OPERATIONS RESEARCH
Vol. 56, No. 6, November–December 2008, pp. 1393–1410
ISSN 0030-364X | EISSN 1526-5463 | 08 | 5606 | 1393

DOI 10.1287/opre.1080.0648

Estimating the Patient's Price of Privacy in Liver Transplantation

Burhaneddin Sandıkçı
The University of Chicago Booth School of Business, Chicago, Illinois 60637, burhan@uchicago.edu

Lisa M. Maillart, Andrew J. Schaefer
Department of Industrial Engineering, University of Pittsburgh, Pittsburgh, Pennsylvania 15260
{maillart@pitt.edu, schaefer@ie.pitt.edu}

Oguzhan Alagoz
Department of Industrial and Systems Engineering, University of Wisconsin, Madison, Wisconsin 53706,
alagoz@engr.wisc.edu

Mark S. Roberts
Department of General Internal Medicine, University of Pittsburgh, Pittsburgh, Pennsylvania 15213,
robertsm@upmc.edu

In the United States, patients with end-stage liver disease must join a waiting list to be eligible for cadaveric liver transplantation. Due to privacy concerns, the details of the composition of this waiting list are not publicly available. This paper considers the benefits associated with creating a more transparent waiting list. We study these benefits by modeling the organ accept/reject decision faced by these patients as a Markov decision process in which the state of the process is described by patient health, quality of the offered liver, and a measure of the rank of the patient in the waiting list. We prove conditions under which there exist structured optimal solutions, such as monotone value functions and control-limit optimal policies. We define the concept of the patient's price of privacy, namely, the number of expected life days lost due to the lack of complete waiting list information. We conduct extensive numerical studies based on clinical data, which indicate that this price of privacy is typically on the order of 5% of the optimal solution value.

Subject classifications: dynamic programming/optimal control: applications, Markov; health care: treatment.
Area of review: Special Issue on Operations Research in Health Care.
History: Received January 2007; revisions received July 2007, October 2007, December 2007; accepted January 2008.

1. Introduction

End-stage liver disease (ESLD), which includes diseases such as primary biliary cirrhosis and hepatitis, is the 12th leading cause of death in the United States (National Center for Health Statistics (NCHS) 2006) in large part because transplantation is the only available therapy for ESLD patients. The vast majority of ESLD patients join a waiting list of patients who are eligible for transplantation from cadaveric donors. As seen in Figure 1, the disparity between the demand for and supply of cadaveric livers is large and growing, which results in a significant number of patient deaths while waiting. Indeed, there are currently over 17,000 patients on the liver transplant waiting list and, over the past 12 years, this list has grown by approximately 1,000 patients per year (United Network for Organ Sharing (UNOS) 2006a). All of these facts motivate the need for research on better management of this scarce resource.

When a cadaveric liver becomes available, UNOS, the organization that administers the organ allocation activities in the United States, offers the liver to patients using an allocation mechanism. For each cadaveric liver, this mechanism assigns priorities to patients based on disease severity of the patient, geography, and physiologic compatibility between the donor and the potential recipient (e.g., blood type and size). UNOS uses a patient's total waiting time at her current level of health or worse as a final tie breaker among otherwise identical patients (see UNOS 2006b for more details on the matching mechanism). Often multiple potential recipients are notified concurrently because UNOS considers the final decision of whether or not to use the offered liver to be "the prerogative of the transplant surgeon and/or the physician responsible for the care of the patient" (UNOS 2006b, p. 3-26). As Howard (2002) reports, surgeons reject low-quality organs for healthy patients in the hope that they may receive a better organ offer in the future. Several characteristics of the donor may affect the perceived quality of the donated organ such as length of intensive care unit stay and antecedents of hypertension (Cuende et al. 2005); existence and degree of steatosis (Salizzoni et al. 2003); race, height, and involvement in a cerebravascular accident (Feng et al. 2006); and age, blood type, and gender (Roberts et al. 2004). Indeed, despite the scarcity of donated organs, almost half of the liver offers are rejected by the first surgeon to whom the

Figure 1. Recent trends in U.S. liver transplantation (1995–2004).

