*American Economic Review 2019, 109(11): 4026–4070*
*https://doi.org/10.1257/aer.20180771*

# Market Failure in Kidney Exchange[†]

*By* Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton R. Featherstone, and Ömer Karaduman*

*We show that kidney exchange markets suffer from market failures whose remedy could increase transplants by 30 to 63 percent. First, we document that the market is fragmented and inefficient; most transplants are arranged by hospitals instead of national platforms. Second, we propose a model to show two sources of inefficiency: hospitals only partly internalize their patients' benefits from exchange, and current platforms suboptimally reward hospitals for submitting patients and donors. Third, we calibrate a production function and show that individual hospitals operate below efficient scale. Eliminating this inefficiency requires either a mandate or a combination of new mechanisms and reimbursement reforms.* (*JEL* D24, D47, I11)

The kidney exchange market in the United States enables approximately 800 transplants per year for kidney patients who have a willing but incompatible live donor. Exchanges are organized by matching these patient–donor pairs into swaps that enable transplants. Each such transplant extends and improves the patient's quality of life and saves hundreds of thousands of dollars in medical costs, ultimately creating an economic value estimated at more than one million dollars.[1] Since monetary compensation for living donors is forbidden and deceased donors

*Agarwal: Department of Economics, MIT, 77 Massachusetts Avenue, E52-406, Cambridge, MA 02139, and NBER (email: agarwaln@mit.edu); Ashlagi: Stanford University, Huang Engineering Center 262, 475 Via Ortega, Stanford, CA 94305 (email: iashlagi@stanford.edu); Azevedo: Wharton, 1455 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104 (email: eazevedo@wharton.upenn.edu); Featherstone: Wharton, 1352 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104 (email: claytonf@wharton.upenn.edu); Karaduman: Department of Economics, MIT, 77 Massachusetts Avenue, E52-300, Cambridge, MA 02139 (email: omerk@mit.edu). Thomas Lemieux was the coeditor for this article. First version: September 20, 2017. We thank Alex Garza, Sarah Taranto, and Jennifer Wainwright of the United Network for Organ Sharing; Jonathan Kopke and Michael Rees of the Alliance for Paired Kidney Donation; Cathi Murphey and Adam Bingaman of Methodist Hospital at San Antonio; and Garet Hill of the National Kidney Registry for their expertise and access to data. We thank Glenn Ellison, Amy Finkelstein, Parag Pathak, Al Roth, Jean Tirole, Utku Ünver, Mike Whinston, and seminar participants for helpful discussions. Abigail Ostriker provided excellent research assistance. Azevedo acknowledges support from the Alfred P. Sloan Foundation (FG2018-10387) and Wharton's Dean's Research Fund, and Agarwal and Ashlagi from the National Science Foundation (SES-1729090 and SES-1254768, respectively). Disclaimer: the data reported here have been supplied by the United Network for Organ Sharing (UNOS) as the contractor for the Organ Procurement and Transplantation Network (OPTN). The interpretation and reporting of these data are the responsibility of the authors 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/aer.20180771 to visit the article page for additional materials and author disclosure statements.

[1] Kidney exchange is among a handful of recent innovations that both improve health care delivery and save costs (see Chandra and Skinner 2012). Transplantation roughly doubles the life expectancy of patients with end-stage renal disease and is cheaper than the alternative treatment of dialysis. Medicare provides nearly universal coverage, irrespective of age, for patients with end-stage renal disease, which comprises about 7 percent of Medicare's annual budget (see United States Renal Data System 2016). The cost savings of transplantation relative to dialysis alone have been estimated to be over $270,000 (see Section I).

are in increasingly short supply, kidney exchange markets play an important role in mitigating the shortage of organs available for transplant.[2]

This paper shows that, despite significant success, US kidney exchange suffers from market failures that result in hundreds of lost transplants per year. Our descriptive evidence shows that the market is fragmented and operates inefficiently. The inefficiency arises from two standard market failures. First, kidney exchange platforms use inefficient mechanisms: hospitals are not rewarded for submitting high social value patients and donors to the platform. Second, there are significant agency problems: hospitals face most of the costs of participating in national platforms but receive only a fraction of the benefits. These market failures are serious, but fixable. We show how to combine theory and data to design efficient mechanisms, and discuss policies such as reimbursement reforms and mandates. Our estimates suggest that fixing these problems would generate between 240 to 500 additional transplants per year (25 to 63 percent of the current total).

Our argument has three parts. First, we use administrative datasets to show that the market is fragmented, inefficient, and shows signs of agency problems. Second, we develop a simple model to explain the market failures and propose solutions. Third, we combine the model and data to estimate the magnitude of the inefficiencies and to design practical alternative mechanisms and policies.

The first part documents three key facts using data on all transplants in the United States and proprietary data from the three largest US kidney exchange platforms. First, the market is highly fragmented. Instead of most transactions being arranged by a few large platforms, 62 percent of kidney exchange transplants involve patients and donors from the same hospital. Second, we find direct evidence of inefficient exchanges in the market. Kidney exchanges performed within hospitals often transplant kidneys from easy-to-match donors to easy-to-match patients, a practice which existing theory has shown to be inefficient (Roth, Sönmez, and Ünver 2007). Third, hospital behavior is inconsistent with pure maximization of patient welfare. Evidence suggests that hospitals are sensitive to the financial and administrative transaction costs of participating in kidney exchange, even though these costs are small relative to the social value of transplants. Many hospitals do not participate in national platforms, and even when they do, the typical hospital does not conduct all kidney exchanges through a national platform.

The second part develops a model to explain these facts and design policy responses. Although kidney exchange markets do not directly use monetary incentives to acquire organs, we can analyze them with standard neoclassical producer theory. A kidney exchange platform produces a final good (transplants) from intermediate goods (submissions of patients and donors) supplied by a competitive fringe (hospitals) according to a production function. This model is motivated by three key institutional features. First, hospitals are the key decision makers steering participants toward kidney exchange (Roth, Sönmez, and Ünver 2005; Ashlagi and Roth

[2] There are over 97,000 patients currently waiting for a kidney from a deceased donor, but less than one-fifth are expected to be transplanted in the next year. Becker and Elías (2007) argues that the wait-list could be completely eliminated if there were monetary compensation for live donors. However, this type of transaction is widely panned by bioethicists, and almost all countries forbid it. The National Organ Transplantation Act prohibits compensating donors to acquire organs in the United States, but explicitly allows for kidney exchange through the Charlie W. Norwood Living Organ Donation Act.

2014; Rees et al. 2009). Second, due to biological compatibility constraints, some patients and donors generate considerably more transplants than others when they join a platform. Third, the structure of optimal matches makes transplants a natural numéraire good. Platforms can effectively transfer transplants from one hospital to another by choosing which hard-to-match patients to match.

Much of the economics of kidney exchange is determined by the shape of the production function. Returns to scale determine how large a platform must be to attain efficiency, which tells us whether a fragmented market can be efficient. Marginal products determine the values of different types of patients and donors to the platform, which are key factors in designing efficient mechanisms.

Theorem 1 shows that inefficiency comes from the two market failures we discussed. First, platforms use inefficient mechanisms. When a hospital submits a patient or a donor to a platform, current mechanisms reward hospitals according to the *probability* with which that hospital's patient is matched. But the theorem shows that, to maximize hospital welfare, hospitals should be rewarded with the *marginal product* of their submissions (the expected number of additional transplants enabled), plus a small adjustment term. Because existing platforms do not reward hospitals based on the marginal product of their submissions, even a hospital that maximizes the number of own-patient transplants has to perform socially inefficient matches. This problem can be addressed by using a points mechanism that rewards hospitals according to marginal products. Without making the connection between kidney exchange and neoclassical producer theory, it is not obvious that this market failure exists, much less how to fix it. The second market failure is that hospital objectives may differ from pure social welfare maximization, a problem which we refer to as an agency problem. For example, hospitals may participate too little in kidney exchange because they face most of the costs but only receive a fraction of the benefits. This problem can be addressed with subsidy policies and mandates.

The third part of this paper combines theory and data to quantify inefficiency in the market and to suggest policy responses. To do so, we recover the production function using administrative data from the largest US kidney exchange platform and detailed information on matching algorithms and operational procedures.

The production function yields three sets of results. First, we measure the returns to scale and estimate the inefficiency from market fragmentation. We find that the largest kidney exchange platform is well above the minimum efficient scale, while almost all single-hospital platforms are far below the efficient scale. We estimate that the gains from moving all production to the efficient scale is at least 200 transplants per year, and likely closer to 500. These improvements correspond to an economic value of between $240 million and $500 million annually, of which approximately one-quarter is due to savings on health care costs. Thus, consistent with the descriptive evidence and the shape of the production function, fragmentation has a large efficiency cost. Under certain assumptions, mandating that hospitals participate in national kidney exchange platforms could realize most of these gains. Our results suggest that the market can support two to three national platforms at close to the efficient scale. Therefore, a stronger mandate that additionally restricts participation to a single platform would not generate further gains.

Second, we use the estimated production function to design more efficient mechanisms. Optimal mechanisms should reward submissions approximately according

to marginal products, while current mechanisms reward submissions according to probabilities of matching. We find that marginal products are considerably different from probabilities of matching, which implies that existing mechanisms are far from optimal. We discuss how optimal points mechanisms based on our estimates could be used to improve hospital incentives.

Third, we study the importance of the two market failures. The loss in hospital welfare due to the inefficient mechanism depends on the wedge between current and optimal rewards, and on the elasticity of supply from hospitals. We have estimated the wedges and the marginal products, but our data do not have enough information to estimate supply elasticities credibly. Therefore, we calculate this deadweight loss under a broad range of assumptions on elasticities. Except under the most extreme of these, the deadweight loss is significant but lower than the inefficiency due to market fragmentation. Hence, both the current mechanism and agency problems contribute significant inefficiency in the market. This finding suggests a two-pronged policy approach: improve mechanisms and encourage participation. This approach is particularly appealing if a mandate is not politically feasible.

*Relation to the Literature*.—We build upon several earlier contributions. In kidney exchange, Roth, Sönmez, and Ünver (2004, 2007) and a large subsequent literature have studied optimal matching algorithms and the technology of kidney exchange. In particular, Roth, Sönmez, and Ünver (2007) calculates the marginal products of different types of pairs in a simplified theoretical model. Our empirical estimates of marginal products are, to our knowledge, the first test of their theoretical predictions. We find qualitatively similar marginal products for most types, with differences because the empirical model accounts for features like immune sensitivity and matching frictions. Motivated by anecdotal evidence, other papers in the kidney exchange literature have theoretically analyzed the problem of hospital participation. Roth, Sönmez, and Ünver (2005) showed that hospitals may have incentives to match patients outside of an exchange and that this may compromise efficiency. Ashlagi and Roth (2014) proposes a mechanism that addresses this issue in a stylized static model by only considering exchanges in which all patients that a hospital can transplant on its own are matched.[3] Hospital rewards in their mechanism differ from long-run marginal products in a dynamic and stochastic setting when patients and donors are registered over time. Thus, Theorem 1 suggests that even if their mechanism provides good incentives in static models, it can be inefficient in practice.[4] Rees et al. (2012) advocates for reimbursement policy reform based on the argument that costs of kidney exchange are a barrier to participation in the market.

We also draw on other areas of economics that are not traditionally used in market design. Theorem 1 describes linear rewards that maximize hospital welfare. The key idea is based on analogies to linear commodity taxation (Ramsey 1927) and regulation of multi-product monopolists (Boiteux 1956). Our theoretical contribution is to apply these classic ideas and proofs to kidney exchange. A difference is that

[3] Toulis and Parkes (2015) proposes an alternative algorithm in the same lines as Ashlagi and Roth (2014).

[4] Hajaj et al. (2015) also proposes a dynamic "credit mechanism" for kidney exchange; however, their mechanism rewards each hospital based on the total number of pairs that it submits, irrespective of type. In this way, their suggested rewards differ from the marginal-product rewards that we suggest in Section IIIC.

we recognize the possibility that hospital welfare and social welfare differ, which we model as a general wedge between the two, as in the modern behavioral public finance literature (Farhi and Gabaix 2017). Our paper is also related to the industrial organization literature on platforms. Platforms bring together market participants, much like kidney exchange platforms. A key issue in this literature is the structure of optimal prices, commonly studied for two-sided platforms (Rochet and Tirole 2003, Weyl 2010). Our analysis of kidney exchange involves pricing to only one side of the market: hospitals. Finally, we measure the total number of transplants lost due to market fragmentation using a misallocation analysis. This exercise is similar in spirit to the literature documenting and quantifying the efficiency implications of productivity dispersion (see Bartelsman and Doms 2000, Hsieh and Klenow 2009).

## I. Background and Data

### A. *Basics of Kidney Exchange*

End-stage renal disease (ESRD) afflicts more than half a million Americans. The disease is almost universally covered by Medicare, even for patients under the age of 65. The Medicare ESRD program accounts for 7 percent of its budget, mostly spent on patients undergoing dialysis (United States Renal Data System 2016). The preferred treatment for ESRD patients is transplantation, which increases the quality and length of life by several years and is cheaper than dialysis. Transplantation saves approximately $270,000 per Medicare beneficiary and even more for privately insured patients (Wolfe et al. 1999, Irwin et al. 2012, Held et al. 2016). Moreover, the health risks to living donors are small. Taken together, these facts indicate that a living-donor kidney transplant has large economic value. Held et al. (2016) estimates this value at $1.1 million using a detailed cost-benefit analysis.[5]

There is a severe shortage of organs for transplantation. Each year, approximately 13,000 patients are transplanted using organs from deceased donors and another 5,500 from living donors. Demand far outstrips this supply with approximately 35,000 patients added to the deceased donor kidney wait-list in each of the recent few years. The shortage has resulted in the kidney wait-list growing to almost 100,000 patients, with about 8,000 patients per year dying or being categorized as too sick to transplant.[6] Monetary compensation cannot be used to address this shortage because of ethical and legal reasons: it is forbidden in almost every country, including the United States (Becker and Elías 2007).

Kidney exchange is an innovative way to ameliorate this shortage (Roth, Sönmez, and Ünver 2004; Sönmez and Ünver 2013). It serves patients who have a willing live donor with whom they are not biologically compatible. Such patients can swap donors with others in the same situation, enabling transplants for many patients.

[5] Most of the $1.1 million comes from gains in quality-adjusted life-years (QALYs), valued at $200,000. Even if each QALY is valued at only $100,000, the estimated economic value only drops to $660,000. This drop is less than 50 percent because the cost savings on dialysis are also significant. In 2014, Medicare paid $87,638 per year per dialysis patient but only $32,586 in post-transplant costs per year per patient (United States Renal Data System 2016, chapters 7 and 11).

[6] Statistics taken from https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/ (accessed December 21, 2017).

These swaps are organized by *kidney exchange platforms* that match patients and donors registered with them. The platforms receive three types of *submissions*. The most common type is a *pair*, consisting of a patient and a living donor. The second type is an *altruistic donor*, who is willing to donate a kidney to a stranger without requiring a transplant for an associated patient. Finally, there are some *unpaired patients*, who do not have a willing live donor.

Platforms organize transplants in two ways. The first, called a *cycle*, involves a set of pairs. The kidney from one pair's donor is transplanted into the patient in the next pair, and so on, until the cycle is closed. All transplants are carried out simultaneously to reduce the risk that a pair donates a kidney without also receiving one. Cycles are usually limited to at most three pairs due to logistical constraints. The second type, called a *chain*, is initiated when an altruistic donor donates to a patient in an incompatible pair. The donor from this pair can then continue the chain by donating to the next pair, and so on, until the chain terminates with an unpaired patient. Chains can be very long in principle because transplants do not have to be performed simultaneously, easing medical logistics.[7] However, our data from the National Kidney Registry (NKR), the largest US kidney exchange platform, show that most chains involve four to five transplants. Initially, cycles were the most common type of transaction, but chains became more important over time and now account for about 90 percent of transplants.

There are two types of biological compatibility constraints on kidney transplants: *blood-type* and *tissue-type* compatibility (Danovitch 2009). A donor is blood-type incompatible with a patient if the donor has a blood antigen that the patient lacks. There are two blood antigens, known as A and B. Blood type is A or B if the blood has only the A or the B antigen, respectively, AB if it has both, and O if it has neither. A donor is tissue-type incompatible with a patient if the donor has human leukocyte antigens (HLA) to which the patient has an immune response.[8] The most common measure of *sensitization*, that is, how likely a patient is to reject a transplant due to tissue-type incompatibility, is the *Panel Reactive Antibody* (PRA) score. A patient's PRA is between 0 and 100 and denotes the percentage of a representative population of donors with whom a patient is tissue-type incompatible. Because this measure depends on the choice of representative population, the NKR's algorithm uses an alternative measure tailored to its own pool called *match power*. It measures, for a given recipient (donor), the fraction of donors (recipients) on the platform that are both blood-type and tissue-type compatible.

[7] Rees et al. (2009) reports on an early chain involving ten transplants conducted over the course of eight months. One reason chains can be executed over a long period of time is that donors rarely renege (Cowan et al. 2017). This trust allows chains in which a patient can receive a transplant before her related donor donates to the next patient in the chain.

[8] Each patient has a list of antibodies to some, possibly large, subset of HLA antigens. If the recipient has an antibody to one of the donor kidney's antigens, the recipient's immune system will attack the kidney, leading to immediate rejection. A recipient is tissue-type compatible with a donor's kidney if she has no antibodies corresponding to the major HLA antigens of the donor's kidney (Danovitch 2009). Note that recent developments in desensitization techniques have allowed some of these incompatibilities to be overcome (Orandi et al. 2014).

