patients nationwide. Within each group, patients were ordered according to points awarded based on patient and donor characteristics with ties broken by how long the patient had waited. Points were awarded to patients that had highly sensitive immune systems, that were pediatric, and that were immunologically well-matched to the donor.

The reform aimed to alleviate the inefficiency, organ waste, and inequity that resulted, in part, from a growing wait list that emphasized waiting time (see Israni et al. 2014). Despite the long wait list, about 20 percent of medically viable kidneys are discarded each year. Reports by the kidney committee suggest that the goal of the reform was to allocate organs to patients that would benefit most from them, to reduce waste, and to avoid hurting any specific group of

* Agarwal: MIT and NBER (email: agarwaln@mit.edu); Ashlagi: Stanford University (email: iashlagi@stanford.edu); Somaini: Stanford University and NBER (email: soma@stanford.edu); Waldinger: MIT (email: dwalding@mit.edu). Agarwal and Ashlagi acknowledge support from the National Science Foundation grants SES-1729090 and SES-1254768, respectively. Disclaimer: The data reported here have been supplied by UNOS as the contractor for the Organ Procurement and Transplantation Network (OPTN). The interpretation and reporting of these data are the responsibility of the author(s) and in no way should be seen as an official policy of or interpretation by the OPTN or the US Government.

[†] Go to <https://doi.org/10.1257/pandp.20181079> to visit the article page for additional materials and author disclosure statement(s).

patients.¹ The new system prioritizes patients in the top quintile of predicted post-transplant survival probability for kidneys in the lowest quintile of estimated risks of post-transplant organ failure. It also increased national organ sharing for patients that have the most sensitive immune systems.

II. Dynamic Mechanism Design

We now show that it is important to consider dynamic incentives when designing a wait list mechanism because design trade-offs depend on the preferences of agents.

A. A Model of Decisions on a Wait List

We consider a model similar to the one described in Bloch and Cantala (2017). There are $n \geq k$ agents waiting for an object, where k is a positive constant. Each day, an object arrives and can be offered to $l < k$ agents. Offers are made sequentially, according to a predefined priority order. Each object must be assigned immediately and is wasted if none of the top l agents accepts. Each agent makes an accept-reject decision. She is assigned the object if she accepts it, and is then removed from the wait list. Agents that reject the object remain on the wait list, and may accept a future offer. We assume that waiting incurs a per-period cost of c .

An optimal decision rule recommends acceptance if and only if the value of receiving an assignment is higher than the value of waiting for a future offer. Specifically, let α be the value of an object to an agent, $s \in S$ denote the agent's position on the wait list, and let F_s be the cumulative distribution function of the value of objects offered to an agent in position s . This equilibrium distribution, F_s , depends on the strategies of other agents on the wait list. We suppress this dependence from the notation and focus on the optimal decision of each agent. This optimal decision can be written as

$$(1) \quad a^* = 1\{\alpha > V(s; M)\},$$

where $V(s; M)$ is the expected net present value of continuing to wait, and M denotes the mechanism in use.

¹We obtained the reports of the committee from the Communications Office at the United Network for Organ Sharing (UNOS).

The acceptance thresholds depend on the agent's position s and the mechanism M because these quantities influence the value of waiting. Agarwal et al. (2018) show how to estimate the values of various object attributes in a similar dynamic choice model using data on accept-reject decisions and knowledge of the mechanism. They apply their methods to data on decisions from the deceased donor kidney wait list.

This approach is in contrast to the one taken in KPSAM, which assumes that acceptance behavior is invariant to the mechanism, that is $a^* = 1\{\alpha > V\}$, where V does not depend on M or on s . It is easy to see that the simpler model is likely to yield incorrect predictions if the mechanism influences the value of waiting.

B. Preferences and Design

We now compare the first come first served (FCFS) mechanism, which orders patients based on how long they have waited, with a lottery mechanism (LM). We assume that LM randomly orders the k agents that have waited the longest (see Leshno 2017 for a similar mechanism). We present results for two different preference models. In both models, values are drawn from the cumulative distribution function F . The first, which we call *vertical preferences*, assumes that the value of each object is common across agents. The second, which we call *horizontal preferences*, assumes that agents' values for a given object are drawn independently. For simplicity, we assume that F has support only on the unit interval $[0, 1]$ and focus on the case when $l = 2$, so that the object is wasted if the top two agents decline it.