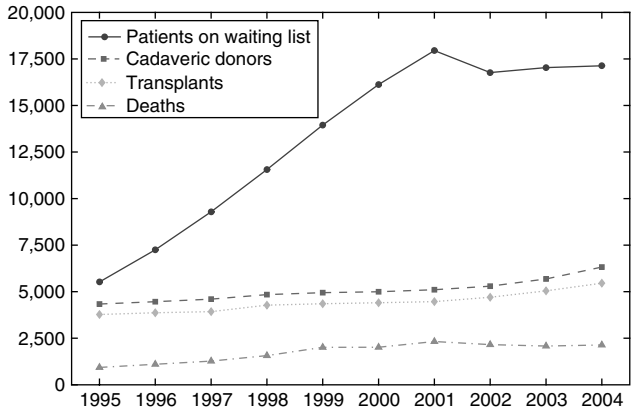

Note. Data source: Department of Health and Human Services (HHS 2005).

offer is made (Howard 2002). The optimization of this accept/reject decision, within the confines of the current allocation system, is the focus of this paper. The decision-making process is assumed to be joint between the patient and an agent (such as the patient's physician and/or surgeon) who acts in the patient's best interest. However, for expositional simplicity, we refer to the decision maker as the patient throughout the rest of the paper.

Several researchers consider the organ accept/reject problem from an individual patient's perspective (Ahn and Hornberger 1996; Alagoz et al. 2004, 2007a, 2007b; David and Yechiali 1985; Hornberger and Ahn 1997; Howard 2002), or from the society's perspective (David and Yechiali 1990, 1995; Righter 1989; Roth et al. 2004; Su and Zenios 2005; Zenios 1999, 2002; Zenios et al. 1999, 2000), or from a joint perspective (Su and Zenios 2004, 2005, 2006). Most of this literature makes unrealistic assumptions such as patient health does not change over time, each organ may be offered to at most one patient, new patients do not arrive, listed patients do not die, all patients are homogeneous, organ quality does not deteriorate, offered organs cannot be declined, and all patients have the same pretransplant life expectancy. We refer the reader to Alagoz (2004) for a detailed discussion of the organ transplantation literature. Of this body of work, the most relevant is Alagoz et al. (2007a). They present a Markov decision process (MDP) model in which the state is described by patient health and organ quality. For each possible state, provided an offer is made, the patient chooses to either accept or reject the offer so as to maximize her total expected reward. Their approach captures the effects of the waiting list *implicitly* through the organ arrival probabilities, which are assumed to be a function of patient health. Under the current liver allocation policy, however, the frequency and the quality of liver offers made to an individual patient are significantly affected by the physiology and the geographic location of the other patients on the waiting list. Therefore, the model analyzed in this paper captures the effects of the waiting list in a more *explicit* fashion and lends insight into how the composition of the list impacts optimal accept/reject decisions.

Historically, much of the relevant information about the patients on the waiting list has been hidden due to privacy concerns. Although this practice ensures some level of confidentiality, it also forces patients to make accept/reject decisions with incomplete information. To alleviate this problem, UNOS now publishes coarse (yet still incomplete) descriptions of the waiting list on their website (UNOS 2006a). For example, one can learn the number of patients in specific ranges of disease severity in a specific geographic area. Anything less than a complete description of the waiting list, however, still results in some loss to the patient because this lack of knowledge may result in suboptimal decisions. We call this loss due to incomplete rank knowledge the patient's *price of privacy.* This quantity may also be interpreted as the value of obtaining rank information. We leverage our model analysis to provide quantitative estimates of this price of privacy in terms of overall life expectancy. That is, we do not advocate any particular change in the privacy policy; rather, we wish to quantify the costs of this privacy to the patient.

We emphasize that revealing the ranks of the patients, in general, is not equivalent to revealing their identities; and, in this paper, we use the term "privacy" to refer to the former as opposed to its general association with the latter. However, we realize that if there are very few patients registered in a particular geographic area, then rank revelation may inadvertently reveal patient identities. One possible remedy is to design a system that sends private signals to potential recipients that includes their own rank information only.