### B. *Institutional Features and the Economics of Kidney Exchange*

There are three institutional features that are crucial for the economics of kidney exchange. First, kidney exchange takes place both in large, national platforms and within individual hospitals. There are three major national platforms currently operating in the United States: the National Kidney Registry (NKR), which is the largest; the Alliance for Paired Kidney Donation (APD); and the United Network for Organ Sharing (UNOS) Kidney Paired Donation Pilot Program. These large platforms match patients using optimization software that maximizes a weighted number of transplants. They differ in terms of exact algorithms and operational details.[9] Once an exchange has been arranged, standard medical procedures are used to conduct the transplant with no known differences across the platforms. Platforms also do not influence a patient or donor's surgeon choice. Besides these major platforms, there are small regional platforms and individual hospitals that also organize kidney exchanges.

Hospitals are not forced to participate in platforms. In fact, most hospitals that participate in large national platforms also match patients outside those platforms. When multi-hospital kidney exchange platforms were conceived, hospital participation was assumed to be all or nothing. As the market developed, it became clear that platforms effectively reward hospitals with transplants in order to receive submissions. For this reason, most platforms explicitly reward hospitals that submit altruistic donors by matching one of their unpaired patients.[10]

The second important institutional feature is that biological compatibility creates substantial variation in the social value of different submissions. One reason for this variation is blood-type compatibility. To simplify exposition, assume that there are only two blood types, O and A. These two types together are a significant majority of patients and donors in the United States. Denote a pair with patient blood type X and donor blood type Y as X–Y, and let $q_{\text{X–Y}}$ be the number of such pairs in a pool. Assume that $q_{\text{A–O}} < q_{\text{O–A}}$, which is the empirically relevant case.[11] For this simplified case, Roth, Sönmez, and Ünver (2007) showed that, in the large market limit, the number of transplants that can be performed, $f(\boldsymbol{q})$, is

$$f(\boldsymbol{q}) = 2 \cdot q_{\text{A–O}} + 1 \cdot (q_{\text{A–A}} + q_{\text{O–O}}) + 0 \cdot q_{\text{O–A}}. \tag{1}$$

This result follows because A–A and O–O pairs can be matched with pairs of the same type. Roth, Sönmez, and Ünver (2007) calls these pairs *self-demanded*. Self-demanded pairs have a marginal product of 1, in the sense that they generate 1 additional transplant when they join the pool. However, an O–A pair can only be transplanted using a cycle with one of the valuable A–O pairs. Thus, there will be many leftover O–A pairs that can only be transplanted if more A–O pairs

[9] See Abraham, Blum, and Sandholm (2007); Ashlagi et al. (2019); Anderson et al. (2014); Dickerson, Procaccia, and Sandholm (2012); and Agarwal et al. (2018).

[10] Until recently, matching an unpaired patient in return for submitting an altruistic donor has been the only form of reward used by these platforms. After the end of our sample, the NKR started experimenting with a more complex rewards system to encourage the registration of easy-to-match patients and donors. To our knowledge, it is not based on the marginal products.

[11] This fact is confirmed for patients and donors registered in the NKR. See Table 2.

join the pool. A–O pairs are called *over-demanded* and have a marginal product of 2. O–A pairs are called *under-demanded* and have a marginal product of 0. An under-demanded pair competes with another under-demanded pair and adds no value to the pool. Roth, Sönmez, and Ünver (2007) showed that this qualitative pattern holds even in a model with all possible blood types.

Current platform rules largely ignore this variation in the social value of submissions, inducing hospitals to perform socially inefficient matches. Consider a hospital with two over-demanded pairs. The hospital could perform a pairwise exchange to conduct two transplants. However, if the hospital submits both pairs to the platform, then in expectation, the hospital receives a number of transplants equal to twice the probability that one of them is matched. According to our data, this probability is 0.8, so the hospital expects only 1.6 transplants from submitting, which pushes it to match its patients outside the platform. However, each pair the hospital submits to the platform generates its marginal product, which the Roth, Sönmez, and Ünver (2007) model puts at 2. This suggests that the platform could generate four transplants if the hospital would submit both its pairs. Using a more realistic empirical model, we estimate just under three additional transplants (Section IV). Either way, matching these two pairs within the hospital is socially inefficient despite the hospital acting in the best interest of its patients.

An important corollary of Roth, Sönmez, and Ünver's (2007) results is that transplants are a natural numéraire in a kidney exchange platform. Because hospitals have a large number of under-demanded pairs, it is easy for a platform to transfer transplants from one hospital to another without compromising efficiency, simply by choosing which under-demanded pairs to match.

The third important institutional feature is that hospitals do not necessarily maximize a utilitarian measure of the welfare of the patients and third-party payers who they represent. We refer to such behavior as a broadly defined agency problem, since hospitals incur most of the transaction costs of kidney exchange. The social value from one transplant is more than $1,000,000, of which the majority is savings in health care costs and gains in quality-adjusted life years. But, hospital revenues are between $100,000 to $160,000 per transplant.[12] Variable profits are likely much smaller. Thus, even socially insignificant transaction costs of performing kidney exchange through a platform can be important for hospitals. Conversations with hospital staff indicate that participation in kidney exchange platforms involves logistical and administrative hassle in addition to direct costs arising from biological testing and platform fees.[13] Previous surveys and interviews have found that these transaction costs are commonly cited barriers to participation (Ellison 2014, American Society of Transplant Surgeons 2016). Besides costs, hospitals may also have behavioral reasons for not perfectly maximizing patient welfare. For example, there is

[12] See Held et al. (2016) and United States Renal Data System (2013). The revenues include payments for surgery teams, drugs, equipment, and capital.

[13] Platforms require extensive biological testing, which is particularly complicated because donors and patients are in different hospitals. Platforms also charge fees, which are paid by hospitals. NKR charges an annual fee of about $10,000 plus about $4,000 per transplant. See National Kidney Registry (2016) for NKR's fees, and Rees et al. (2012) and Wall, Veale, and Melcher (2017) for a broader discussion of the kidney exchange costs borne by hospitals.

considerable heterogeneity regarding hospital sophistication: some hospitals use optimization software to match patients while others do so manually.

### C. *Data*

We assembled two datasets for this paper. The first, the transplant dataset, records all kidney exchange transplants in the United States. We use this dataset to document fragmentation, inefficiency, and participation in the kidney exchange market. The second, the NKR dataset, records all patients and donors who registered with the largest US kidney exchange platform, the NKR. We use this dataset to estimate a transplant production function.

The transplant dataset consists of anonymized records of every kidney transplant conducted in the United States between January 1, 2008 and December 4, 2014. We obtained this dataset from the Organ Procurement and Transplantation Network (OPTN), a contractor for the US Department of Health and Human Services.[14] The OPTN dataset includes each transplant's date and location; whether it is part of a kidney exchange; the age, sex, weight, height, body mass index (BMI), blood type, and HLA antigens of the donor and recipient; and the unacceptable antigens and days on dialysis of the recipient. See online Appendix C for details.

Although a comprehensive source for data on transplants, the only field in the OPTN dataset that specifically pertains to kidney exchange is an indicator for which transplants were part of such an exchange. Therefore, the OPTN dataset does not identify which, if any, multi-hospital kidney exchange platform organized a given transplant.

To address this limitation, we separately obtained anonymized records of all transplants organized by each of the three largest multi-hospital kidney exchange platforms in the United States: NKR, APD, and UNOS. By merging the data from these platforms with the OPTN data, we identified which transplants were organized through NKR, APD, UNOS, or other avenues. This merge is not straightforward because all of our datasets are anonymized. Fortunately, the rich biological data allow us to match transplants across datasets on the blood type, sex, and HLA antigens of the recipient and donor; and the date and location of the transplant. See online Appendix C for more details. We were able to match approximately 94 percent of transplants at these platforms to the corresponding OPTN data with a high degree of certainty.[15]

The transplant dataset contains information on transplants that were performed, but not on the pool of patients and donors that were available for kidney exchange. This information is needed to estimate a platform's transplant production function. Therefore, we assembled the NKR dataset. It records all patients and donors who

[14] 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 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.

[15] Of the matches, 90 percent were within 1 day on the transplant date, within 5 years on donor and recipient age, and agreed on the hospital where the transplant was conducted as well as the blood type, sex, and all six major human leukocyte antigen (HLA) alleles relevant for kidney transplantation (2 alleles each at the HLA-A, B, and DR loci) for both the donor and recipient.

registered with the NKR between April 2, 2012 and December 4, 2014. This dataset is sourced from the administrative records the NKR uses to organize transplants. It includes the registration date, blood type, age, sex, and HLA antigens for both patients and donors. It also records whether the patient or donor left NKR's system, and the date and reason for departure (transplantation or otherwise). In addition, it includes information on pairings between donors and patients (if any), unacceptable antigens, and all the restrictions a patient places on which organs are acceptable. These fields allow us to determine the set of transplants the NKR considers acceptable and medically feasible. We also have detailed data on how the transplants were organized, including the donors and patients involved, and the chain or cycle configuration. Online Appendix C provides details on how we assembled the NKR dataset.

## II. Descriptive Evidence

We now document three key facts: the kidney exchange market is highly fragmented, this fragmentation leads to inefficiency, and there is evidence of agency problems.

### A. *Fragmentation*

We first document that the market is highly fragmented. Most kidney exchange transactions are matched internally by individual hospitals, as opposed to by large, national kidney exchange platforms. A kidney exchange transplant is defined as *within hospital* if the donor's operation took place in the same hospital as the patient's, and *across hospitals* if the donor's and patient's operations took place in different hospitals.[16] We also classify transplants based on which platform coordinated the exchange: NKR, APD, or UNOS. Transplants that were not organized by one of these platforms are classified as being performed by other platforms, including single-hospital programs and small regional platforms.

Figure 1 shows that the market is highly fragmented. The three largest multi-hospital platforms together only account for a minority share of the kidney exchange market. Of all kidney exchange transplants, 62 percent are within-hospital transplants that are not facilitated by the NKR, APD, or UNOS. Over 100 hospitals performed kidney exchanges outside these three platforms during this period.

Unlike the dominance of within-hospital exchanges in the overall market, a large majority of the transplants facilitated by multi-hospital platforms are across hospitals. This contrast between the overall market and the platforms is striking as the platforms do not prioritize across-hospital exchanges as a rule; such exchanges are a by-product of maximizing the total number of transplants. This suggests that coordinating across hospitals has potential gains.

[16] The common practice is to transport the organ after recovery instead of transporting the donor and recovering the organ elsewhere. Conversations with surgeons suggest that the primary motivation for this practice is to safeguard the donor's interests: she has built relationships at her hospital, and the donation surgery requires extensive pre-planning and follow-up care.

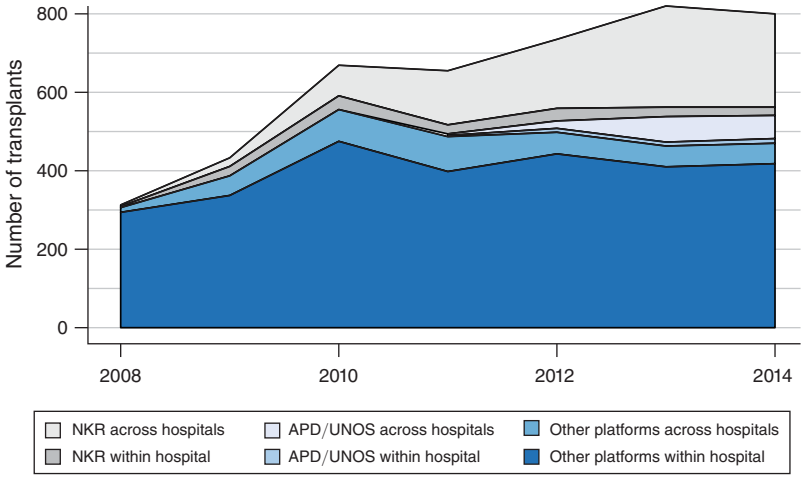


Figure 1. Market Fragmentation and Trends in Kidney Exchange

*Notes:* The figure displays the number of kidney exchange transplants in different categories. The category *Other* represents transplants that were not facilitated by NKR, APD, or UNOS. Single-hospital platforms fall under this category. *Within hospital* and *across hospital* classify a transplant into whether the donor's hospital was the same as the patient's hospital.

Figure 1 also shows that the total number of kidney exchange transplants grew from about 400 in 2008 to about 800 in 2014.[17] However, overall market growth seems to have slowed in recent years. The total number in 2017 remains at around 800,[18] well below some estimates of the potential size of the market (Bingaman et al. 2012, Massie et al. 2013).

The growth in kidney exchange between 2010 and 2014 is concurrent with the NKR becoming the dominant kidney exchange platform. The NKR accounted for 33.1 percent of all kidney exchange transplants in 2014 and facilitated more than 5 times as many transplants as the APD and UNOS combined.[19] The importance of the NKR during our sample period motivates our focus on the platform in the subsequent sections.

### B. *Evidence of Inefficiency*

Market fragmentation creates inefficiency if there are increasing returns to scale and hospitals are operating below efficient scale. We now present direct evidence of hospitals conducting exchanges that are inefficient from a social perspective.

One easily detectable inefficiency is a transplant between a blood-type O donor and a non-O patient. As explained in Roth, Sönmez, and Ünver (2007) and in Section I, O donors are scarce while O patients are abundant. If all transplants are of equal social value, optimal matches in a large market should only transplant organs

[17] Our data for the NKR extend until December 4, 2014. This censoring may account for the slight drop in transplants in the last year of this figure.

[18] Source: https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/ (accessed December 21, 2017).

[19] The APD has grown in recent years, significantly closing the gap.

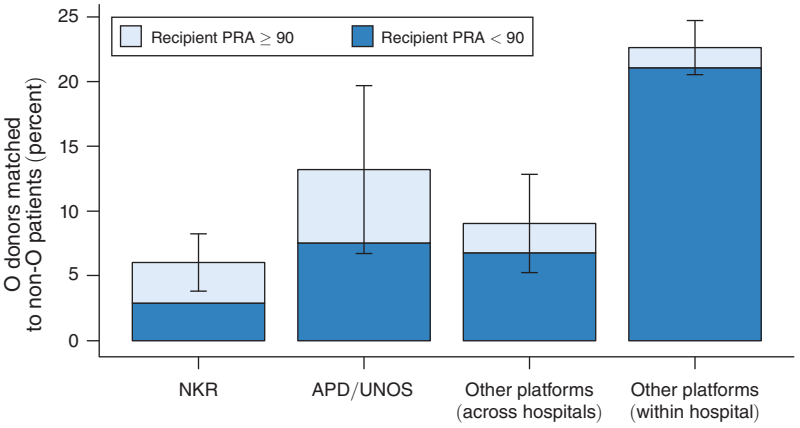


FIGURE 2. EVIDENCE OF HOSPITALS PERFORMING INEFFICIENT MATCHES

*Notes:* The bars display the percentage of O donors whose kidneys were transplanted into non-O patients for different categories of transplants. *Other* indicates a transplant not organized by NKR, APD, or UNOS. This category includes transplants organized by single-hospital platforms. *Within hospital* and *across hospital* classify a transplant into whether the donor hospital was the same as the patient hospital. The colors decompose this total into highly sensitized patients (PRA $>$ 90) and non-highly sensitized patients. The error bars depict 95 percent confidence intervals for the totals.

from O donors to O patients because O patients cannot accept other blood types.[20] The exception to this rule is for a highly sensitized patient, that is, one with a very high PRA. The platform might want to match an O donor with such a patient if it were the only way to get her transplanted.

Figure 2 displays the fraction of O donors that are used to transplant non-O patients, categorized into NKR transplants, APD/UNOS transplants, across-hospital transplants at other platforms, and within-hospital transplants at other platforms. Among NKR transplants, only 6.5 percent of O donors are used for non-O patients. In contrast, among within-hospital transplants outside the three platforms, this figure is 22.8 percent. The difference is statistically significant ($p < 0.01$) and constitutes strong evidence that hospitals often perform inefficient matches outside the platform. The figures for APD, UNOS, and across-hospital transplants at other platforms are in between these two categories, but much closer to the NKR.

An alternative explanation for inefficient matching is that within-hospital transplants use O donors to help highly sensitized patients who would otherwise remain untransplanted. However, Figure 2 shows that almost none of the potentially inefficient transplants in the *Other* (*within hospital*) category involve highly sensitized patients. In contrast, about one-half of the potentially inefficient NKR transplants involve highly sensitized patients.

This exercise treats the value of all transplants as equal, irrespective of patient or donor blood type and whether the exchange was organized through a platform.

[20] Strictly speaking, efficiency as discussed here means maximizing the total number of transplants. However, transplanting an O donor to a non-O patient is also likely to be Pareto inefficient: if a pairwise exchange between two over-demanded A–O pairs were replaced by two A–O to O–A exchanges, then all parties would be better off (assuming the under-demanded O–A pairs would otherwise be unmatched).