1. Vertical Preferences.—FCFS: In equilibrium, the agent in position $i \in \{1, 2\}$ will accept the object if and only if the value of the object exceeds a cutoff α_i . The value of waiting for the agent in position 1 solves

$$V(1) = \int_{\alpha_1}^1 x dF + F(\alpha_1) V(1) - c.$$

Because this agent must be indifferent between accepting an object with value α_1 and continuing, $V(1) = \alpha_1$. Therefore, α_1 solves

$$(2) \quad \int_{\alpha_1}^1 (x - \alpha_1) dF = c.$$

Similarly, the value of waiting for the agent in the second position solves

$$V(2) = (1 - F(\alpha_1)) V(1) + \int_{\alpha_2}^{\alpha_1} x dF + F(\alpha_2) V(2) - c.$$

Using the expression $V(i) = \alpha_i$ and equation (2), we have that α_2 solves

$$(3) \quad \int_{\alpha_2}^1 (x - \alpha_2) dF = 2c.$$

The fraction of objects wasted in this model is $F(\alpha_2)$, and the expected value of each assigned object is $E[\alpha | \alpha > \alpha_2]$.

LM: Observe that all agents use the same cutoff in equilibrium, α_{LM} . Because values are perfectly correlated, an agent accepts the object only if she was (randomly) chosen to be at the top of the list. Therefore, the expected waiting time before the next offer is k and the threshold α_{LM} solves

$$(4) \quad \int_{\alpha_{LM}}^1 (x - \alpha_{LM}) dF = kc.$$

Comparison: Observe that $\int_{\alpha}^1 (x - \alpha) dF$ is decreasing in α . Therefore, waste is lower in the lottery mechanism because $\alpha_2 \geq \alpha_{LM}$ for all $k > 2$. However, the value conditional on assignment is higher under FCFS. A social planner that accounts for waiting costs and the value of assignments may prefer lowering waste if the waiting list is long enough or if c is high enough. Indeed, an agent's refusal in FCFS can impose a negative externality on agents lower on the list because a rejected object is of lower quality and waiting time increases.

2. *Horizontal Preferences.*—**FCFS:** The cutoff for the agent in the first position, α_1 , is identical to the vertical case. The value of waiting for the agent in the second position is given by

$$V(2) = (1 - F(\alpha_1)) V(1) + F(\alpha_1) \left[\int_{\alpha_2}^1 x dF + F(\alpha_2) V(2) \right] - c.$$

The cutoff α_2 now solves

$$(5) \quad F(\alpha_1) \int_{\alpha_2}^1 (x - \alpha_2) dF + (1 - F(\alpha_1)) (\alpha_1 - \alpha_2) = c.$$

LM: The value of waiting for the top k agents solves

$$V = \frac{1}{k} (1 + F(\alpha)) \left[\int_{\alpha}^1 x dF + F(\alpha) V \right] + \left[1 - \frac{1}{k} (1 + F(\alpha)) \right] V - c.$$

The first term represents the case when an agent receives an offer, and the second term represents the remaining case. Setting $V = \alpha_{LM}$, the cutoff α_{LM} solves

$$(6) \quad (1 + F(\alpha)) \int_{\alpha}^1 (x - \alpha) dF = kc.$$

A comparison of equations (4) and (6) shows that each agent is more selective if preferences are horizontal. However, waste might be lower because one of many agents may accept the object.

Comparison: Again, it can be shown that $\alpha_{LM} \leq \alpha_2 \leq \alpha_1$ if $k > 2$. However, an agent rejecting an offer does not indicate that the object is undesirable. Therefore, agents at the top of the list can exert a positive externality on other agents by rejecting offers.