The patients registered at transplant centers that have a large market share in a geographic location may know (through their physicians) more than the revealed coarse descriptions about the waiting list. They may even be able to identify the precise ranks of the patients registered in the same location. Our model provides an accurate representation of such situations.

The public revelation of the waiting list information naturally gives rise to a gaming environment, where each patient, when making their own decisions, has to consider the possible decisions of other patients on the waiting list. To the best of our knowledge, the only organ transplantation paper to consider such game-theoretic aspects is Su and Zenios (2004). They model the kidney transplant waiting system as an M/M/1 queue with exponential reneging that represents patient death. They assume a homogeneous patient population, which enables a detailed competitive equilibrium analysis, and discuss the effects of different queuing disciplines such as first-come-first-served and last-come-first-served on the system performance. Unfortunately, such a queuing model is inappropriate for liver

transplantation because patient priorities evolve dynamically with their health, and different livers may induce different priorities due to differences in geography and blood type. Furthermore, acknowledging the heterogeneity of the patients leads to an intractable asymmetric multiplayer nonzero-sum stochastic game. We refer the interested reader to Vieille (2002) and Neyman and Sorin (2003) for a recent review of the state-of-the-art in stochastic games.

Although an analysis capturing the competition among patients and clinical realism simultaneously is desirable, we focus, in this paper, on modeling the patient's perspective by amplifying clinical realism and suppressing the competition among patients. This approach comes with some limitations: (1) the estimates quantified in this paper refer to the individual gain of a single patient who is provided with full rank information; (2) the current analysis does not modify the decisions of other patients who are assumed to act without rank information; and (3) our approach may underestimate or overestimate the true price of privacy, therefore it may be viewed as a heuristic for quantifying the true price of privacy.

Several researchers have approached various aspects of the organ transplantation problem using queuing models (e.g., Zenios 1999; Zenios et al. 2000; Su and Zenios 2004, 2006). Most of the queuing models are concerned with kidney transplantation. We refer to Zenios (2004) for a recent review of queuing-based models in kidney transplantation. However, the liver transplant waiting list is much more complex than a simple queue (UNOS 2006b, Howard 2001) because the priorities assigned to patients are a function of geography and health, and this fact renders queuing approaches to liver transplantation inappropriate.

A related stream of research analyzes equilibria in queueing models. We refer the interested reader to Hassin and Haviv (2003) and Altman (2005) for recent surveys of this literature. This stream of research focuses on questions like when to join a queue, which queue to join when there are multiple queues, and what priority level to purchase when different priorities are allowed. All of these questions are related to customers' decisions at the time of their arrivals. However, in our case, patients are prioritized at the time of an organ arrival according to the liver allocation policy. Therefore, they do not have any choice at the time of their arrival but rather a prerogative to refuse an organ (service) offer. Furthermore, we also allow patient priorities to change over time and we explicitly use the rank information of the patient to make this decision. A methodologically similar paper that models the decision making (although the customers are not the decision makers) using rank information is Swani et al. (2001), which formulates an MDP model to find optimal replacement policies for a single motion picture exhibitor ignoring the competition between theater chains. They contend with providing a numerical analysis of the model without attempting any structural results. In their competitive equilibrium analysis within the kidney allocation system, Su and Zenios (2004) also use the rank information of the patients to characterize the rank-dependent threshold policies. However, these characterizations are heavily influenced by the assumption of homogeneous patients, which we cannot justify in the context of this paper.

To summarize, we are concerned with three different system scenarios: (a) the system in which every patient has partial rank information as in the current allocation system and behaves optimally using this information, (b) the system analyzed in this paper, in which only one special patient has full rank information and all other patients behave as they do now, and (c) a proposed system, in which every patient has full rank information and acts optimally. The analysis of system (a) is still an open issue. The best available representation of system (a) is the model of Alagoz et al. (2007a). We compare our model (system (b)) to the model of Alagoz et al. (2007a) to obtain an estimate of the true price of privacy, which ideally would be computed by comparing system (a) to the proposed system (c). If the waiting list information is provided to everyone as in the proposed system, then there may be a complete shift of equilibrium, but our model falls short of identifying this new equilibrium.