Table 1—Summary Statistics for Kidney Exchange Transplants

| | NKR | APD/UNOS | Other platforms | |
|---|---|---|---|---|
| | | | Across hospital | Within hospital |
| Observations | 1,118 | 198 | 341 | 2,719 |
| Patient blood type (%) | | | | |
| A | 34.7 | 36.4 | 37.2 | 37.1 |
| B | 19.0 | 21.2 | 17.6 | 17.0 |
| AB | 5.7 | 3.5 | 7.0 | 5.7 |
| O | 40.6 | 38.9 | 38.1 | 40.2 |
| Donor blood type (%) | | | | |
| A | 36.8 | 35.4 | 37.5 | 33.4 |
| B | 18.2 | 20.2 | 14.7 | 13.8 |
| AB | 3.9 | 1.5 | 6.7 | 2.9 |
| O | 41.1 | 42.9 | 41.1 | 49.9 |
| Panel reactive antibody (PRA) (sensitization) | | | | |
| Mean | 35.0 | 43.0 | 30.4 | 17.6 |
| Standard deviation | 39.7 | 40.8 | 37.5 | 30.8 |
| Percent >90 | 16.4 | 20.6 | 12.0 | 5.1 |
| Transplant outcomes and quality measures | | | | |
| Donor age | | | | |
| Mean | 44.1 | 44.6 | 44.1 | 43.2 |
| Standard deviation | 11.8 | 11.1 | 11.3 | 11.8 |
| Donor body mass index (BMI) | | | | |
| Mean | 26.5 | 27.0 | 26.6 | 26.5 |
| Standard deviation | 4.0 | 4.0 | 4.1 | 4.2 |
| Donor height (cm) | | | | |
| Mean | 169.4 | 168.0 | 169.6 | 169.3 |
| Standard deviation | 9.8 | 9.6 | 10.3 | 9.8 |
| Donor weight (kg) | | | | |
| Mean | 76.3 | 76.3 | 76.9 | 76.3 |
| Standard deviation | 15.1 | 13.9 | 15.4 | 15.1 |
| Tissue type mismatch (0–6) | | | | |
| Mean | 4.2 | 4.2 | 4.2 | 4.4 |
| Standard deviation | 1.3 | 1.4 | 1.2 | 1.2 |
| Mean days on dialysis | | | | |
| Mean | 1,026.6 | 1,048.4 | 1,063.1 | 969.1 |
| Standard deviation | 1,088.1 | 848.1 | 1,269.5 | 990.9 |

*Note:* Sample of all kidney exchange transplants between January 1, 2008 and December 4, 2014.

The total number of transplants provides a transparent and economically relevant measure of market outcomes. Although not reported, our data show that indicators of the life-years benefit of kidney exchange, such as patient or donor age, do not significantly differ by blood type.

The remaining concern is that transplants organized through a platform may differ in quality or costs. However, Section IA argues that transaction costs of conducting transplants through a platform are negligible relative to the value of transplants lost by matching O donors to non-O patients. Moreover, Table 1 shows that there are no substantial differences in donor or match quality across our platform types. One reason patients considering a multi-hospital platform need not worry about donor quality is that the platforms allow patients and doctors to specify donor acceptability criteria. They also allow patients to refuse proposed transplants if the donor is unsuitable. The only noticeable difference in Table 1 is that patients who receive

a transplant through a platform typically spend only two more months on dialysis than do patients who receive a within-hospital transplant. Given the average patient wait time of about 32 months, this difference represents only an 8 percent increase. The longer waiting time at the platforms should be expected because, as we discuss below, patients transplanted through the platform are, on average, harder to match.

If each of these inefficient transplants from O donors to non-O patients comes at the cost of one other transplant, as in the Roth, Sönmez, and Ünver (2007) model, then achieving the level of efficiency obtained by the NKR would have resulted in about 250 additional transplants between 2008 and 2014.[21] The advantage of considering only the clearly inefficient transplants is that the results provide transparent evidence of inefficiency. The total inefficiency can be much larger.

### C. *Hospital Participation and Evidence of Agency Problems*

Why do hospitals fail to participate more in national platforms? We start by documenting key facts about hospital behavior and argue that hospitals do not purely maximize the number of transplanted patients. Instead, they seem to maximize complex and heterogeneous objectives, including, but not limited to, profits and patient welfare.

*Descriptive Evidence*.—We focus on participation behavior at the NKR because it is the primary multi-hospital kidney exchange platform during our sample period (Table 1). Panel A of Figure 3 depicts the extensive margin of participation among hospitals conducting kidney exchange transplants. A hospital is considered an NKR participant if it conducted at least one transplant through the NKR during our sample period. The figure is a binned scatterplot of the fraction of hospitals that participate in the NKR versus hospital size (measured by the total number of kidney transplants performed, both living and deceased).[22] Panel B depicts the intensive margin of participation. The vertical axis in this scatterplot is the fraction of kidney exchange transplants that a hospital performs through the NKR. The results are qualitatively similar if we consider participation at any of the three largest kidney exchange platforms because the APD and UNOS are relatively small during our sample period.

The figures reveal four key facts about participation. First, both the extensive and intensive margins are important drivers of market fragmentation. Only 41.4 percent of hospitals participate in the NKR. Within those participating hospitals, only 59.1 percent of transplants are conducted through the NKR. These results are qualitatively similar if participation in any of the three national platforms (NKR, APD, or UNOS) is considered because only a few hospitals participate in multiple platforms.[23] Second, larger hospitals are considerably more likely to participate in

[21] Table 1 shows that within-hospital platforms have a larger gap between the fraction of O donors and patients than the NKR. The difference in this gap, multiplied by the number of within-hospital transplants, is a measure of transplants lost due to inefficient use of O donors in within-hospital transplants.

[22] This broad measure of size limits the endogenous effect of NKR participation on hospital size since deceased donor and direct living-donor transplants form the bulk of kidney transplants. During our sample period, the total number of kidney transplants has remained stable relative to the growth in kidney exchange.

[23] Only 10 hospitals out of the 64 that participated in the NKR between 2012 and 2014 also conducted a transplant through the APD or UNOS. Most of these hospitals did not do so in all three years and conducted only a handful of transplants through the other platforms.

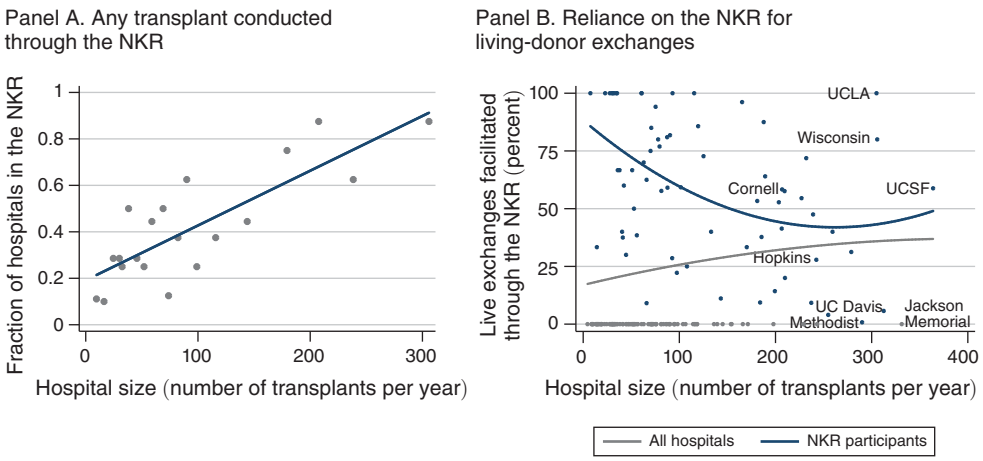


FIGURE 3. HETEROGENEITY IN PARTICIPATION IN THE NKR

the NKR. The probability of participating at all is about 80 percent for a hospital that performs approximately 250 transplants per year but only about 30 percent for a hospital that performs about 50 transplants per year (panel A of Figure 3). Third, conditional on participating, large hospitals conduct more of their matches outside the platform (panel B). Although size positively correlates with the fraction of kidney exchange transplants performed in the NKR, the relationship is negative if we focus exclusively on hospitals that participate at all (panel B). Fourth, there is a high degree of heterogeneity in intensive margin participation. Even among hospitals with similar size, participation varies considerably (panel B). For example, among the five transplant hospitals that perform more than 300 transplants per year, one does not participate at all (Jackson Memorial), one has a participation rate close to zero (UC Davis Medical Center), one has a rate in the 50–60 percent range (UCSF Medical Center), and two have rates greater than 75 percent (UCLA Medical Center and the University of Wisconsin Hospital).

The data also provide information on the characteristics of patients submitted to the NKR and of the patients transplanted by each hospital, categorized by how the transplant was facilitated. Tables 1 and 2 reveal three main facts.

First, the NKR receives submissions that are very hard to match compared to the general population (Table 2). The blood types of both altruistic and paired donors skew away from O donors and toward A donors relative to the US population. The deceased donor population has about 45 percent O donors and 40 percent A donors. In contrast, patients in pairs are disproportionately likely to have blood type O (58.6 percent), and their related donors are unlikely to have blood type O (31.9 percent). Only a small fraction of pairs (13.8 percent) are over-demanded. Interestingly, unpaired patients are much more likely to have an easy-to-match blood type (the majority having A). The average PRA for patients registered with the NKR is 47.6, which corresponds to tissue-type incompatibility with almost one-half of the reference population.

Second, the NKR transplants patients who are considerably harder to match than patients transplanted by single hospitals (Table 1). Approximately 40 percent of the patients and donors transplanted through the NKR are blood type O. The PRA of

TABLE 2—SUMMARY STATISTICS FOR NKR SUBMISSIONS

| | Altruistic donors | Pairs | Unpaired patients |
|---|---|---|---|
| Observations | 164 | 1,265 | 501 |
| Patient blood type (%) | | | |
| A | — | 23.8 | 51.1 |
| B | — | 15.0 | 16.0 |
| AB | — | 2.6 | 19.0 |
| O | — | 58.6 | 14.0 |
| Donor blood type (%) | | | |
| A | 44.5 | 44.4 | — |
| B | 14.0 | 18.5 | — |
| AB | 3.7 | 5.2 | — |
| O | 37.8 | 31.9 | — |
| Match power | | | |
| Recipient/pair | | | |
| Mean | — | 0.218 | 0.431 |
| Standard deviation | — | 0.210 | 0.392 |
| Donor | | | |
| Mean | 0.279 | 0.258 | — |
| Standard deviation | 0.162 | 0.159 | — |
| Panel reactive antibody (PRA) (sensitization) | | | |
| Mean | — | 48.8 | 44.4 |
| Standard deviation | — | 41.1 | 45.1 |
| Pair type (%) | | | |
| Over-demanded | — | 13.8 | — |
| Under-demanded | — | 41.9 | — |

*Notes:* A pair is over-demanded if the patient is blood-type compatible with the related donor, but not of the same blood type. Under-demanded pairs are either blood-type O patients without blood-type O donors or are blood type-AB donors without blood-type AB patients. Sample is all patients and donors registered in the NKR between April 2, 2012 and December 4, 2014.

the patients transplanted through the NKR is approximately 35, and about one in six patients have a PRA above 90. These statistics are similar for across-hospital kidney exchanges not facilitated by the NKR and transplants facilitated by APD or UNOS. In contrast, among within-hospital kidney exchanges not conducted by a large platform, almost 50 percent of the donors are blood type O, but only 40 percent of the patients. The average PRA of patients transplanted through within-hospital exchanges is only 18. This is almost one-half of the mean PRA for patients transplanted through one of the three national platforms.

Third, transplants on all platforms look similar in donor quality measures that do not affect compatibility, such as weight, body mass index, and age (Table 1). This supports our equal treatment of all transplants for welfare calculations, irrespective of whether they are facilitated through a national platform.

*Implications for Hospital Behavior.*—The facts above have implications for different hypotheses about hospital behavior. In the discussion that follows, we approximate total patient welfare with the total number of transplants because, as we argued in Section I, kidney exchange transaction costs are small relative to the benefits of transplantation.

The first hypothesis is that hospitals maximize the total welfare of all patients in the system, regardless of the hospital to which a patient belongs. This hypothesis is strongly rejected by several features of the data, such as the evidence of socially inefficient matches (Figure 2).

A second hypothesis is that hospitals only maximize the welfare of their own patients. This hypothesis was investigated theoretically by Ashlagi and Roth (2014), which argues that hospitals will try to match as many of their patients internally as possible and only submit the remaining patients to a multi-hospital kidney exchange platform. This hypothesis fits some qualitative patterns in the data, but not others. For example, it explains why larger hospitals in the NKR perform fewer transplants through the platform: these hospitals have more opportunities to match patients outside the platform (panel B of Figure 3). However, it does not explain why many hospitals do not participate in a national platform at all, even though all hospitals likely have patients who cannot be matched.[24] Moreover, many small hospitals do not participate in the NKR, even though these hospitals, due to their size, are least likely to find matches within the hospital. The patterns suggest that hospitals respond to fixed costs of participating in kidney exchange platforms, even though these transaction costs are small relative to benefits to patients and cost savings from dialysis to insurers.

A third hypothesis is that hospitals are profit maximizers. This hypothesis is consistent with the fact that small hospitals are less likely to participate in the NKR (panel A of Figure 3) because the fixed costs of participation may not compensate for the gains in profits from additional transplants. However, this theory alone cannot fully explain the large variation in the degree of participation, especially among large hospitals. For example, transplant coordinators at Cornell Medical Center (a large hospital with a high NKR participation rate) report that a primary reason for participating is the importance of contributing to a national kidney exchange platform (Ellison 2014).

Taken together, the evidence on hospital participation suggests that hospitals maximize complex and heterogeneous objectives. This finding is consistent with the anecdotal evidence on kidney exchange reviewed in Section I, as well as the standard view in health care economics (Arrow 1963) and more recent findings about the behavior of health care providers (Kolstad 2013, Clemens and Gottlieb 2014).

The facts about selection into which patients and donors are submitted to the NKR also indicate that two of these theories, maximizing profits and maximizing their own patients' welfare, can explain many hospitals' behavior. These theories' shared implication is that pairs submitted to national platforms are negatively selected, in the sense of being hard to match. In both cases, a hospital only submits a pair to a platform if an internal match is not possible. Unfortunately, we cannot directly test this prediction because we do not have data on the entire pool of patients available to individual hospitals. But, it is reassuring that the results on selection do not falsify the two theories that best fit the participation behavior.

To summarize, these findings have two important implications. First, there is clear evidence of agency problems, as we defined broadly in Section I. Second, the data indicate that none of the simple models describe the behavior of all hospitals.

## III. Theory

The evidence above shows that kidney exchange markets are fragmented, which leads to real efficiency loss. We now build a model similar to one of a traditional

[24] Recall that over-demanded pairs are typically scarce. We will see in Section IV that even the NKR is able to match only about 50 percent of its donors.

market in which the platform procures submissions (donors and patients) from hospitals and rewards these hospitals with transplants. We use the model to explain how inefficiency arises, to quantify it, and to develop responses. In particular, under certain assumptions, policy responses like participation mandates and providing incentives can improve outcomes.

## A. *Model*

A kidney exchange platform procures submissions from hospitals and rewards them with transplants. The platform's ability to produce transplants is described by a *production function f*. We consider *types of submissions* $i = 1, \ldots, I$. A (column) vector of *submissions* $\mathbf{q} = (q_i)_{i=1}^{I}$ in $\mathbb{R}_+^I$ specifies a quantity $q_i$ of each submission type available to the platform, where $\mathbb{R}_+$ is the set of non-negative real numbers. Given a vector of submissions $\mathbf{q}$, the platform can produce $f(\mathbf{q})$ transplants. The model can be interpreted as either static or as a steady state from a dynamic model. We will use the steady-state interpretation in the empirical analysis. All variables are measured in flows (i.e., transplants per year).

The production function $f(\mathbf{q})$ summarizes what matches are possible. Roth, Sönmez, and Ünver (2007) calculated the production function using a simple model that we described in Section I. Since that paper assumed that all submissions are pairs and that only blood-type compatibility matters, its model has $I = 16$ types. Our analysis applies both to such theoretically tractable production functions as well as to more complex ones. Section IV uses an empirical production function that allows submissions to differ by whether they are patient–donor pairs, altruistic donors, or unpaired patients, and by a host of variables including blood type and antigen and antibody profiles. Thus, the number of types $I$ is potentially large.

We say that the production function $f$ has *constant returns to scale at* $q$ if its elasticity with respect to scale at $\mathbf{q}$ is equal to 1. That is, $\left((\alpha / f(\alpha \mathbf{q})) \cdot (\partial f(\alpha \mathbf{q}) / \partial \alpha)\right)\big|_{\alpha=1} = 1$, which is equivalent to $\nabla \mathbf{f}(\mathbf{q}) \cdot \mathbf{q} = f(\mathbf{q})$. The Roth, Sönmez, and Ünver (2007) model considers a large platform with constant returns to scale. Our empirical production function in Section IV will measure the returns to scale for the NKR.

The platform produces transplants using submissions provided by *hospitals* $h = 1, \ldots, H$. Hospitals are rewarded for these submissions with transplants. We assume these rewards are linear in submissions and anonymous. That is, there exists a (row) *vector of rewards* $\mathbf{p} = (p_i)_{i=1}^{I}$ in $\mathbb{R}^I$ where the i-th component denotes the (expected) number of transplants awarded to the hospital per submission of type $i$. The units of $p_i$ are transplants per submission. A hospital that submits a flow $\mathbf{q}^h$ in $\mathbb{R}_+^I$ of submissions receives a flow $\mathbf{p} \cdot \mathbf{q}^h$ of transplants, where $\cdot$ represents matrix multiplication. Since all transplants that are performed must be allocated to some hospital, a platform must satisfy the constraint that $f\left(\sum_h \mathbf{q}^h\right) = \sum_h \mathbf{p} \cdot \mathbf{q}^h$.

This linear reward schedule is a good approximation of current platforms' rules because their matching algorithms maximize a weighted sum of the number of matches without considering which patients and donors are submitted by each hospital (Sönmez and Ünver 2013, Anderson et al. 2015). That is, when a hospital submits an additional pair, the probability that the platform matches a different pair

from the same hospital does not significantly change. Therefore, the current reward for submitting a type $i$ pair is equal to the probability $p_i$ that the pair is matched.