C. A Numerical Example

Table 1 presents a numerical example with $\alpha \sim U[0, 1]$, $c = 1/6$, and $k = 3$. Under the chosen parameters, the LM results in the same waste and match value conditional on assignment (MV) irrespective of whether preferences are vertical or horizontal. Moving from LM to FCFS results in a higher match value and higher waste. However, waste increases less and match value increases more under horizontal preferences than under vertical preferences. Indeed, the unconditional expected value obtained from each object (EV) is higher in FCFS than in LM if preferences are horizontal, but the reverse is

TABLE 1—FCFS VERSUS LOTTERY

	Cutoffs	Waste	MV	EV
Vertical				
FCFS	(0.42, 0.18)	0.18	0.59	0.48
LM	0	0	0.5	0.5
Horizontal				
FCFS	(0.42, 0.309)	0.13	0.69	0.60
LM	0	0	0.5	0.5

Note: The cutoffs for FCFS are described as the pair (α_1, α_2) , while LM has only one cutoff.

true if preferences are vertical. If preferences are vertical, FCFS is dominated by LM.

Higher waste can outweigh the matching benefits of FCFS if preferences are horizontal. Due to the increase in waste when moving from LM to FCFS, the expected wait between assignments increases from 1 to $1/(1-0.13) \approx 1.15$. The social costs of waiting therefore increase by $n \times c \times 0.15$. Therefore, FCFS outperforms LM if $n \leq 4$, but not otherwise.² Taken together, these results show that a social planner's decision between these two mechanisms ultimately depend upon the nature of primitives.

The result that FCFS may better match agents to more preferred objects is not specific to this example. Arnosti and Shi (2017) also find that FCFS produces higher match value than LM. Indeed, Bloch and Cantala (2017) show that FCFS yields better match value than any mechanism that gives agents that have waited longer weakly higher priority. However, this result depends on the nature of primitives as shown in Leshno (2017), who studies a model with agents that have preferences for a specific type of object. He shows that it may be optimal to run a lottery among agents at the "top positions" to influence selectivity and reduce misallocation. Similarly, using a model with stochastic arrivals and vertical preferences, Su and Zenios (2004) show that a last come first served mechanism reduces waste and improves social welfare relative to FCFS if preferences are vertical.

²The variable n is the expected equilibrium queue length, which is bounded if agents depart exogenously without an assignment. The increased waste could result in an endogenously longer queue.

These examples assume that the planner does not have information about how much different agents value various object types. It is easy to construct examples in which the optimal mechanism prioritizes agents based on observed predictors of value. Similarly, mechanisms that prioritize agents based on predictable differences in waiting costs can improve welfare.

III. Evidence on Dynamic Incentives

A testable implication of the model is that agents with a low option value of waiting are more likely to accept an offer of a given quality. This is in contrast to static choice models such as the one used in KPSAM. We now present descriptive evidence consistent with this hypothesis.

A. Data

This study uses data from the Organ Procurement and Transplantation Network (OPTN). The OPTN data system includes data on all donors, wait-listed candidates, and transplant recipients in the United States submitted by the members of the Organ Procurement and Transplantation Network (OPTN). The Health Resources and Services Administration (HRSA), US Department of Health and Human Services provides oversight to the activities of the OPTN contractor. The primary dataset on the wait list, the Potential Transplant Recipient (PTR) dataset, contains the offers made and patient decisions. This dataset is drawn from the backbone system used to coordinate offers and decisions.

We restrict attention to data on the kidney wait list and to acceptance decisions between January 1, 2010 and December 31, 2013 by all patients registered in Donor Service Areas (DSAs) that used the standard allocation rules. Except in cases of a perfect tissue-type match, allocation takes place based on geography, with DSAs constituting the smallest unit. Our analysis covers 37 out of the 58 DSAs in the United States. Together, these DSAs account for 59 percent of the patient population.

Table 2 describes the data, which include offers from 17,811 donors to 105,536 patients. The total number of offers is close to 17 million, but most of these offers are screened out by pre-set acceptance criteria set by patients (we

TABLE 2—SUMMARY STATISTICS

	Observations
Donation service areas (DSAs)	37
Transplant centers	144
Donors	17,811
Patients	105,536
Offers	16,981,773
Offers that met screening criteria	3,593,988
Acceptances	31,385
Acceptances resulting in transplant	22,154

Notes: Some acceptances do not result in a transplant due to a final biological test called a crossmatch. A positive crossmatch indicates that the patient would reject the organ.

simulate the mechanism in order to recover all offers, whether they were screened out or not). While most accepted offers result in a transplant, some do not because a final immunological test can fail. In these cases, the organ is assigned to the next highest priority patient that accepted it and passes the test. We treat all acceptances as indicative of the value of assignment exceeding the value of waiting.