The rest of this paper is organized as follows. Section 2 presents the Markov decision process model formulation, which expands the state space of the Alagoz et al. (2007a) model to include the patient's rank. Section 3 establishes analytical conditions that yield desirable structural results for the optimal policy. Section 4 discusses the price of privacy in more detail and presents the results of clinically driven numerical experiments. Finally, the paper concludes in §5 by summarizing the contributions and pointing out possible extensions.

2. The MDP Model

We start this section by assuming that patients are self-interested agents and do not explicitly consider the possible actions other candidates may take in making their own decisions. Such an assumption is indispensable for analytical and computational tractability of the resulting model. Consider an ESLD patient who must decide (with her physician and/or surgeon) whether to accept or reject a liver offered for transplantation so as to maximize her total expected discounted reward. We assume that the patient makes this decision at discrete time periods. If a liver is not offered in a particular period, then the patient is forced to "wait" until the next period. The nontrivial decision to be optimized is when there is a liver offered. If the patient chooses to wait, then she accrues an intermediate reward which is a function of her current health status, and faces the same problem at the next time period, provided she lives. If, on the other hand, the patient chooses to "accept" the offer, then she receives a lump-sum terminal reward (e.g., the expected discounted posttransplant survival or quality-adjusted survival). This terminal reward is a function of the patient's current health status as well as the quality of the accepted

liver. By choosing to transplant an offered liver, the patient terminates the process.

Among the organ acceptance models that consider the effect of the waiting list, Alagoz et al. (2007a) implicitly models the waiting list through the organ arrival probabilities. At the other extreme, a *fully explicit* model of the waiting list would track the health, location, blood type, and waiting time of all patients in the list. Our approach balances the additional complexity associated with incorporating information about the waiting list with practical considerations such as model calibration and solution time.

Furthermore, under the current liver allocation mechanism, the priorities assigned to patients are not only determined by the characteristics of the patients in the waiting list, but also by the characteristics of the donated liver. Therefore, even in a hypothetical environment in which all the patients' characteristics are constant and no new patients arrive, a currently listed patient may be assigned different priorities for different livers. For example, the priority of a patient for a liver donated in the same geographic service area as she is registered may be significantly different than her priority for an identical liver donated in a different service area. Furthermore, even for two different livers donated in the same geographic area, the patient may be assigned different priorities depending on her blood type compatibility with the donated livers. The current allocation mechanism partitions the United States into approximately 60 geographic areas of varying sizes, population densities, donation rates, and realized transplants, each of which is served by a single Organ Procurement Organization (OPO), and assigns three blood type compatibility levels to patients. Incorporating these two factors alone to model the rank of a patient over all possible livers would lead to a dramatic increase in the size of the state space.

As a compromise, we define a patient's "rank" as a scalar, namely, *the rank of the patient among all patients' expected priorities*, where the expectation is taken over all possible livers. To illustrate this definition, consider a hypothetical example with four patients and two livers. The characteristics of these patients are given in Table 1. Assume that both livers are blood type A; however, Liver 1 is procured in OPO 1, whereas Liver 2 is procured in OPO 2. All else held equal, the order of patient priorities based on the current UNOS policy, from first to last, would be a-b-c-d and c-a-b-d for Livers 1 and 2, respectively. The expected priorities and the rank of expected priorities of these patients are as shown in the respective columns of Table 1. For the sake of exposition, from now on, we use "rank" to mean the rank of the patient among all patients' expected priorities.

Table 1. Patient characteristics for the hypothetical example.

Patient	Geographic area	Blood type	Expected priority	Rank of expected priorities
a	1	A	1.5	1
b	1	AB	2.5	3
c	2	A	2	2
d	3	O	4	4

We define a state, s, of the MDP model to be composed of the triplet (h, ℓ, k), where h is the patient's health status, ℓ is the quality of the liver being offered, and k is the rank of the patient. We assume that the components of the state can take on the following values: $h \in \Omega = \{1, 2, \ldots, H\}$, where the quality of health is decreasing as h increases; $\ell \in \Phi = \{1, 2, \ldots, L+1\}$, where the quality of the liver is decreasing as ℓ increases, and $L+1$ represents no liver being offered; and $k \in \Psi = \{1, 2, \ldots\}$, where the patient moves further from the top of the waiting list as k increases.