We assume that hospital utility equals the number of transplants received from the platform minus a *private cost* of submissions, $C^h(\mathbf{q}^h)$, measured in transplant units. For instance, if a hospital maximizes the number of its own patients that are transplanted, then $C^h(\mathbf{q}^h)$ is the number of within-hospital transplants that the hospital must forgo in order to submit $\mathbf{q}^h$. However, $C^h(\mathbf{q}^h)$ could also include a weight on additional hospital profits from within-hospital transplants. Our analysis will remain agnostic about the specific form of $C^h(\mathbf{q}^h)$, except when explicitly noted.

Welfare is defined over an *allocation* $(\mathbf{q}^h)_{h=1}^H$ that specifies the quantity of pairs supplied by each hospital. We will use two welfare notions, both of which use transplants as a numéraire because platforms can effectively transfer transplants between hospitals by choosing which under-demanded submissions to match (see Section I).

The first notion is *hospital welfare* $W^H(\mathbf{q}^1, \ldots, \mathbf{q}^H)$, which is the total welfare measured from the point of view of hospitals. Hospital welfare equals the total number of transplants produced (which is the same number of transplants that hospitals receive) minus the private costs. That is,

$$W^H(\mathbf{q}^1, \ldots, \mathbf{q}^H) = f\left(\sum_h \mathbf{q}^h\right) - \sum_{h=1}^H C^h(\mathbf{q}^h). \tag{2}$$

This is a compelling notion of welfare if the goal is to help key market participants (hospitals, in this case) achieve their objectives.

Hospital welfare is not compelling if hospitals do not purely maximize patient and insurer welfare. As discussed in Sections IA and II, there is anecdotal and empirical evidence of such agency problems. For this reason, we also consider a utilitarian welfare measure, which we term social welfare.[25]

Define $SC^h(\mathbf{q}^h)$ as the *social cost* for hospital $h$ to supply a vector $\mathbf{q}^h$ submissions. If there are agency problems, then social and private costs are different, and there is an agency *externality* from hospital $h$'s submissions because $C^h(\mathbf{q}^h) \neq SC^h(\mathbf{q}^h)$. For example, $C^h(\mathbf{q}^h)$ is larger than $SC^h(\mathbf{q}^h)$ if hospital $h$ acts as though the financial and logistical costs of participating in kidney exchange platforms are significant relative to the private value of a transplant. The externality represents the benefits to stakeholders other than the hospital itself. We will refer to any such wedge as a (broadly defined) agency problem.[26] In the particular case where there are no agency problems, we have $C^h(\mathbf{q}^h) = SC^h(\mathbf{q}^h)$ for all $h$. Define *social welfare* to be

$$SW(\mathbf{q}^1, \ldots, \mathbf{q}^H) = f(\mathbf{q}) - \sum_{h=1}^H SC^h(\mathbf{q}^h).$$

[25] The theory does not make specific assumptions about social welfare. In the empirical application, we will consider social welfare to be equal to the number of transplants performed because the costs of organizing exchanges and transplants are small relative to the social benefit of a transplant.

[26] The wedge between $C^h$ and $SC^h$ includes all reasons why hospitals' objectives deviate from social goals. But, it can be decomposed into behavioral reasons why hospitals do not maximize social welfare (such as inattention) and differences in payoffs if hospitals were rational. We follow the behavioral public finance literature (Farhi and Gabaix 2017) in letting the wedge include all of these differences. We use the term agency for two reasons. First, Section I argues that hospitals represent patients and insurers, and that the additional costs of participating in kidney exchange are negligible when compared to the value of a transplant. Second, survey and anecdotal evidence suggests that these costs influence decisions, indicating that agency problems are important. We cannot rule out behavioral biases resulting in suboptimal hospital behavior.

Define first-best hospital welfare as the supremum of $W^H$ and first-best social welfare as the supremum of $SW$.

Given these primitives, for a vector of rewards $\mathbf{p}$, the *hospital supply* of hospital $h$ is

$$\mathbf{S}^h(\mathbf{p}) \;=\; \underset{\mathbf{q}^h \in \mathbb{R}^I_+}{\operatorname{argmax}} \left[\mathbf{p} \cdot \mathbf{q}^h - C^h\left(\mathbf{q}^h\right)\right].$$

Define the *aggregate cost*, $C(\mathbf{q})$, to be the minimum sum of hospital private costs necessary to ensure that hospitals supply $\mathbf{q} \equiv \sum_h \mathbf{q}^h$ in aggregate. Let the *aggregate supply* correspondence be

$$\mathbf{S}(\mathbf{p}) \;=\; \underset{\mathbf{q} \in \mathbb{R}^I_+}{\operatorname{argmax}} \left[\mathbf{p} \cdot \mathbf{q} - C(\mathbf{q})\right].$$

We assume that the production function, social and private costs, and aggregate cost functions are defined over all non-negative real vectors and are smooth. The maximum of each hospital's objective is attained for some quantity for every vector of rewards. Further, assume that aggregate cost is strictly convex.

Online Appendix B shows that aggregating individual hospital supplies yields $\mathbf{S}(\mathbf{p})$. Denote the *aggregate inverse supply* with $\mathbf{P}_S(\mathbf{q}) \;=\; \left\{\mathbf{p} \in \mathbb{R}^I \mid \mathbf{q} \in \mathbf{S}(\mathbf{p})\right\}$. Further, online Appendix B shows that, for strictly positive $\mathbf{q}$, the aggregate inverse supply is single-valued and $\mathbf{P}_S(\mathbf{q}) \;=\; \nabla \mathbf{C}(\mathbf{q})$, where we use the convention that gradients are row vectors. This result is similar to how firms supply at price equal to marginal cost in a competitive market.

### B. *Illustrative Example: Agency and the Wedge between Private and Social Costs*

Our model of the kidney exchange market is framed in terms of transplants as a numéraire, and captures agency problems as a wedge between private and social costs. We now present a particular example to clarify these two features of the model. The specific assumptions in this section are not necessary for our results.

Let $K^h\left(\mathbf{q}^h\right)$ be the monetary costs borne by hospital $h$ of sending $\mathbf{q}^h$ submissions to a kidney exchange platform. These can include platform fees, costs of rearranging the hospital's schedule around the platform, and funds for hiring additional transplant coordinators (see Section IA). Let $T^h\left(\mathbf{q}^h\right)$ be the flow of kidney exchange transplants that hospital $h$ forgoes when submitting $\mathbf{q}^h$ to the platform because it cannot match these patients and donors internally.[27]

To combine the monetary costs and the transplant costs of submitting, we need a rate of exchange between the two. Let hospitals value each transplant at $v$ dollars, which includes profits and the value that hospitals place on transplanting their patients. Gross revenues from a transplant are approximately \$150,000 (United States Renal Data System 2013, Held et al. 2016). For illustrative purposes, take $v$ to be \$50,000, which represents a generous 50 percent markup on costs. In transplant units, hospital $h$'s cost function is

$$C^h\left(\mathbf{q}^h\right) \;=\; T^h\left(\mathbf{q}^h\right) + \frac{K^h\left(\mathbf{q}^h\right)}{v}.$$

[27] If the hospital cannot supply a quantity $\mathbf{q}^h$, define $T^h\left(\mathbf{q}^h\right)$ as infinity.

The private value of a transplant just discussed does not account for any benefits that fall to non-hospital stakeholders. Such benefits include the value a patient has for a transplant beyond the value the hospital places on it and the savings in health care costs to insurers. Hospitals contract with these agents, but may not account for all of their benefits because they are not incentivized to do so. For this reason, the social value of a transplant may differ from the private value to a hospital, creating agency problems.

Let society value transplants at $V > v$ dollars. The cost-benefit analysis in Held et al. (2016) places $V$ at \$1.1 million.[28] But, even one-half of that value is much higher than hospital profits. This fits our model with social costs

$$SC^h(\mathbf{q}^h) = T^h(\mathbf{q}^h) + \frac{K^h(\mathbf{q}^h)}{V}.$$

Hence, the wedge between private and social costs equals

$$C^h(\mathbf{q}^h) - SC^h(\mathbf{q}^h) = \left(\frac{1}{v} - \frac{1}{V}\right) \cdot K(\mathbf{q}^h).$$

The difference is how much more hospitals care about the costs of participating in a kidney exchange platform than society does, measured in transplant units.

To develop intuition for this wedge's magnitude, assume that the monetary cost is linear in the number of submissions, i.e., $K^h(\mathbf{q}^h) = k\sum_{i=1}^{I} q_i^h$. Then, the wedge is roughly $(k/v)\sum_{i=1}^{I} q_i^h$, where the approximation holds because the social value of a transplant $V$ is much larger than the monetary cost $k$. The wedge is large because it depends on the platform participation costs borne by the hospitals as a fraction of a transplant's private value, not its social value.

For example, if $k$ is \$10,000 and $v$ is \$50,000, then the wedge is $k/v = 0.20$ transplants per submission. Hospitals compare this wedge to the rewards vector $\mathbf{p}$, which is equal to the probability of matching various submissions in the current mechanism. In effect, the wedge creates an incentive for the hospital to not submit a patient or donor to a national platform. The calculation above suggests that, because of agency problems, rewards have to be 20 percentage points higher in order to induce a given submission. Therefore, it is likely that agency problems are an important part of the kidney exchange market.

### C. *Optimal Incentives*

The following theorem describes optimal rewards.

THEOREM 1 (Optimal Rewards): *Consider a vector of rewards* $\mathbf{p}$ *and an allocation* $(\mathbf{q}^h)_{h=1}^H$ *with strictly positive aggregate quantity* $\mathbf{q}$ *that maximizes hospital welfare subject to all hospitals choosing* $\mathbf{q}^h \in \mathbf{S}^h(\mathbf{p})$ *and subject to the total*

[28] Even if a kidney exchange patient would have otherwise received an organ from a deceased donor, the kidney exchange enables some other patient to receive the deceased donor organ. Therefore, the social benefit of each kidney exchange transplant should still be the same as the gain from one transplant.

*rewards allocated being the same as the number of transplants produced, that is,* $f(\mathbf{q}) = \mathbf{p} \cdot \mathbf{q}$. *Then:*

(*i*) *The platform rewards each type of submission with its marginal product minus an adjustment term. Specifically,*

$$\mathbf{p} = \nabla \mathbf{f}(\mathbf{q}) - \mathbf{A}(\mathbf{q}),$$

*where*

$$\mathbf{A}(\mathbf{q}) = \left( \frac{\nabla \mathbf{f}(\mathbf{q}) \cdot \mathbf{q} - f(\mathbf{q})}{\mathbf{q}' \cdot \mathbf{DP}_S(\mathbf{q}) \cdot \mathbf{q}} \right) \mathbf{q}' \cdot \mathbf{DP}_S(\mathbf{q}),$$

*and* $\mathbf{DP}_S(\mathbf{q})$ *is the Jacobian matrix of the inverse supply.*

(*ii*) *If the production function has constant returns to scale at* $\mathbf{q}$, *then the reward for each type of submission is equal to its marginal product,* $\mathbf{p} = \nabla \mathbf{f}(\mathbf{q})$. *Moreover, the allocation* $\left(\mathbf{q}^h\right)_{h=1}^H$ *attains first-best hospital welfare.*

(*iii*) *If, in addition, social cost is equal to private cost* ($C^h(q^h) = SC^h(q^h)$ *for all* $h$), *then this allocation attains first-best social welfare.*

The first part shows that to maximize hospital welfare, the reward for each submission must be approximately equal to its marginal product. To build intuition, ignore the constraint that the platform cannot allocate more transplants than it produces. The platform is similar to a firm that produces a consumption good (transplants) using intermediate goods (submissions). Since the supply of intermediate goods equates prices $\mathbf{p}$ with marginal costs $\nabla \mathbf{C}$, the first-order condition for maximizing hospital welfare implies that efficient prices $\mathbf{p}$ must equal to marginal products $\nabla \mathbf{f}$.

The only complication is the constraint that a platform cannot allocate more rewards than the number of transplants produced, which binds if $f$ does not exhibit constant returns to scale. If $f$ exhibits increasing returns to scale, then the average product is lower than the marginal product. In this case, $\mathbf{A}(\mathbf{q}) \cdot \mathbf{q}$, the optimal level of shading relative to marginal products, summed over all flows, equals $\nabla \mathbf{f}(\mathbf{q}) \cdot \mathbf{q} - f(\mathbf{q})$. Thus, the optimal level of shading depends on the returns to scale. The formula for $\mathbf{A}(\mathbf{q})$ shows that the platform should shade more aggressively on submissions with less elastic supply.[29] Our formula is similar to standard formulas in optimal linear commodity taxation (Ramsey 1927), regulation of monopolies (Boiteux 1956), and optimal pricing (Lerner 1934). The proofs are identical for kidney exchange platforms, but provide new insight in an application where monetary prices are not paid to acquire submissions.

[29] To see why, consider the case when the cross-elasticities of supply are zero so that $\mathbf{DP}_S$ is a diagonal matrix. Then, for each type $i$, the reward is marked down from marginal product according to $(1/p_i)\left((\partial f/\partial q_i(\mathbf{q})) - p_i\right) = \lambda/\varepsilon_i$, where $\varepsilon_i$ is the own-price supply elasticity and $\lambda$ is the Lagrange multiplier on the constraint that all transplants produced must be given out as rewards, that is, $f(\mathbf{q}) = \mathbf{p} \cdot \mathbf{q}$.

The theorem shows that current platform rules are inefficient. Instead of rewarding submissions with their marginal products, current rules reward submissions with the probability of being transplanted. Therefore, there is a wedge between the social and private benefits of submissions. Under current rules, a hospital chooses between serving its own patients or providing a service to the system as a whole. A clear example of this dilemma, described in Section IB, is of a hospital with two over-demanded pairs. This hospital could match the pairs internally instead of submitting them to a platform, but doing so would cause the type of inefficiency documented in Section II.

The second part of the theorem shows that, when returns to scale are constant, the optimal mechanism rewards submissions exactly according to marginal products. The adjustment term in this case equals zero, and optimal rewards achieve first-best hospital welfare. As we will show in Section IV, this case is empirically relevant because the NKR is well within the region of approximately constant returns to scale. Therefore, optimal mechanisms can be calculated in practice by estimating marginal products.

Moreover, when there are constant returns to scale, there is no need to consider nonlinear rewards because we can achieve first-best hospital welfare with linear rewards. One approach for using these results in practice is to introduce a dynamic points mechanism. For each submission, a platform can credit points to a hospital equal to the marginal product. Then a point can be subtracted whenever a hospital conducts a transplant. The platform performs optimal matches with a constraint that no hospital's point balance falls below a certain level. We discuss the design of this type of mechanism in Section VA. Another approach is to use a mandate that forces hospitals to make the efficient level of submissions. We discuss this type of policy response in Section VC.

The third part of the theorem states that if the production function exhibits constant returns to scale and there are no agency problems, then the optimal mechanism achieves first-best social welfare. This result clarifies that there are two possible sources of inefficiency: suboptimal platform incentives (i.e., rewards that deviate from marginal products, $\mathbf{p} \neq \nabla \mathbf{f}$) and agency problems (hospitals failing to fully internalize the welfare of the parties they represent, $C^h \neq SC^h$).

Figure 4 depicts these two market failures in a graphical framework reminiscent of supply and demand. It ignores the adjustment term because estimates in Section IV will show that NKR is at approximately constant returns to scale. The horizontal axis plots aggregate supply $\mathbf{q}$. The vertical axis plots marginal products, social costs, and social benefits. The current vector of rewards, which is equal to the probability of matching each pair, is denoted by $\mathbf{p}_0$. The current quantity supplied, $\mathbf{q}_0$, is determined by the aggregate supply curve evaluated at current rewards. The curve $\nabla SC(\mathbf{q})$ is the marginal aggregate social cost if hospitals choose privately optimal quantities given rewards $\mathbf{P}_S(\mathbf{q})$.[30]

The hospital-optimal quantity $\mathbf{q}^*$ equates $\nabla \mathbf{f}$ with marginal aggregate private costs. Thus, the first market failure is that the platform gives inefficient incentives. The second market failure is that there are agency problems because hospitals do

[30] Formally, $SC(\mathbf{q}) \equiv \sum_{h=1}^{H} SC^h\big(\mathbf{S}^h(\mathbf{P}_S(\mathbf{q}))\big)$ is the reward-moderated social cost. The figure assumes that each individual hospital supply is uniquely defined and that $SC(\mathbf{q})$ is differentiable. See online Appendix A.2.

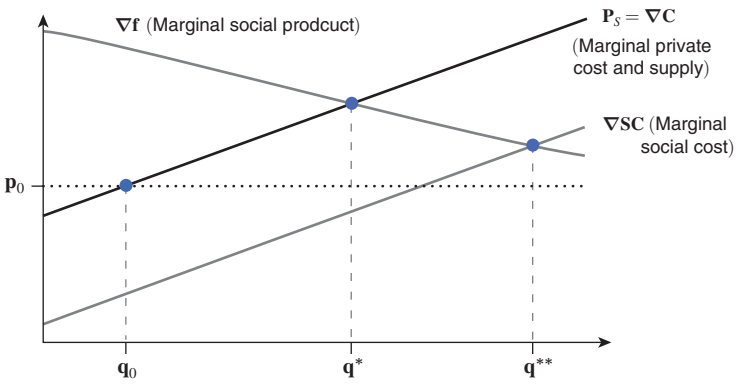


FIGURE 4. THE TWO SOURCES OF MARKET FAILURE

*Notes:* The horizontal axis represents aggregate quantity of submissions into the kidney exchange platform. The curves represent the marginal product of submissions, $\nabla\mathbf{f}(\mathbf{q})$, the marginal private cost of submissions from the perspective of hospitals, $\nabla\mathbf{C}(\mathbf{q})$ (which is equal to $\mathbf{P}_S(\mathbf{q})$, the inverse aggregate supply), and the marginal social cost of submissions, $\nabla\mathbf{SC}(\mathbf{q})$ (see footnote 30). Both axes represent $I$-dimensional vectors. The figure depicts the current quantity, with agency problems and a suboptimal mechanism, $\mathbf{q}_0$; the quantity from a hospital-optimal mechanism, but with agency problems, $\mathbf{q}^*$; and the first-best quantity, with an efficient mechanism and no agency problems, $\mathbf{q}^{**}$.

not fully account for the costs and benefits of other parties (i.e., $C^h \neq SC^h$). The aggregate quantity $\mathbf{q}^{**}$ maximizes social welfare subject to hospitals responding optimally given a rewards vector. It attains the first-best social and hospital welfare if we also solve agency problems so that $C^h = SC^h$, which makes the two welfare notions coincide.