Surprisingly, less than 0.2 percent of all offers are accepted by patients, but the total number of transplants indicate that the organs from most donors are transplanted. As Agarwal et al. (2018) discuss, this low acceptance rate is a result of relatively undesirable organs being offered to several thousand patients. Indeed, the acceptance rate amongst the top 10 offers is 18.1 percent, and still higher at the very top of the list.

B. Empirical Strategy and Results

The ideal empirical strategy would compare patients that have the same characteristics, but face different option values for exogenous reasons. One strategy would be to use variation induced by differences in priority while holding all patient characteristics constant. Unfortunately, we are not aware of any such sources of variation in the kidney wait list.

Instead, we use biological compatibility to study how variation in offer rates and option values affects acceptance decisions. A patient that is likely to be biologically incompatible with a large number of donors should have a lower option value of waiting. If dynamic incentives are important, we would expect patients with sensitive immune systems to be

less selective. The main concern with this strategy is that immune sensitization also influences the value from a transplant. Sensitized patients are less likely to be healthy, and to benefit from a transplant relative to remaining on dialysis (see Naji et al. 2017). This fact would most likely negatively bias the estimated relationship between acceptance rates and patient immune sensitivity. We use the standard measure of patient sensitivity, Calculated Panel Reactive Antibodies (CPRA). CPRA is the percentage of donors from a representative pool with whom a patient is expected to be tissue-type incompatible.

Under our hypotheses, one should expect CPRA to be negatively correlated with expected offer rates and positively correlated with acceptance rates. Figure 1, panel A, shows that offer rates decline with CPRA, despite the fact that patients above a CPRA of 80 percent receive higher priority. Figure 1, panel B, shows an increasing acceptance rate in CPRA, particularly for high CPRA patients. The main confounding factors are that the priority system and other patient and donor characteristics can affect the value of a transplant. While this concern cannot be fully addressed, one can assess the extent to which this relationship is robust to the inclusion of a rich set of covariates.

Table 3 presents estimates from a linear probability model of acceptance as a function of CPRA, controlling for a variety of patient, donor, and match-specific characteristics. Column 1 controls only for an indicator of whether the patient is completely unsensitized. The coefficient on CPRA is positive and statistically significant. Column 2 controls for all patient-specific indicators of priority types. This reduces the estimated effects of CPRA because the CPRA thresholds of 20 percent and 80 percent are cutoffs in the priority system. Columns 3 and 4 add controls for additional patient and donor characteristics. The CPRA coefficient estimates are a little lower, but still positive and statistically significant. Column 5 controls more flexibly for interactions between CPRA and indicators of tissue-type similarity because high CPRA patients with sensitive immune systems may differentially prefer kidneys that their bodies are less likely to reject. These interactions barely affect the CPRA coefficient estimate. These results are consistent with dynamic incentives influencing acceptance decisions.