For convenience, we add two absorbing states, Δ and ∇, to represent the dead and transplanted states, respectively. Therefore, the state space of the model becomes

$$\mathcal{S} = \mathcal{S}' \cup \{\Delta\} \cup \{\nabla\},$$

where

$$\mathcal{S}' = \{(h, \ell, k) \mid h \in \Omega, \ell \in \Phi, k \in \Psi\}.$$

The set of possible actions in state $s \in \mathcal{S}$ is

$$\mathcal{A}_s = \begin{cases} \{W\} & \text{if } s = \Delta \text{ or } s = \nabla \\ & \text{or } s = \{(h, \ell, k) \mid \ell = L+1\}, \\ \{T, W\} & \text{otherwise,} \end{cases}$$

where W stands for rejecting the offer and waiting for one more period and T stands for accepting the offer and transplanting. The immediate rewards for each possible state-action pair (s, a), such that $s \in \mathcal{S}$ and $a \in \mathcal{A}_s$, is given by

$$r(s, a) = \begin{cases} 0 & \text{if } s = \Delta \text{ or } s = \nabla, \\ r_W(h) & \text{if } s \in \mathcal{S}' \text{ and } a = W, \\ r_T(h, \ell) & \text{if } a = T. \end{cases}$$

The patient does not accumulate any additional rewards once she is dead or transplanted. Furthermore, a pretransplant patient who chooses to wait for one more period receives an intermediate reward of $r_W(h)$, which is a function of the patient's health status only. Finally, if the patient chooses to transplant, she receives a lump-sum reward of $r_T(h, \ell)$, which is a function of both the patient's health status and the quality of the liver being offered.

The final component of the MDP model is the transition probabilities. When the patient chooses the transplant action, she transitions to the transplanted state, ∇, with probability one (i.e., $P\{\nabla \mid s, T\} = 1$ for all $s \in \mathcal{S}$ such that

$T \in \mathcal{A}_s$). If, on the other hand, the patient chooses to wait, then the transition probabilities are defined as

$$P\{s' \mid s=(h,\ell,k), W\} = \begin{cases} p(h',\ell',k' \mid h,\ell,k) & \text{if } s'=(h',\ell',k'), \\ 1-\sum_{s'\in\mathcal{S}'} p(h',\ell',k' \mid h,\ell,k) & \text{if } s'=\Delta, \\ 0 & \text{if } s'=\nabla. \end{cases}$$

We assume that the transition probabilities and rewards are stationary. We further assume that $p(h',\ell',k' \mid h,\ell,k) = \mathcal{H}\{h' \mid h\}\cdot\mathcal{K}\{k' \mid k\}\cdot\mathcal{L}\{\ell' \mid k'\}$ for all $h, h' \in \Omega$, $k, k' \in \Psi$, and $\ell, \ell' \in \Phi$, where $\mathcal{H}\{h' \mid h\}$ is the probability that the patient's health status will be h' at time $t+1$ given that her health at time t is h, $\mathcal{K}\{k' \mid k\}$ is the probability that the patient's rank will be k' at time $t+1$ given that her rank at time t is k, and $\mathcal{L}\{\ell' \mid k'\}$ is the probability that the patient will be offered an organ of quality ℓ' at time $t+1$ given that her rank at time $t+1$ is k'. We assume that the transitions among health states and the transitions among rank states are independent. Admittedly, the patient's rank at time $t+1$ depends on her health at time $t+1$. However, for analytical and computational tractability, we choose to include the dependency on her rank at time t only. Because patient rank is the primary indicator of patient health, an additional dependency on health would not change the values of these probabilities significantly. Finally, we define the *rank transition probability matrix*, $\mathcal{K}$, as $\mathcal{K} = [\mathcal{K}\{k' \mid k\}]_{\forall k,k'\in\Psi}$, the *health probability matrix*, $\mathcal{H}$, as $\mathcal{H} = [\mathcal{H}\{h' \mid h\}]_{\forall h,h'\in\Omega}$, and the *liver offer probability matrix*, $\mathcal{L}$, as $\mathcal{L} = [\mathcal{L}\{\ell' \mid k'\}]_{\forall k'\in\Psi,\forall\ell'\in\Phi}$. We emphasize that $\sum_{k'\in\Psi}\mathcal{K}\{k' \mid k\} = \sum_{\ell'\in\Phi}\mathcal{L}\{\ell' \mid k'\} = 1$ for all $k, k' \in \Psi$. We interpret $1-\sum_{h'\in\Omega}\mathcal{H}\{h' \mid h\}$ as the probability of dying when the patient's health is h.