In the example above, agency problems can be solved by reimbursing hospitals for the costs of kidney exchange through the platform $K^h(q^h)$. Note that it is not possible to reach the first-best social welfare by only improving the mechanism because there are often not enough transplants to set rewards that are high enough to solve agency problems (see Proposition A2 in the online Appendix for details). Alternatively, under certain assumptions, a regulator can simply enforce the first-best allocation with a mandate.

These arguments suggest two alternative approaches for policy. One is to implement a mandate that, for example, requires that hospitals submit all patients and donors to a national platform. Under certain assumptions, this is socially optimal. Another approach is both to have exchanges use efficient mechanisms and to create policies that solve agency problems. This approach avoids the political barriers that a mandate might face and is also optimal, under certain assumptions. Section V elaborates on these issues.

An upshot of this analysis is that, much like in traditional markets, many key questions about kidney exchange depend on the production function, which we turn to next.

## IV. Production Function Estimates and Results

We now recover the production function using data from the dominant kidney exchange platform during our sample period, the NKR (see Table 1). We use these

estimates to measure the inefficiency due to market fragmentation, calculate the rewards in an optimal mechanism, and assess the efficiency gain from moving to an optimal mechanism.

### A. *An Engineering Approach*

Production functions are commonly estimated using data on inputs and outputs from several firms. The key econometric challenges in this literature are endogeneity in the chosen inputs and selection in the set of operating firms (see Marschak and Andrews 1944, Olley and Pakes 1996). Unfortunately, this approach is not appropriate in our setting for three reasons. First, the standard methods are best suited for low-dimensional production functions that only depend on a few inputs, such as capital and labor. In our case, the inputs are different submission types, which are numerous (they vary by blood type, sensitization, and whether the patient or donor is paired). Second, commonly used functional forms, such as Cobb-Douglas, restrict all inputs to be substitutes, a property that may not be appropriate in the kidney exchange context.[31] Third, the standard methods depend on a panel dataset with inputs and outputs of multiple firms and exogenous variation of inputs. However, we only have data from a single, large platform.

We circumvent these econometric issues by using an engineering approach based on detailed institutional knowledge and administrative data on the processes involved in organizing kidney exchange. This approach was first introduced within economics by Chenery (1949), but fell out of favor due to its complexity and because some relevant inputs (e.g., managerial resources) were often difficult to measure (Walters 1963). However, it is particularly well suited for our study because we have detailed data on NKR's operations, and the composition and biological compatibility of its patient pool.[32]

We simulate the various steps involved in organizing kidney exchange to evaluate the number of transplants, $f(\mathbf{q};\boldsymbol{\theta})$, that can be produced with a flow of submissions, $\mathbf{q}$, given parameters, $\boldsymbol{\theta}$. The simulation is dynamic, with each period representing one day. There are four steps that take place: hospital submissions, transplant proposal, final review and transplantation, and departure. The parameters governing the first and last steps are directly estimated from the NKR data; the parameters involved in the second step are known; and the parameters from the third step are calibrated to fit observed transplantation probabilities for various patient and donor types, as well as the average length of chains. Our estimation and calibration methods are described in the following paragraphs, with details provided in online Appendix D.

(i) **Submissions:** Hospitals submit patients and donors, either individually or in pairs, to the platform. These submissions are added to the current pool of patients and donors already registered with the exchange. Patients and doctors, at this time, can submit minimal acceptance criteria for a donor.

[31] In a transplant production function, one input type may either be a substitute or a complement with another. For example, an A–O pair and an O–A pair are complementary since they together produce one pairwise kidney exchange. However, an O–A pair is a substitute input for an O–O pair in this setting.

[32] One of us (Ashlagi) developed the matching software for several platforms and has worked with NKR.

Submissions arrive according to a Poisson process with arrival rates given by a vector $\mathbf{q}$, whose dimension equals the number of submission types, $I$. The NKR baseline is represented by $\mathbf{q}_0$. We estimate the daily arrival rate of each submission type $i$ as average number of arrivals per year. An identical arrival process with Poisson arrival rates $\mathbf{q}$ allows us to calculate the production function at other arrival rates $\mathbf{q} \neq \mathbf{q}_0$.

Our exercises will start by treating each submission as a separate type ($I = 1{,}930$). We will then aggregate types to best predict probabilities of matching and marginal products using biological characteristics that are relevant for kidney exchange (e.g., blood type and patient PRA).

(ii) **Transplant Proposal:** Each day, the NKR identifies an optimal weighted set of potential exchanges within the stock of patients and donors. This algorithm incorporates four constraints. First, none of the proposed transplants should be (known to be) biologically incompatible or ruled out by preset acceptance criteria. These constraints are directly observed in the data. Second, no donor or recipient can be involved in more than one transplant. Third, a donor in a pair is only asked to donate an organ if her associated recipient is proposed a transplant. Finally, the cycle size is limited, because of the logistical difficulties inherent in organizing many simultaneous surgeries.[33]

The parameters of this algorithm are the weights used by the NKR for a transplanting a given donor and patient and the maximum cycle size. Consistent with NKR policy and observed data, we prohibit all cycles of length four or greater. The weights are known to one of the authors (Ashlagi) and are detailed in online Appendix D. They prioritize unlikely matches in an attempt to utilize hard-to-match donors and transplant hard-to-match patients whenever possible. The weights typically only break ties between two matches with the same number of transplants in favor of retaining patients and donors who are more likely to match in the future.

(iii) **Final Review and Transplantation:** Each proposed transplant is reviewed by doctors, patients, and donors, and approved before it is performed. Both approval and biological testing can take several days. Moreover, patients and donors in proposed transplants that are under review on a given day are excluded from the maximal matching algorithm on that day. This step also involves a final set of blood tests to ensure biological compatibility.[34] Cycles in which any patient refuses or is found to be incompatible with the proposed

[33] Formally, the NKR maximizes $\sum_{jk} c_{jk} w_{jk} x_{jk}$ by picking $x_{jk} \in \{0,1\}$, where $x_{jk} = 1$ denotes a proposed transplant from donor $k$ to patient $j$; $w_{jk}$ is the weight given to the transplant by the NKR; and $c_{jk}$ is an indicator for whether the transplant is biologically compatible and acceptable. This problem is subject to three additional constraints. First, no donor or patient is involved in more than one transplant, i.e., $\sum_j x_{jk} \leq 1$ and $\sum_k x_{jk} \leq 1$. Second, if donor $k$ and patient $j$ belong to a pair, then $x_{j'k} = 1$ for some $j'$ only if $x_{jk'} = 1$ for some donor $k'$. To write the third constraint, note that a cycle of length $n$ is an ordered tuple, $(j_1, j_2, \ldots, j_n)$ where $x_{j_k j_{k+1}} = 1$ for $k < n$ and $x_{j_n j_1} = 1$. We impose the constraint $n \leq 3$. Because there are a very large number of cycle length constraints, we first solve a relaxed problem without this last constraint and iteratively add the constraints to prohibit large cycles. Online Appendix D provides further details.

[34] Failures are recorded by setting $c_{jk} = 0$ (see previous footnote) for future iterations if the donor $k$ was refused by patient $j$.

donor are abandoned. NKR usually abandons chains in which the second patient cannot be transplanted. For other chains, all proposals until the first failure are consummated. The donor belonging to the final patient–donor pair in such a chain may initiate new chains in the future, much like an altruistic donor. This donor is often referred to as the "bridge" donor. Consistent with NKR policy, unpaired patients are prioritized according to the net difference between altruistic donors and unpaired patients previously transplanted by the patient's hospital.

This step results in frictions within the system that reduce transplantation rates (Agarwal et al. 2018). The parameters that govern these frictions are the time required for each of the two approval steps, the probability that a proposed transplant is abandoned in each step, and the duration for which a bridge donor is retained in the pool before donating her kidney to a patient on the deceased donor list.

Unfortunately, we do not have detailed data on which transplants were refused, how often transplants were aborted due to biological testing, or how long each review phase takes. Additionally, the NKR does not seem to have clear-cut algorithmic policies on how to use bridge donors. Chains would be indefinitely long if bridge donors were allowed to initiate new chains forever but too short if bridge donors were not used. Although cases of donors reneging are rare (Cowan et al. 2017), platforms try to transplant bridge donors quickly, to an unpaired patient if necessary, to them.

We calibrate these parameters by simulating our model to find values that most closely replicate the match probabilities, durations, pool size, and chain lengths observed in our data. We match average values of each of these variables (except for chain length) for altruistic donors, patient–donor pairs, and unpaired patients, each partitioned by coarse blood type categories.[35] Moments based on match probabilities and durations are motivated by their close relationship to the flow of transplants produced and the incentives for hospitals. Pool size seeks to capture aggregate information about the productivity of the NKR. Finally, chain lengths capture the importance of altruistic donors.

Our simulations suggest that the best fit to these moments is, for both the acceptance and the biological testing phases, a two-week period and a one-fifth failure rate. Reducing the failure rates in simulations primarily increases chain length and transplantation rates, while reducing the duration of either phase increases the transplantation rates without having a large effect on chain length. For the bridge donor policy, we find that a hold period of 30 days best fits the data.

Details on the fit of our calibrated parameters are provided in online Appendix D.5. Further, online Appendix E repeats all of our analyses under alternative parameters to examine robustness of our results.

[35] In principle, we could have estimated these parameters using simulated minimum distance. However, a simulation for each parameter value can take weeks, making optimization over the parameter set infeasible.

**Departure:** Patients and donors often depart the NKR without a transplant. A patient and his/her associated donor may leave the platform because the patient dies, becomes too sick to transplant, or receives a kidney transplant elsewhere. Therefore, we need to estimate the probability that a patient or a donor leaves the NKR without a transplant.[36]

We estimate an exponential hazard model of departures by maximum likelihood, using the registration and transplantation dates (if transplanted) for each patient and donor, and data on how long the patient or donor was registered in the NKR without a transplant.[37] The departure rates in the model depend on the fraction of donors (patients) ever registered with the NKR who are compatible with a patient (donor), blood type, and age at registration. Online Appendix D.2.2 presents the estimates for the model.

This procedure allows us to evaluate a transplant production function for any vector of inputs $\mathbf{q}$ by simulating each of these events for each calendar day. Given any initial pool of patients and donors in the NKR, these simulations generate a Markov chain with a sequence of registrations, transplants, and departures. We initialize the NKR pool with the set of patients and donors registered on April 1, 2012, and burn-in 2,000 simulation days in each run. The dependence on the initial pool eventually fades away. We compute the time average of the total number of transplants to estimate $f$ by

$$\hat{f}(\mathbf{q}) = \frac{1}{T}\sum_{t=1}^{T} y_t,$$

where $T$ is the total number of days simulated and $y_t$ is the total number of transplants in period $t$ of our simulation. In what follows, we report estimates based on an average of 100 simulations. Standard errors are calculated using the non-overlapping batch means estimator described in online Appendix D.4.

### B. *Returns to Scale and Misallocation*

*Returns to Scale*.—We first document the estimated returns to scale in the transplant production function by evaluating it for submission vectors $\mathbf{q}$ with the same composition as the NKR (i.e., scalar multiples of $\mathbf{q}_0$). We then consider average products, where scale is measured by the total flow of donors submitted per year, which

[36] Our approach will treat all donor departures as a lost opportunity for a transplant if a better design can use that donor for a transplant. To validate this assumption, we tried to determine the outcome of paired patients who leave the NKR without a transplant by matching them to the OPTN data on all living- and deceased-donor transplants. Our ability to follow these patients is not perfect, but approximately three-quarters of patients could be perfectly matched on the HLA-A, B, and DR loci; gender; and blood type. A majority of patients either remained untransplanted or received a deceased-donor transplant, effectively crowding out a kidney from another patient. Of those who received a living-donor transplant, most received direct donations and the vast majority did not utilize a multi-hospital kidney exchange platform. These facts support our treatment of departures as an appropriate approximation.

[37] Specifically, the departure rate for registration $j$ is given by $\lambda_{g_j}\exp(z_j\beta)$, where $g_j$ denotes whether $j$ is an altruistic donor, a patient–donor pair, or an unpaired patient; $\lambda_{g_j}$ is a group-specific constant departure risk; $z_j$ denotes a vector of characteristics for $j$; and $\beta$ is a conformable vector of coefficients. We fit with the censored observations of departure times for each registration in the NKR, where censoring can occur because we only observe a lower bound for the departure time if $j$ was transplanted or remained in the NKR pool at the end of our sample period.

we denote as $x(\mathbf{q})$. This measure is also the flow of transplants that a platform could perform if all donors were utilized. Figure 5 plots the average product measured as $f(\mathbf{q})/x(\mathbf{q})$.

The figure shows that there are increasing returns to scale, but that productivity eventually plateaus. With a scale of donor arrivals per year, the NKR is well within the region of approximately constant returns to scale.[38,39] The NKR has an average product of 0.51 transplants per donor, which varies only slightly once the scale is sufficiently large. A platform that is one-half the size of NKR has an average product of 0.46 transplants per donor, while a platform that is double the size has an average product of 0.547 transplants per donor. Therefore, the market can operate at a high level of efficiency even if there are a handful of competing platforms. These estimates suggest that mergers of sufficiently large platforms would have small effects on efficiency.

Next, we use these estimates to calculate whether individual-hospital platforms operate at an efficient scale. Recall that within-hospital transplants collectively account for the majority of kidney exchanges. A challenge with this exercise is that we observe neither the number nor the composition of patients and donors available to a hospital. To make progress, assume for the moment that hospitals have the same production technology and composition as the NKR. Further, assume that hospitals conducting within-hospital transplants do not participate in the NKR. Under these assumptions, one can use an individual hospital's observed rate of kidney exchange transplants to infer its scale. Specifically, let $y^h$ be the flow of within-hospital kidney exchange transplants conducted at hospital $h$. We estimate the flow $x^h$ of donors available to hospital $h$ as the flow necessary to produce $y^h$ with the same composition and technology as the NKR. That is, $x^h$ solves $y^h = \hat{f}\left(x^h \cdot \mathbf{q}_0/x(\mathbf{q}_0)\right)$, where $\mathbf{q}_0$ is the flow of submissions received by the NKR. The potential biases and robustness of our results to altering these assumptions are discussed in Section IVB.

This exercise suggests that almost all individual hospitals operate far below the efficient scale. The histogram in Figure 5 shows the estimated distribution of hospital scale. The median hospital has a scale of 9 donor arrivals per year. The ninetieth percentile is 27 donor arrivals per year. The largest, Methodist Hospital in San Antonio, has a scale of 109 donor arrivals per year. The average product at these

[38] The error bars use a non-overlapping batch-means estimator from a simulation of 20,000 periods. These errors account for sampling variance in the patient and donor registration process, and simulation error in the transplant proposal, refusal, and departure process, but hold the parameters of the simulation model fixed, ignoring first-stage estimation error. Accounting for this source of uncertainty is burdensome because we do not have a closed-form solution for the asymptotic variance of the production function. Moreover, a bootstrap procedure that simulates the entire production function several times is computationally prohibitive. To assess whether first-stage estimation error is important, we simulated the production function at NKR's scale and composition using ten bootstrapped estimates of hazard rate parameters. The resulting standard error in the average product across these draws was only 0.01 transplants per donor.

[39] We assessed whether this finding is driven by our approach to calculating the production function by conducting two complementary exercises. First, we directly investigated the returns to scale at the NKR by examining the relationship between the number of transplants conducted by the NKR per quarter and the number of donors submitted per quarter. Our estimates are noisy and do not rule out constant returns to scale. This finding is not sensitive to other partitions of time or the use of moving averages. Unfortunately, the variation in the number of donors registered is limited because our registration data starts in April 2012. Second, we carried out a simplified static simulation of a kidney exchange platform with no matching frictions and only a few types of pairs. As in Figure 5, the returns to scale rapidly increase and plateau, consistent with the model in Roth, Sönmez, and Ünver (2007), where the number of transplants produced grows linearly in the platform's size. Details of our simulation and the descriptive evidence are available upon request.