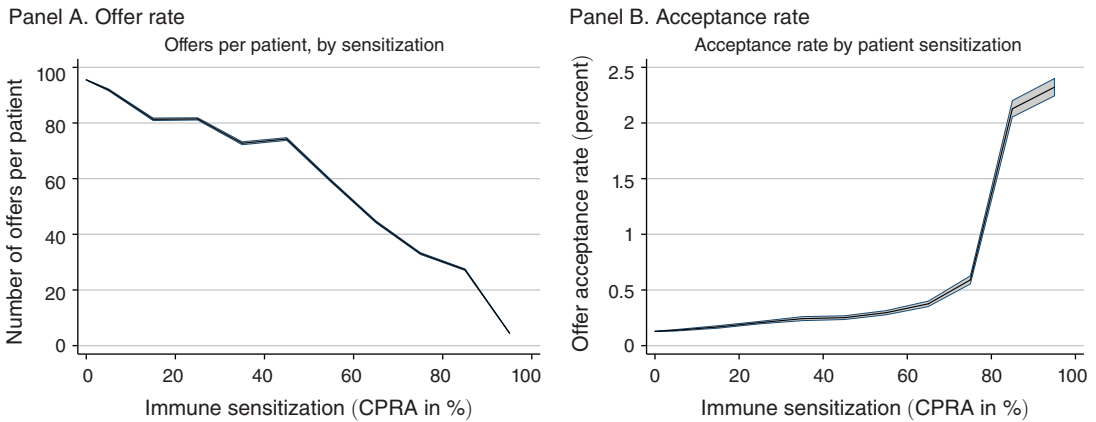


FIGURE 1. OFFER AND ACCEPTANCE RATE BY CPRA

TABLE 3—EVIDENCE ON RESPONSE TO DYNAMIC INCENTIVES

	Dependent variable: offer accepted				
	(1)	(2)	(3)	(4)	(5)
Calculated panel reactive antibodies (CPRA)	0.0141 (0.000389)	0.00466 (0.000271)	0.00382 (0.000262)	0.00239 (0.000215)	0.00232 (0.000214)
Variables affecting priority		X	X	X	X
Patient characteristics			X	X	X
Donor and match characteristics				X	X
Interaction between CPRA and # HLA mismatches					X
Observations	16,981,773	16,981,773	16,981,773	16,981,773	16,981,773
R ²	0.002	0.005	0.007	0.234	0.234

Notes: CPRA is measured on a [0,1] scale at the time of the offer. Column 1 controls for a CPRA = 0 indicator. Column 2 adds indicators for CPRA ≥ 0.2, CPRA ≥ 0.8, and age < 18, as well as waiting time indicators and linear controls for 1–3, 3–5, and >5 years. Column 3 adds other patient characteristics. Column 4 adds donor and match characteristics. Column 5 adds interactions between CPRA and # HLA mismatches. Patient characteristics are indicators for age 18–35, 35–50, and 50–65; indicators for blood type, diabetes, and the patient’s transplant center; and linear controls and indicators for dialysis time 1–3, 3–5, 5–10, and >10 years. Donor characteristics are linear age; linear creatinine clearance with indicators for 0.6–1.8 and >1.8; and indicators for diabetes, cardiac death (DCD), and expanded criteria donor (ECD). Match characteristics are linear in # HLA mismatches; indicators for zero HLA mismatch, 0 and 1 DR mismatch, identical blood type, offer year, and local donor; linear controls for (+) and (–) age difference; and interactions between local and zero-HLA mismatch, local and donor age, donor over 40 and pediatric patient, donor over 55 and patient age 18–35, and donor over 60 and patient age 35–50 and over 50. Standard errors, clustered by donor, are in parentheses.

REFERENCES

Agarwal, Nikhil, Itai Ashlagi, Michael Rees, Paulo Somaini, and Daniel Waldinger. 2018. “An Empirical Framework for Sequential Assignment: The Allocation of Deceased Donor Kidneys.” Unpublished.

Arnosti, Nick, and Peng Shi. 2017. “How (Not) to Allocate Affordable Housing.” Unpublished.

Bloch, Francis, and David Cantala. 2017. “Dynamic Assignment of Objects to Queuing Agents.” *American Economic Journal: Microeconomics* 9 (1): 88–122.

Israni, A. K., N. Salkowski, S. Gustafson, J. J. Snyder, J. J. Friedewald, R. N. Formica, X. Wang, et al. 2014. “New National Allocation Policy for Deceased Donor Kidneys in the United States and Possible Effect on Patient

- Outcomes.” *Journal of the American Society of Nephrology* 25 (8): 1842–48.
- Leshno, Jacob D.** 2017. “Dynamic Matching in Overloaded Waiting Lists.” Unpublished.
- Naji, M., A. D. Stanton, O. Ekwenna, G. Mitro, M. Rees, and J. Ortiz.** 2017. “Alemtuzumab Equalizes Short Term Outcomes in High Risk PRA Patients: Long Term Outcomes Suffer.” *Journal of Clinical Experimental Transplantation* 2 (2).
- Su, Xuanming, and Stefanos Zenios.** 2004. “Patient Choice in Kidney Allocation: The Role of the Queuing Discipline.” *Manufacturing and Service Operations Management* 6 (4): 280–301.