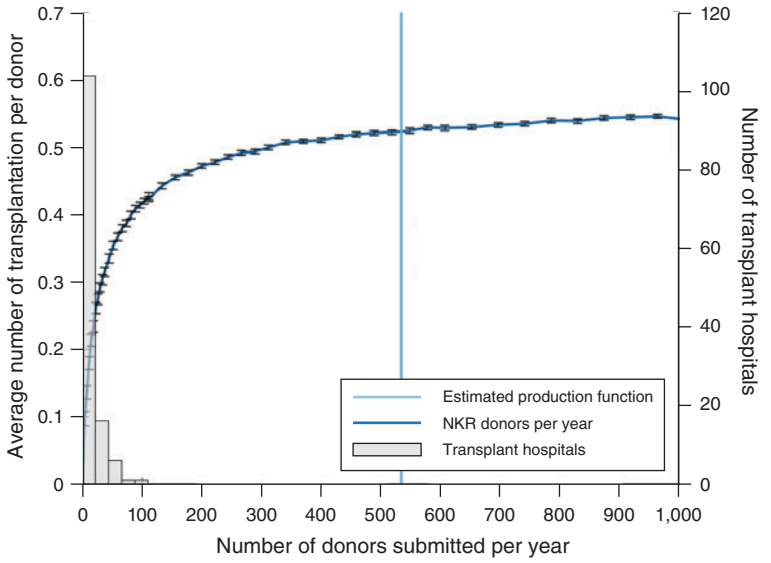


FIGURE 5. PRODUCTION EFFICIENCY VERSUS SCALE

*Notes:* The line plot represents the average product of a kidney exchange platform versus its scale. The histogram is based on the estimated scale of various hospitals. The left vertical axis represents average products, defined as the share of pairs and altruists who are transplanted. The right vertical axis is the scale for the histogram. The horizontal axis represents scale, measured as the yearly arrival rate of pairs and altruists. The error bars on the estimated production function show a 95 percent confidence interval. The plot uses the baseline parameters and the pool composition from the NKR.

scales is 0.16, 0.29, and 0.42 transplants per donor, respectively. Thus, at our estimated production function, even the largest single-hospital platform does not operate at an efficient scale. UNOS and APD both have an estimated average product of approximately 0.40 transplants per donor. Hence, the implied efficiency losses are considerable, even for the largest platform other than the NKR. These results are consistent with the evidence presented in Section IIB that hospitals often perform matches that are socially inefficient and that UNOS and APD are also less efficient than the NKR.

*Misallocation: Inefficiency Due to Small Production Scale*.—We start by using the baseline approach described above to estimate inefficiency due to market fragmentation. That is, we estimate how many additional transplants would be performed if the entire kidney exchange market functioned at NKR's efficiency. We use a hospital's estimated scale to calculate the difference in average product between the hospital and NKR. Because NKR operates at constant returns to scale, this difference multiplied by the hospital scale is the total number of transplants that are lost due to the hospital conducting kidney exchange at an inefficiently small scale. The aggregate lost transplants equals the total deadweight loss because our social welfare function is the total number of transplants nationwide. The estimated deadweight loss presented in Table 3 shows that 500.1 transplants are lost per year due to market fragmentation (panel A, column 1). This number is large relative to the 800 transplants conducted through kidney exchange each year. Furthermore, the economic value of these lost transplants exceeds $500 million per year based on the

Held et al. (2016) estimates of a transplant's value. The cost savings alone are on the order of $150 million per year.

This baseline approach is simple but suffers from four potential biases. First, the composition of submissions in hospitals may differ from that of the NKR. We assess robustness to this assumption by estimating inefficiency using patient and donor compositions from three different groups of hospitals: all hospitals (column 1), hospitals in the top quartile of intensive margin participation rate (column 2), and hospitals in the bottom quartile (column 3).[40] If heterogeneity in participation policies is correlated with composition, then estimates from a production function using patients and donors from each of these groups helps assess robustness to potential compositional differences between single-hospital platforms and the NKR. For example, one hypothesis is that hospitals in the top quartile of participation should submit a less selected pool if heterogeneity in participation rates is driven primarily by hospital policy.[41] Comparing estimates across these three groups suggests that overall inefficiency exceeds 430 transplants per year under alternative assumptions on composition.

Second, our baseline approach assumes that all within-hospital transplants are produced by hospitals in isolation from the rest of the market. The bias due to hospitals that also participate in national platforms does not have a clear direction. We address this issue by disaggregating the efficiency losses by whether a hospital participates in the NKR, APD, or UNOS and by the fraction of the hospital's paired kidney exchanges that are conducted through the NKR. If we restrict attention to the 96 hospitals that do not participate in NKR, the efficiency loss in column 1 is 243.5 transplants per year (panel C, excluding the NKR row). Some of these hospitals participate in UNOS or APD and may be producing transplants at a more efficient scale. Even if we assume that each of the hospitals that participate in UNOS or APD produce transplants at the estimated scales for the two platforms, we still estimate that the deadweight loss in column 1 would be 143.7 transplants per year.[42] However, this extremely conservative calculation is likely at slack for two reasons. First, even among the non-NKR hospitals that participate in either UNOS or APD, two-thirds of kidney exchange transplants are performed within hospital (panel C), while the deadweight loss lower bound of 143.7 assumes that all transplants are produced at the APD/UNOS scale. Second, it ignores deadweight loss from hospitals that participate in NKR. Among the set of NKR participants, just the 17 hospitals in the lowest quartile of fraction of transplants performed in NKR contribute an efficiency

[40] We measure participation rate as donors submitted to NKR divided by donors submitted to NKR plus donors transplanted in a within-hospital exchange.

[41] Some large hospitals submit most of their pairs to NKR as a matter of policy. Smaller hospitals that have incurred the costs of becoming an NKR member often submit all of their patients to the platform because they have few possibilities of organizing an exchange independently. These facts suggest that the group of hospitals in the top quartile should submit a less selected pool if heterogeneity in these policies and hospital size is not correlated with composition. It is also possible that hospitals that participate in the NKR exclusively do so precisely because they have a particularly hard to match pool. While it is difficult to test these hypotheses without data on the patients and donors available to a hospital, our data indicate that hospitals with high NKR participation rates submit easier to match patients in terms of PRA (online Appendix Figure E16). This limited evidence points to the former arguments as more important.

[42] The deadweight loss from hospitals that do not participate in any of the three national platforms alone is 125.6 transplants per year. For hospitals that participate only in UNOS or APD, the deadweight loss is 18.1 transplants per year, assuming that all kidney exchange transplants from these hospitals are produced at the scale of the platform in which they participate.

Table 3—Total Efficiency Loss

| | Number of hospitals | Kidney exchange transplants per year | Within hospital kidney exchange transplants per year | Efficiency loss: Additional kidney exchange transplants (1) | (2) | (3) |
|---|---|---|---|---|---|---|
| *Panel A. All hospitals* | | | | | | |
| All hospitals | 164 | 800.5 | 465.4 | 500.1 | 432.4 | 644.7 |
| *Panel B. By hospital size* (*number of KE transplants per year*) | | | | | | |
| Top quartile | 42 | 598.8 | 358.3 | 257.1 | 213.5 | 343.1 |
| Second quartile | 48 | 143.2 | 73.4 | 139.7 | 122.2 | 169.7 |
| Third quartile | 40 | 45.7 | 27.7 | 76.6 | 73.4 | 101.0 |
| Bottom quartile | 34 | 12.7 | 6.0 | 26.8 | 23.4 | 31.0 |
| *Panel C. By platform membership* | | | | | | |
| NKR | 68 | 580.5 | 297.2 | 256.6 | 217.7 | 335.5 |
| Only UNOS or APD | 45 | 133.0 | 90.7 | 117.9 | 101.9 | 149.8 |
| None | 51 | 86.9 | 77.6 | 125.6 | 112.8 | 159.4 |
| *Panel D. By NKR participation rate* (*fraction of KE transplants facilitated through the NKR*) | | | | | | |
| Top quartile | 17 | 65.2 | 8.2 | 16.9 | 15.5 | 21.5 |
| Second quartile | 17 | 102.3 | 27.0 | 47.9 | 42.9 | 60.0 |
| Third quartile | 17 | 196.7 | 98.2 | 88.7 | 75.7 | 114.6 |
| Bottom quartile | 17 | 216.2 | 163.8 | 103.1 | 83.6 | 139.4 |

*Notes:* Column 1 assumes that the typical transplant hospital has a composition of patient-donor pairs and altruistic donors given by the average registration in the NKR. Column 2 assumes a composition based on hospitals in the top quartile of NKR participation. Column 3 assumes a composition based on hospitals in the lowest quartile of participation rate. Transplants per year is calculated using data between April 1, 2012 and December 4, 2014.

loss of 103.1 transplants per year (panel D). In summary, despite potential bias due to some hospitals participating in large platforms, this decomposition suggests that a loss of 240 transplants per year is a conservative estimate for the costs of market fragmentation.

Third, hospitals may use a different matching technology than the NKR. For example, Bingaman et al. (2012) reports that Methodist Hospital in San Antonio, which is now perhaps the most sophisticated single-hospital program, initially used a Microsoft Access Database and that their algorithm was "stratified by ABO compatibility and then by HLA compatibility." Such algorithms are less efficient than the linear-programming algorithms used by the NKR.[43] On the other hand, single-hospital programs face simpler logistical constraints, which may increase their productivity vis-à-vis our estimates. The direction of this bias is not unambiguous, but it is more likely that single-hospital platforms are less efficient than our estimated production function.[44]

Fourth, these exercises keep the set of patients and donors interested in kidney exchange fixed. However, this flow is endogenous and affects the magnitude of the deadweight loss. Although the direction of this bias is ambiguous, our baseline approach likely yields a conservative estimate of overall market inefficiency. The

[43] In 2013, Methodist Hospital in San Antonio adopted software written by one of us (Ashlagi).

[44] See Agarwal et al. (2018) for an analysis of how various logistics influence the productivity of a kidney exchange platform. NKR's practices are optimized to maximize the number of transplants given the available patients and donors.

most likely bias comes from hospitals valuing transplants at less than the social value and, due to administrative costs, expending inefficiently low effort in recruiting patients and donors. If incentives were optimal, hospitals might try to recruit more, and more valuable, donors into kidney exchange. Our approach does not account for this margin because we do not observe recruitment efforts; therefore, we are likely to underestimate overall market inefficiency.

Table 3 also points to which types of hospitals produce most of the inefficiency. Consider column 1 and, for the purposes of this decomposition, ignore the biases discussed above. Even though they perform internal exchanges more efficiently, large hospitals account for most of the inefficiency because their market share is higher (panel B). Indeed, 51.4 percent of the losses come from hospitals that have a number of kidney exchange transplants in the top quartile. Moreover, both the intensive and extensive margins of participation are important. A little less than one-half of the efficiency losses are due to hospitals that do not participate in the NKR at all, and one-quarter are from hospitals that do not participate in any of the national platforms (panel C). Among hospitals that do participate in the NKR, a large share of the efficiency loss is due to the hospitals with low participation (panel D).

To summarize, although the baseline estimate of 500.1 lost transplants is potentially biased, a battery of robustness exercises suggest the deadweight loss from market fragmentation is at least 240 transplants a year. Additionally, these estimates do not appear to be driven by compositional differences in the kidney exchange pool. Online Appendix Table E9 further evaluates these results' robustness to alternative choices for the production function parameters that were calibrated. Across various specifications, we continue to find that an estimated 240 lost transplants is conservative.

These results are consistent with our descriptive finding that hospitals often perform inefficient matches, and they are robust to moderately large levels of misspecification. Even the most conservative estimates are significant because the baseline estimates suggest inefficiency of 63 percent of the market size. The robustness of the shape of the production function is driven, in large part, by basic biological compatibility constraints. Finding matches is hard with few patients and donors because of both blood-type and tissue-type compatibility constraints. And, as suggested in the large market limit of Roth, Sönmez, and Ünver (2007), tissue-type compatibility constraints become less of a barrier in a large market.

### C. *Inefficiency of Current Mechanisms*

Theorem 1 shows that optimal rewards are approximately equal to marginal products. That is, $\mathbf{p}^* = \nabla \mathbf{f}(\mathbf{q}^*) - \mathbf{A}(\mathbf{q}^*)$, where $\mathbf{q}^*$ and $\mathbf{p}^*$ are the aggregate quantities and rewards that maximize hospital welfare. We will test this equality at the aggregate supply and rewards in our data.

Current rewards, $\mathbf{p}_0$, equal the probabilities of matching for each kind of submission. These probabilities can be easily estimated from our simulations, and the estimated probabilities closely match those from the data (see online Appendix D.5). Marginal products, $\nabla \mathbf{f}(\mathbf{q}_0)$, can be estimated by numerically differentiating the production function. In principle, calculating the adjustment term requires estimates of the supply elasticity matrix, which is not feasible with our data. But, the adjustment

term is small because returns to scale are approximately constant for NKR's size. Therefore, optimal rewards are approximately equal to marginal products. Formally, Theorem 1 implies that the quantity-weighted average of the adjustment term is given by

$$\frac{\mathbf{A}(\mathbf{q}) \cdot \mathbf{q}}{\|\mathbf{q}\|_1} = \frac{\nabla \mathbf{f}(\mathbf{q}) \cdot \mathbf{q}}{\|\mathbf{q}\|_1} - \frac{f(\mathbf{q})}{\|\mathbf{q}\|_1}.$$

That is, the average level of shading equals the difference between the average marginal product and the average product. Evaluating this formula using the estimated production function and numerical derivatives for each of the 1,930 submission types yields an average shading of only $2.16 \times 10^{-4}$ transplants per submission. In what follows, we simply approximate optimal rewards with marginal products.

Figure 6 plots current rewards (the probabilities of matching, $\mathbf{p}_0$) versus optimal rewards (marginal products, $\nabla \mathbf{f}(\mathbf{q}_0)$). Following Roth, Sönmez, and Ünver (2007), the patient and donor categories are aggregated by under-demanded, over-demanded, and self-demanded types, split by sensitization level.

The marginal products are qualitatively similar to the Roth, Sönmez, and Ünver (2007) theoretical predictions discussed in Section I. The marginal product of an under-demanded pair is 0, both in our estimates and in the model. The estimates differ for other types. For example, the marginal product of an over-demanded pair with low sensitization is 2 transplants per submission in the Roth, Sönmez, and Ünver (2007) model, but 1.34 transplants per submission in our estimates. One reason for this difference is that, in our data, these pairs are only matched with probability 0.80. Our empirical model also refines the predictions from the theoretical models by showing how marginal products vary with sensitization. For example, the marginal products of over-demanded and self-demanded pairs are considerably lower if these pairs are sensitized. These finer results can be important when designing practical mechanisms.

The figure shows a large wedge between current and optimal rewards. If current rewards were optimal, all points on these two figures would be on the 45-degree line. Altruistic donors and over-demanded pairs with low PRA are far below this line. Over-demanded pairs with low sensitization have marginal products of 1.34 transplants per submission, but the probability of matching them is only 0.80. Even more extreme, altruistic O donors have a marginal product of 1.74 transplants per submission, but their probability of matching is only 0.93. Therefore, hospitals are not rewarded enough for submitting these types, which may explain why we see relatively few of them submitted to the NKR. Other submission types are drastically overpriced. Under-demanded pairs with low sensitization have marginal products of approximately 0.05 transplants per submission but have a probability of being matched of around 0.36. Similarly, unpaired patients have low marginal products but a significant probability of being matched. These differences suggest the platform can do considerably better by increasing rewards to the productive and undervalued submissions while reducing rewards to the unproductive submissions.

These marginal products are of independent interest as well. For example, Sönmez, Ünver, and Yenmez (2018) analyzes policies that incentivize compatible pairs to join the platform by prioritizing patients on the deceased donor list in case they need a second transplant in the future. While we cannot directly compute the value of a

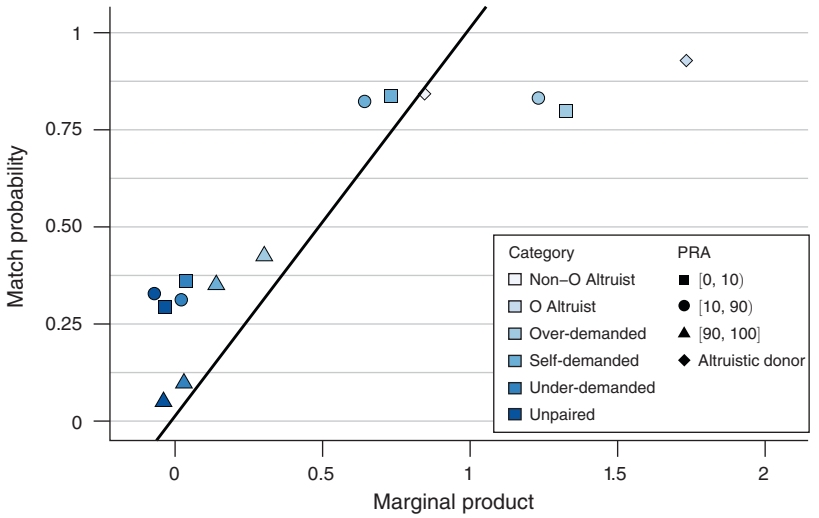


FIGURE 6. PRIVATE VERSUS SOCIALLY OPTIMAL REWARDS FOR SUBMISSION TYPES

*Notes:* The vertical axis is the probability of a submission being matched, which are the private rewards that hospitals receive according to current exchange rules. The horizontal axis plots the marginal product of a submission (in transplants per submission), which equals the social contribution of the submission in terms of transplants. Matching probabilities and marginal products are calculated in the baseline simulation.

compatible pair using the NKR data, such pairs are most similar to over-demanded or self-demanded pairs with low PRA patients. Indeed, our results indicate that these pairs have the highest marginal products.[45] Another result of interest to a platform is that altruistic donors have a marginal product that is much lower than the length of a typical chain, which is four. This is because many patients transplanted through a chain could have otherwise been transplanted through a cycle.

## V. Discussion

### A. *Implementing a Point Mechanism*

Theorem 1 and the small magnitude of the adjustment term $\mathbf{A}(\mathbf{q})$ suggest that platforms should set rewards close to marginal products. We will now show that marginal products are highly predictable using a small number of patient and donor categories. Then, we will discuss the design of point mechanisms that are both approximately efficient and simple enough for practical application.

We use a regression tree to construct categories that best predict marginal products, allowing for the tree to depend on the patient's PRA, submission type (altruistic, patient–donor pair, unpaired patient), and blood type. Figure 7 shows

[45] The average marginal product of self-demanded pairs with PRA less than 50 is 0.74 transplants per submission and the average match probability is 0.83. Over-demanded types with PRA less than 50 have an average marginal product of 1.35 transplants per submission and an average match probability of 0.82. These marginal products may differ from those of compatible pairs because a patient with a compatible donor may be more selective than incompatible pairs.

the categories found by a standard algorithm for finding the best cross-validated predictor for the marginal products. These categories are intuitive, as they split submissions based primarily on submission type, whether the patient/donor is blood type O, and immune sensitivity. The within-category mean marginal products, $\nabla \mathbf{f}$, and probabilities of matching, $\mathbf{p}_0$, are dispersed relative to the (appropriately shrunk) within-category standard deviation. This suggests that marginal products and probabilities of matching are approximated with a small number of categories.

A mechanism that assigns points to a hospital based on these categories can be explained to participants with this tree or a simple table (see online Appendix Table D8). One implementation of the mechanism is as follows. Each time the hospital conducts a transplant for one of its patients, a point is subtracted from its account. Points are awarded when a submission is transplanted, since awarding them at time of submission creates an incentive for shill submissions. To account for the chance a submission is not transplanted, the marginal products should be divided by the probability of matching, $\mathbf{p}_0$, to give the appropriate incentives in expectation. These rewards are denoted by $\mathbf{r}^*$ in Figure 7.

While our analysis suggests that a low-dimensional point mechanism would likely achieve substantial efficiency gains, it abstracts away from several implementation issues. For example, our simplified steady-state model does not specify an extensive-form game, and hence it cannot be used to fully specify optimal mechanisms or to evaluate them. This raises practical and theoretical questions about how to design and implement a dynamic points mechanism. While resolving these details is beyond our present scope, we discuss some of the issues.

In both theory and practice, the point system described above is a natural mechanism for solving this problem. The dynamic mechanism design literature calls such a system a chip, scrip, or token mechanism. Relevant papers include Möbius (2001), Hauser and Hopenhayn (2008), and Abdulkadiroğlu and Bagwell (2013), which consider dynamic favor exchange, and Guo and Hörner (2015), which considers provision of goods to a consumer with stochastic valuations. The general finding of this literature is that token mechanisms, as proposed in Möbius (2001), do better than autarky but not as well as an optimal dynamic mechanism. In fact, token mechanisms are close to first-best if players are patient and there are many time periods. Results in Jackson and Sonnenschein (2007) imply that the inefficiency of token mechanisms declines as the square root of the number of periods (see also Guo and Hörner 2015). Thus, the theoretical literature suggests that point systems, while not exactly optimal, are simple and achieve a high level of efficiency.

Another motivation for using a point mechanism is its simplicity and similarity to fiat money. Similar mechanisms have been previously used in market design applications. For example, Prendergast (2017) describes how a similar mechanism was used to increase the efficiency of food distribution across food banks.

An important issue with point systems is that they require several "plumbing" decisions (Duflo 2017). Should the matching algorithm impose a strict bound on negative balances? If so, what is the optimal minimum balance constraint? A tight constraint provides stronger incentives to hospitals but prior theory suggests that it may reduce efficiency (Möbius 2001; Friedman, Halpern, and Kash

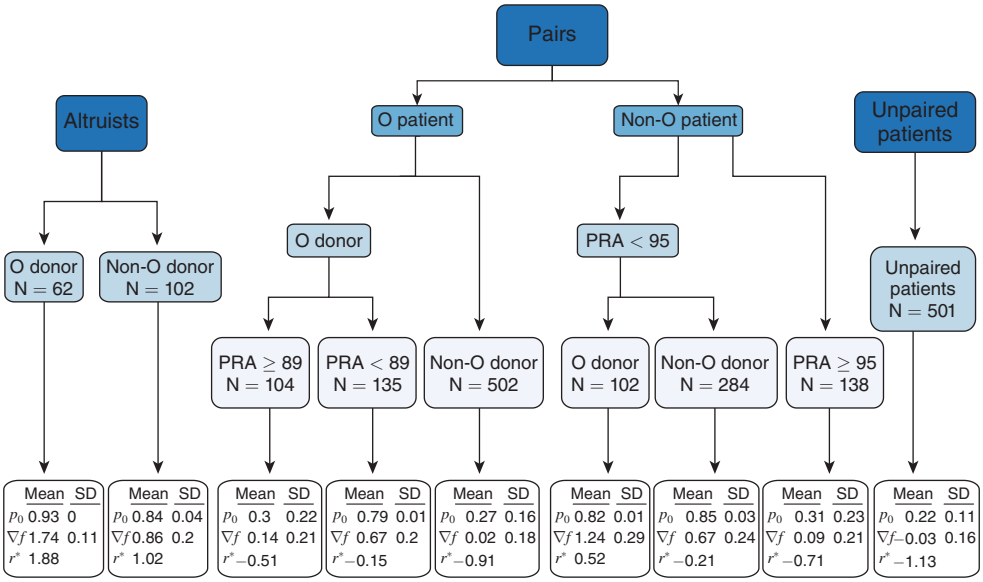


FIGURE 7. REGRESSION TREE FOR MARGINAL PRODUCTS

*Notes:* Categories are determined by regression tree analysis to predict marginal products as a function of whether a submission is a pair, an altruistic donor, or an unpaired patient; blood types; and the patient's PRA. Our procedure followed standard recommendations in Hastie, Tibshirani, and Friedman (2001). Specifically, we used 10-fold cross-validation to pick the penalty parameter on the number of nodes, required each leaf to have at least 20 observations, and pruned a leaf if it did not increase the overall fit by at least 2 percent. Standard errors for the simulations are calculated by following Robert and Casella (2004, chapter 12). The within-category standard deviation is estimated using shrinkage methods recommended in Morris (1983). The match probabilities in the current mechanism are denoted by $\mathbf{p}_0$, the marginal products are denoted by $\nabla\mathbf{f}$, and the optimal rewards at transplantation are denoted by $\mathbf{r}^*$. We calculate $\mathbf{r}^*$ by dividing $\nabla\mathbf{f}$ by $\mathbf{p}_0$, component-wise, and then subtracting 1 for all types except altruistic donors.

2006; Kash, Friedman, and Halpern 2007).[46] How often should marginal products be recalculated as the composition of patients and donors in the platform changes? Recalculating them often is complex and reduces transparency, but recalculating infrequently can reduce efficiency if changes in composition affect marginal products. An analysis of this issue requires a non-steady state dynamic model. However, previous work on kidney exchange and experiments with our simulations suggest that marginal products are not sensitive to the distribution of types. For example, marginal products are constant for a wide range of type compositions in the theoretical model of Roth, Sönmez, and Ünver (2007).

[46] In models where agents can exchange favors over time through a scrip currency, for any amount of scrip per capita, with sufficiently patient players, there are efficient equilibria with a high level of exchange. However, for a fixed discount factor, increasing the amount of scrip only increases efficiency up to a point: too much scrip leads to inefficient equilibria with no exchange. Similarly, Möbius (2001) shows that token mechanisms with less restrictive budget constraints are more efficient, but that these budget constraints must be sufficiently strict relative to the level of impatience in order to give agents incentives to provide favors. These results suggest that it is important to give participants incentive to trade favors, but that as long as this constraint is satisfied, it is more efficient to have point mechanisms with more flexible budget constraints. In kidney exchange, there are often a large number of different maximum cardinality matches, because of the large number of under-demanded pairs. This suggests that efficiency may not be severely compromised even if budget constraints are strict.

As in other market design applications, it is important to consider whether a proposed system is fair and ethically compelling (Roth 2007; Kominers, Teytelboym, and Crawford 2017; Li 2017). It is important to clarify that a points system involves no use of monetary payments to donors. Fundamentally, the point system ensures that, in the long run, the contributions of each hospital to the exchange are commensurate with the transplants performed. Nonetheless, in practice, legal clarification with respect to the National Organ Transplantation Act of 1984 and careful explanation of the system to participants may be necessary.[47]

### B. *Importance of Each Market Failure*

The misallocation analysis of Section IVB yields a conservative lower bound for the deadweight loss of about 240 transplants per year, or 30 percent of the market. The true deadweight loss is potentially much larger, as most specifications yield numbers approximately twice as large. Therefore, it must be the case that at least one of the two market failures (inefficient platform incentives and agency problems) is quantitatively important. We now report suggestive evidence that both market failures are important.

Online Appendix A.1 shows that the deadweight loss from inefficient platform incentives is given by a multidimensional version of the Harberger triangle formula. We cannot calculate this deadweight loss because we do not have data to reliably estimate the elasticity of hospital supply. Nevertheless, the large wedge between the current private and social incentives suggests the deadweight loss is significant unless the supply elasticity is extremely small.

To formalize this point, online Appendix A.1 estimates the deadweight loss in hospital welfare under different assumptions on supply elasticities. We calculated the worst-case deadweight loss given a maximum bound on own elasticity. With own elasticities bounded below by 2, the deadweight loss is over 40 transplants per year. With own elasticities bounded above by 6, the deadweight loss is as high as 100 transplants per year. These estimates are robust to assumptions on cross-elasticities. Because hospitals undervalue transplants if agency problems are important, the loss in social welfare is considerably larger. Specifically, if hospital costs are ignored, we estimate a loss in social welfare of at least 55 transplants per year and at most 120 transplants per year for elasticities between 2 and 6.

These results suggest that addressing inefficient platform incentives has a significant positive impact unless supply is extremely inelastic. The evidence in Section II is typical of markets with elastic supply: most hospitals only register a subset of their patients with the NKR, and many other hospitals do not participate. Both facts are consistent with many hospitals being on the margin, suggesting that hospitals respond to incentives and that supply is at least moderately elastic.[48]

[47] There are two reasons to believe that a points mechanism would be well received. First, previous NKR experimentation with the concept was found to be acceptable in practice. Second, we have discussed this issue with a few kidney exchange platforms, and all feel that a points mechanism is acceptable.

[48] Additionally, Ellison (2014) surveys transplant coordinators and finds three pieces of qualitative evidence that suggest that elasticities are not low. First, hassle costs are an often cited reason for not participating in kidney exchange platforms. Second, many transplant coordinators carefully consider whether to submit particular pairs to a platform and which platforms to partner with, which suggests that many decisions are marginal. Third, coordinators often mention that the time exchanges take to transplant patients is a major concern. Because time to transplant

Moreover, the results also imply that agency problems are important unless supply is extremely elastic. Under the hypothesis that there are no agency problems, hospital welfare equals social welfare, and the optimal mechanism reaches first-best (Theorem 1). Thus, the total deadweight loss in the misallocation analysis must be completely accounted for by the deadweight loss from the Harberger triangle analysis. Yet, even for a high elasticity bound of 6, our social deadweight loss estimate is at most 120 transplants per year. This is still below our lower bound result of 240 from the misallocation analysis. The only way these estimates can overlap is if we have high elasticities and our first-order approximation of deadweight loss is significantly downward-biased. Thus, attributing all the deadweight loss to inefficient platform incentives requires that supply is very elastic, our approximation is sufficiently convex, and the downward bias in the estimated lower bound on inefficiency is small.

While the quantitative magnitudes should be interpreted carefully due to data limitations, the upshot is that policies that address either market failure are likely to be valuable and generate gains on the order of hundreds of transplants per year. But, implementing optimal rewards alone will not eliminate most of the inefficiency. Proposition A.2 in the online Appendix formally shows that, typically, the platform does not produce enough transplants to correct for agency problems.

### C. *Mandate*

The previous subsection showed that agency problems appear to be first order. Hence, a points system of the sort described in Section VA should be complemented with subsidies for platform participation. An alternative to this two-pronged approach is a mandate that requires hospitals to perform all kidney exchanges in large, national platforms. Indeed, participation in a single national platform is required in the United Kingdom, the Netherlands, and Canada (Johnson et al. 2008, De Klerk et al. 2005, Malik and Cole 2014). A perfectly enforced mandate achieves the first-best welfare if all hospitals and platforms have the same production function, and that function exhibits (weakly) increasing returns to scale. While design details of a mandate are beyond the scope of our paper, we now consider four basic issues.

First, there are different types of mandate, with different welfare impacts. One form requires all hospitals to participate fully in one of several national platforms. A more heavy-handed mandate could also require existing national platforms to merge into a single platform. Our results suggest that most of the gains would come from the lighter regulation, since returns to scale are constant once a platform has reached NKR's size. Using the baseline approach from Section IVB, we estimate that the yearly flow of donors in the US kidney exchange market is 1,840 (ranging from 1,711 to 2,115 if we vary the assumptions used in Table 3 to estimate a hospital's scale). These estimates are between three and four times the current NKR scale. Even at four times NKR's scale, the average product of a monopoly kidney exchange with NKR's composition is 0.555 transplants per donor. If instead, there were two platforms twice the size of the NKR, the average product of each would be 0.547 transplants per donor.

depends on the flow of transplants, it is closely related to matching probabilities in the current mechanism. This evidence suggests that coordinators are sensitive to changes in rewards.

Both of these estimates are not much larger than the current average product of 0.51 transplants per donor. Therefore, the gains of mandating a single platform instead of having a few are minimal.

Second, a mandate might dull competition across platforms, slowing the pace of innovation. The history of kidney exchange has several examples of important innovations, such as the widespread use of altruistic pairs in non-simultaneous chains (Rees et al. 2009); global kidney exchange (Rees et al. 2017); voucher programs to encourage donation by giving future priority (Veale et al. 2017; Wall, Veale, and Melcher 2017); and improved operational procedures and matching algorithms (Anderson et al. 2015). These innovations weigh in favor of preserving a market structure with multiple platforms, especially since the efficiency cost of doing so is small.

Third, there is the issue of enforcement. Hospitals could dodge a ban on internal exchanges by reporting them as direct donations. But, even if such a ban were possible, a mandate could have unintended consequences. For instance, it could reduce hospital effort to recruit patients, or in extreme cases, move patients from being transplanted internally to not being transplanted at all. To see how this is possible, consider a model similar to the one in Section III with a single type of submission. Hospital $h$ has a flow of pairs, $\bar{q}^h$, of which it can internally transplant a fraction $t^h$. Through the platform, the per-transplant reward is $p$, which is also the fraction of submissions that get transplanted. The average cost of a submission is $k^h$ dollars, and transplants are valued at $v^h$ dollars.

First, consider hospital behavior without a mandate, as in Section III. Hospital decisions are driven by the (transplant-denominated) private cost,

$$C^h\left(q^h\right) = \begin{cases} \left(t^h + k^h/v^h\right) \cdot q^h & \text{if } q^h \leq \bar{q}^h \\ \infty & \text{if } q^h > \bar{q}^h. \end{cases}$$

Therefore, hospital $h$ submits all its pairs to the platform if $p > t^h + k^h/v^h$; otherwise, it transplants all pairs internally.

Now, consider a mandate that prohibits hospitals from performing internal transplants, but that cannot force hospitals to submit their pairs. The effect of this policy is to reduce the opportunity cost of a submission by $t^h$. This leads hospital $h$ to submit all pairs to the platform if $p > k^h/v^h$; otherwise, it performs *no kidney exchanges at all*, potentially reducing welfare.[49]

Although the stark predictions on hospital behavior are driven by the stylized assumptions, the example illustrates two points. First, enforceability is important. This goes beyond the ability to ban internal exchanges, especially if referrals and

[49] As an illustration, consider the case where the platform can transplant one-half of pairs (i.e., $p = 0.5$, approximately NKR's average product), hospital $h$ can internally match 30 percent of pairs (i.e., $t^h = 0.3$), and its value for a transplant is $v^h = \$50{,}000$. Then, without the mandate, hospital $h$ will submit all pairs to the platform if its cost per submission is less than \$10,000 (i.e., when $k^h < \left(p - t^h\right) \cdot v^h$). With the mandate, this threshold increases to \$25,000 (i.e., $p \cdot v^h$). Therefore, if hospital $h$'s submission cost is less than \$10,000, the mandate makes no difference. If hospital $h$'s submission cost is between \$10,000 and \$25,000, the mandate causes it to submit pairs to the platform that it would have otherwise transplanted internally, improving welfare. But, if hospital $h$'s submission cost is greater than \$25,000, the mandate results in the hospital doing no kidney exchange transplants whatsoever, *reducing* welfare.

patient recruiting are important.[50] Second, mandates can be complementary to other policies. In the example, subsidies to kidney exchange could lower $k^h$ enough to ensure that hospitals submit all pairs, which would allow the mandate to achieve first-best welfare.

Finally, a mandate poses political challenges: it would likely require an act of Congress, and a regulator would have to decide which platform(s) to authorize. All of the currently operating national platforms are private entities, and in fact, only UNOS, the smallest of the three, is formally linked with the US government.

## VI. Conclusion

Kidney exchange improves a patient's quality of life and extends life expectancy while reducing costs. We demonstrate that fragmentation in the US market results in an efficiency loss of between 30 and 63 percent of the roughly 800 kidney exchange transplants performed per year, implying a waste of hundreds of transplants per year.

The inefficiency arises from two standard market failures. First, platforms use inefficient mechanisms that do not reward hospitals according to the marginal products of their contributions. This induces hospitals to perform inefficient within-hospital matches, even if hospitals solely maximize the welfare of their own patients. Second, there are agency problems that make hospitals too sensitive to the costs of participating in kidney exchange platforms.[51] Our analysis shows that both market failures are likely important.

These findings have both short-term policy implications and broader implications for the design of kidney exchange markets. There are two short-term policy implications. First, there could be returns to existing platforms experimenting with point systems. Such systems can be implemented by individual platforms, and doing so will likely help them expand. Second, third-party payers should consider subsidizing kidney exchange at platforms. We argued that hospitals are likely responsive to the cost of participating in kidney exchange platforms, a behavior that leads to significant welfare loss. Subsidies from Medicare and private payers could mitigate this problem. Moreover, our analysis suggests that this two-pronged approach, which addresses the two market failures separately, is likely to be more robust than approaches that address both market failures simultaneously.

Consistent with our results, there are initiatives moving in the direction of these policy changes. The NKR recently started experimenting with a points system through their Center Liquidity Contribution Program. Some private insurers have

[50] Consider a variation of the example above in which hospital $h$ gets a flow $\left(1 - r^h\right)\bar{q}^h$ of pairs actively seeking kidney exchange, and a flow $r^h\bar{q}^h$ of pairs that can be recruited to participate in kidney exchange, but only if the hospital expends effort. Assume a mandate that sees all pairs that ultimately come to the hospital and requires them to be submitted. Let $p^h < k^h/v^h$, so that hospital $h$ prefers to submit as few pairs as possible. Without a mandate, hospital $h$ would have recruited $\bar{q}^h$ pairs and transplanted $t^h \cdot \bar{q}^h$ of them. With the mandate, the hospital does not recruit, and the platform arranges $p\left(1 - r^h\right)\bar{q}^h$ transplants. Hence, the mandate produces $\left(p - p \cdot r^h - t^h\right)\bar{q}^h$ more transplants, which is only positive if $1 - t^h/p > r^h$, that is, if the fraction of pairs that require recruitment is less than 1 minus the ratio of hospital and platform productivity.

[51] This decomposition of market failure sources is consistent with long-standing concerns of surgeons, insurers, platforms, and researchers, and with recent policy changes. Roth, Sönmez, and Ünver (2005) and Ashlagi and Roth (2014) recognized that hospitals may have incentives to match patients internally in static models. Surgeons and insurers have noted that it may be in the interest of insurers to subsidize exchanges and have proposed that they do so (Rees et al. 2012).

started covering the costs of participating in kidney exchange platforms. Our results indicate that there could be large gains from continuing to move in this direction. Further, all platforms could use data-driven rewards systems. Future research can contribute to the design and evaluation of these policies.

The kidney exchange market currently represents 15 percent of all living-donor kidney transplants. Therefore, kidney exchange is important in absolute terms, but is presently small relative to the total shortage of organs discussed in Becker and Elías (2007). There are ongoing efforts towards expanding the set of patients and donors participating in this market. For example, Mathur et al. (2018) reports on the efforts of the National Living Donor Assistance Center to reduce the financial burden on donors by reimbursing some of the costs of donating an organ; Sönmez, Ünver, and Yenmez (2018) proposes incentives for compatible pairs to join kidney exchange platforms; and Nikzad et al. (2017) reports on efforts for expanding kidney exchange to a global scale. The designs suggested by our analysis are complementary to these innovations because fixing the market failures identified above makes better use of the available donors.

Alternatively, a participation mandate might increase welfare under certain assumptions. And, the mandate can be simpler than the short-run policies that we discussed, even though it may be politically more difficult to implement. The United Kingdom, the Netherlands, and Canada have mandated participation in a single national program (Johnson et al. 2008, De Klerk et al. 2005, Malik and Cole 2014). One concern is that mandating participation in a single platform can reduce competition between platforms, which has arguably contributed to innovation. However, our estimates indicate that it would be close to efficient to have a few large platforms in the United States, because most of the potential efficiency gain would come from moving the market from individual hospitals to national platforms, rather than merging the largest platforms.

Although our study focuses on the US kidney exchange market, similar market failures arise when countries with independent kidney exchange platforms try to arrange cross-border exchanges. Such arrangements are being made between the Czech Republic and Austria (Böhmig et al. 2017) and between Italy, Spain, and Portugal.[52] In the latter case, these countries first matched their patients and donors internally before attempting to cooperate. This problem bears resemblance to fragmentation in the US system, suggesting that our insights on returns to scale and optimal rewards may be useful more broadly.

## REFERENCES

**Abdulkadiroğlu, Atila, and Kyle Bagwell.** 2013. "Trust, Reciprocity, and Favors in Cooperative Relationships." *American Economic Journal: Microeconomics* 5 (2): 213–59.

**Abraham, David J., Avrim Blum, and Tuomas Sandholm.** 2007. "Clearing Algorithms for Barter Exchange Markets: Enabling Nationwide Kidney Exchanges." *Proceedings of the 8th ACM Conference on Electronic Commerce*, 295–304.

**Agarwal, Nikhil, Itai Ashlagi, Eduardo Azevedo, Clayton Featherstone, and Ömer Karaduman.** 2018. "What Matters for the Productivity of Kidney Exchange?" *AEA Papers and Proceedings* 108: 334–40.

[52] See http://ekha.eu/blog/italy-and-spain-achieve-first-international-cross-kidney-transplantation/.

**Agarwal, Nikhil, Itai Ashlagi, Eduardo Azevedo, Clayton R. Featherstone, and Ömer Karaduman.** 2019. "Market Failure in Kidney Exchange: Dataset." *American Economic Review*. https://doi.org/10.1257/aer.20180771.

**American Society of Transplant Surgeons.** 2016. "Kidney Paired Donation: Community Perspectives and Best Practices." Unpublished.

**Anderson, Ross, Itai Ashlagi, David Gamarnik, and Yash Kanoria.** 2014. "A Dynamic Model of Barter Exchange." *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1925–1933.

**Anderson, Ross, Itai Ashlagi, David Gamarnik, Michael Rees, Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver.** 2015. "Kidney Exchange and the Alliance for Paired Donation: Operations Research Changes the Way Kidneys Are Transplanted." *Interfaces* 45 (1): 26–42.

**Arrow, Kenneth J.** 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review* 53 (5): 941–73.

**Ashlagi, Itai, Maximillien Burq, Patrick Jaillet, and Vahideh Manshadi.** 2019. "On Matching and Thickness in Heterogeneous Dynamic Markets." *Operations Research* 67 (4): 927–49.

**Ashlagi, Itai, and Alvin E. Roth.** 2014. "Free Riding and Participation in Large Scale, Multi-Hospital Kidney Exchange." *Theoretical Economics* 9 (3): 817–63.

**Bartelsman, Eric J., and Mark Doms.** 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature* 38 (3): 569–94.

**Becker, Gary S., and Julio Jorge Elías.** 2007. "Introducing Incentives in the Market for Live and Cadaveric Organ Donations." *Journal of Economic Perspectives* 21 (3): 3–24.

**Bingaman, A. W., F. H. Wright, Jr., M. Kapturczak, L. Shen, S. Vick, and C. L. Murphey.** 2012. "Single-Center Kidney Paired Donation: The Methodist San Antonio Experience." *American Journal of Transplantation* 12 (8): 2125–32.

**Böhmig, Georg A., Jirí Fronek, Antonij Slavcev, Gottfried F. Fischer, Gabriela Berlakovich, and Ondrej Viklicky.** 2017. "Czech-Austrian Kidney Paired Donation: First European Cross-Border Living Donor Kidney Exchange." *Transplant International* 30 (6): 638–39.

**Boiteux, Marcel.** 1956. "Sur la Gestion des Monopoles Publics Astreints à l'Équilibre Budgétaire." *Econometrica* 24 (1): 22–40.

**Chandra, Amitabh, and Jonathan Skinner.** 2012. "Technology Growth and Expenditure Growth in Health Care." *Journal of Economic Literature* 50 (3): 645–80.

**Chenery, Hollis B.** 1949. "Engineering Production Functions." *Quarterly Journal of Economics* 63 (4): 507–31.

**Clemens, Jeffrey, and Joshua D. Gottlieb.** 2014. "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review* 104 (4): 1320–49.

**Cowan, N., H. A. Gritsch, N. Nassiri, J. Sinacore, and J. Veale.** 2017. "Broken Chains and Reneging: A Review of 1748 Kidney Paired Donation Transplants." *American Journal of Transplantation* 17 (9): 2451–57.

**Danovitch, Gabriel M.** 2009. *Handbook of Kidney Transplantation.* Philadelphia: Lippincott Williams & Wilkins.

**De Klerk, Marry, Karin M. Keizer, Frans H. J. Claas, Marian Witvliet, Bernadette J. J. M. Haase-Kromwijk, and Willem Weimar.** 2005. "The Dutch National Living Donor Kidney Exchange Program." *American Journal of Transplantation* 5 (9): 2302–05.

**Dickerson, John P., Ariel D. Procaccia, and Tuomas Sandholm.** 2012. "Dynamic Matching via Weighted Myopia with Application to Kidney Exchange." Conference on Artificial Intelligence (AAAI).

**Duflo, Esther.** 2017. "Richard T. Ely Lecture: The Economist as Plumber." *American Economic Review* 107 (5): 1–26.

**Ellison, Blake.** 2014. "A Systematic Review of Kidney Paired Donation: Applying Lessons from Historic and Contemporary Case Studies to Improve the US Model." Unpublished.

**Farhi, Emmanuel, and Xavier Gabaix.** 2017. "Optimal Taxation with Behavioral Agents." Unpublished.

**Friedman, Eric J., Joseph Y. Halpern, and Ian Kash.** 2006. "Efficiency and Nash Equilibria in a Scrip System for P2P Networks." *Proceedings of the 7th ACM Conference on Electronic Commerce*, 140–49.

**Guo, Yingni, and Johannes Hörner.** 2015. "Dynamic Mechanisms with Money." Unpublished.

**Hajaj, Chen, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne.** 2015. "Strategy-Proof and Efficient Kidney Exchange Using a Credit Mechanism." *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 921–28.

**Hastie, Trevor, Robert Tibshirani, and Jerome Friedman.** 2001. *The Elements of Statistical Learning.* New York: Springer.

**Hauser, Christine, and Hugo Hopenhayn.** 2008. "Trading Favors: Optimal Exchange and Forgiveness." Collegio Carlo Alberto Notebooks 88.

**Held, P. J., F. McCormick, A. Ojo, and J. P. Roberts.** 2016. "A Cost-Benefit Analysis of Government Compensation of Kidney Donors." *American Journal of Transplantation* 16 (3): 877–85.

**Hsieh, Chang-Tai, and Peter J. Klenow.** 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4): 1403–48.

**Irwin, F. D., A. F. Bonagura, S. W. Crawford, and M. Foote.** 2012. "Kidney Paired Donation: A Payer Perspective." *American Journal of Transplantation* 12 (6): 1388–91.

**Jackson, Matthew O., and Hugo F. Sonnenschein.** 2007. "Overcoming Incentive Constraints by Linking Decisions." *Econometrica* 75 (1): 241–57.

**Johnson, Rachel J., Joanne E. Allen, Susan V. Fuggle, J. Andrew Bradley, and Chris Rudge.** 2008. "Early Experience of Paired Living Kidney Donation in the United Kingdom." *Transplantation* 86 (12): 1672–77.

**Kash, Ian A., Eric J. Friedman, and Joseph Y. Halpern.** 2007. "Optimizing Scrip Systems: Efficiency, Crashes, Hoarders, and Altruists." *Proceedings of the 8th ACM Conference on Electronic Commerce,* 305–15.

**Kolstad, Jonathan T.** 2013. "Information and Quality When Motivation Is Intrinsic: Evidence from Surgeon Report Cards." *American Economic Review* 103 (7): 2875–910.

**Kominers, Scott D., Alex Teytelboym, and Vincent P. Crawford.** 2017. "An Invitation to Market Design." *Oxford Review of Economic Policy* 33 (3): 541–71.

**Lerner, A. P.** 1934. "The Concept of Monopoly and the Measurement of Monopoly Power." *Review of Economic Studies* 1 (3): 157–75.

**Li, Shengwu.** 2017. "Ethics and Market Design." *Oxford Review of Economic Policy* 33 (4): 705–20.

**Malik, Shafi, and Edward Cole.** 2014. "Foundations and Principles of the Canadian Living Donor Paired Exchange Program." *Canadian Journal of Kidney Health and Disease* 1 (1): 6.

**Marschak, Jacob, and William H. Andrews.** 1944. "Random Simultaneous Equations and the Theory of Production." *Econometrica* 12 (3/4): 143–205.

**Massie, Allan B., Sommer E. Gentry, Robert A. Montgomery, Adam A. Bingaman, and Dorry L. Segev.** 2013. "Center-Level Utilization of Kidney Paired Donation." *American Journal of Transplantation* 13 (5): 1317–22.

**Mathur, Amit K., Jiawei Xing, David M. Dickinson, Patricia H. Warren, Kimberly A. Gifford, Barry A. Hong, Akinlolu Ojo, and Robert M. Merion.** 2018. "Return on Investment for Financial Assistance for Living Kidney Donors in the United States." *Clinical Transplantation* 32 (7): e13277.

**Möbius, Markus M.** 2001. "Trading Favors." Unpublished.

**Morris, Carl N.** 1983. "Parametric Empirical Bayes Inference: Theory and Applications." *Journal of the American Statistical Association* 78 (381): 47–55.

**National Kidney Registry.** 2016. "Member Center Terms & Conditions." http://www.kidneyregistry.org/docs/NKR_MC_Terms_Conditions.pdf.

**Nikzad, Afshin, Mohammad Akbarpour, Michael A. Rees, and Alvin E. Roth.** 2017. "Financing Transplant Costs of the Poor: A Dynamic Model of Global Kidney Exchange." Unpublished.

**Olley, G. Steven, and Ariel Pakes.** 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–97.

**Orandi, B. J., J. M. Garonzik-Wang, A. B. Massie, A. A. Zachary, J. R. Montgomery, K. J. Van Arendonk, M. D. Stegall, et al.** 2014. "Quantifying the Risk of Incompatible Kidney Transplantation: A Multicenter Study." *American Journal of Transplantation* 14 (7): 1573–80.

**Prendergast, Canice.** 2017. "The Allocation of Food to Food Banks." Unpublished.

**Ramsey, F. P.** 1927. "A Contribution to the Theory of Taxation." *Economic Journal* 37 (145): 47–61.

**Rees, M. A., T. B. Dunn, C. S. Kuhr, C. L. Marsh, J. Rogers, S. E. Rees, A. Cicero, et al.** 2017. "Kidney Exchange to Overcome Financial Barriers to Kidney Transplantation." *American Journal of Transplantation* 17 (3): 782–90.

**Rees, M. A., J. E. Kopke, R. P. Pelletier, D. L. Segev, M. E. Rutter, A. J. Fabrega, J. Rogers, et al.** 2009. "A Nonsimultaneous, Extended, Altruistic-Donor Chain." *New England Journal of Medicine* 360 (11): 1096–1101.

**Rees, M. A., M. A. Schnitzler, E. Y. Zavala, J. A. Cutler, A. E. Roth, F. D. Irwin, S. W. Crawford, and A. B. Leichtman.** 2012. "Call to Develop a Standard Acquisition Charge Model for Kidney Paired Donation." *American Journal of Transplantation* 12 (6): 1392–97.

**Robert, Christian P., and George Casella.** 2004. Monte Carlo Statistical Methods. New York: Springer.

**Rochet, Jean-Charles, and Jean Tirole.** 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990–1029.

**Roth, Alvin E.** 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives* 21 (3): 37–58.

**Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver.** 2004. "Kidney Exchange." *Quarterly Journal of Economics* 119 (2): 457–88.

**Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver.** 2005. "Transplant Center Incentives in Kidney Exchange." Unpublished.

**Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver.** 2007. "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences." *American Economic Review* 97 (3): 828–51.

**Sönmez, Tayfun, and M. Utku Ünver.** 2013. "Market Design for Kidney Exchange." In *The Handbook of Market Design*, edited by Nir Vulkan, Alvin E. Roth, and Zvika Neeman, 93–137. Oxford: Oxford University Press.

**Sönmez, Tayfun, M. Utku Ünver, and M. Bumin Yenmez.** 2018. "Incentivized Kidney Exchange." Unpublished.

**Toulis, Panos, and David C. Parkes.** 2015. "Design and Analysis of Multi-Hospital Kidney Exchange Mechanisms Using Random Graphs." *Games and Economic Behavior* 91: 360–82.

**United States Renal Data System.** 2013. *2013 USRDS Annual Data Report: Epidemiology of Kidney Disease in the United States.* Vol. 2. Bethesda, MD: National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases.

**United States Renal Data System.** 2016. *2016 USRDS Annual Data Report: Epidemiology of Kidney Disease in the United States.* Vol. 2. Bethesda, MD: National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases.

**Veale, Jeffrey, Alexander Capron, Nima Nassiri, Gabriel Danovitch, H. Gritsch, Amy Waterman, Joseph Del Pizzo, et al.** 2017. "Vouchers for Future Kidney Transplants to Overcome 'Chronological Incompatibility' between Living Donors and Recipients." *Transplantation* 101 (9): 2115–19.

**Wall, Anji, Jeffrey Veale, and Marc Melcher.** 2017. "Advanced Donation Programs and Deceased Donor-Initiated Chains—2 Innovations in Kidney Paired Donation." *Transplantation* 101 (12): 2818–24.

**Walters, A. A.** 1963. "Production and Cost Functions: An Econometric Survey." *Econometrica* 31 (1/2): 1–66.

**Weyl, E. Glen.** 2010. "A Price Theory of Multi-Sided Platforms." *American Economic Review* 100 (4): 1642–72.

**Wolfe, Robert A., Valarie B. Ashby, Edgar L. Milford, Akinlolu O. Ojo, Robert E. Ettenger, Lawrence Y. C. Agodoa, Philip J. Held, and Friedrich K. Port.** 1999. "Comparison of Mortality in All Patients on Dialysis, Patients on Dialysis Awaiting Transplantation, and Recipients of a First Cadaveric Transplant." *New England Journal of Medicine* 341 (23): 1725–30.