Given the discount rate $\lambda \in [0,1)$, the optimal solution to this problem can be obtained by solving the Bellman optimality equations (Puterman 1994)

$$v(h,L+1,k) = r_W(h) + \lambda \sum_{(h',\ell',k')} \mathcal{H}\{h' \mid h\}\cdot\mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h',\ell',k') \quad \forall h\in\Omega,\ \forall k\in\Psi, \quad (1)$$

and

$$v(h,\ell,k) = \max\left\{ r_T(h,\ell), r_W(h) + \lambda \cdot \sum_{(h',\ell',k')} \mathcal{H}\{h' \mid h\}\mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h',\ell',k') \right\} \quad \forall h\in\Omega,\ \forall k\in\Psi,\ \ell\in\Phi\setminus\{L+1\}. \quad (2)$$

The value associated with the absorbing states of death and posttransplant, Δ and ∇, is zero by construction and these states are, therefore, excluded from Equations (1) and (2).

3. Structural Properties

In this section, we establish several structural properties of the MDP model formulated in §2. Specifically, we identify conditions on the parameters that guarantee structured value functions and optimal policies. In addition to their analytical elegance, such results may provide deeper insight into the overall problem and help devise computationally faster solution approaches.

The following assumptions hold throughout.

Assumption 1 (AS1). *$r_W(h)$ is nonincreasing in h.*

Assumption 2 (AS2). *$r_T(h,\ell)$ is nonincreasing in both h and ℓ.*

Assumption 1 implies that the intermediate reward of waiting does not increase as the patient deteriorates. Similarly, (AS2) implies that the posttransplant reward does not increase as the patient deteriorates and/or the quality of the liver degrades.

Proposition 1 establishes the intuitive fact that it is always better to be offered a higher-quality organ. The proof is obvious and is therefore omitted.

Proposition 1. *$v(h,\ell,k)$ is monotonically nonincreasing in ℓ for any $h\in\Omega$ and $k\in\Psi$.*

Similar to Alagoz et al. (2007a), we define a *liver-based control-limit optimal policy* to be a policy among the optimal policies that, for a given health state h, and rank k, distinguishes a critical liver state ℓ^* and prescribes "transplant" for all livers $\ell \leqslant \ell^*$ and "wait" for all livers $\ell > \ell^*$.

Theorem 1. *There exists a liver-based control-limit optimal policy for all $h\in\Omega$ and $k\in\Psi$.*

Proof. For any given liver quality $\ell < L$, it suffices to prove that if $a^*(h,\ell+1,k) = T$, then $a^*(h,\ell,k) = T$ for all $h\in\Omega$ and $k\in\Psi$. For any given $h\in\Omega$ and $k\in\Psi$, if $a^*(h,\ell+1,k) = T$, then

$$\begin{aligned} v(h,\ell+1,k) &= r_T(h,\ell+1) \\ &\geqslant r_W(h) + \lambda \sum_{h'}\sum_{k'}\sum_{\ell'} \mathcal{H}\{h' \mid h\}\cdot\mathcal{K}\{k' \mid k\} \\ &\quad \cdot \mathcal{L}\{\ell' \mid k'\}\cdot v(h',k',\ell'). \end{aligned}$$

Because $v(h,\ell,k) = \max\{r_T(h,\ell), r_W(h) + \lambda\sum_{h'}\sum_{k'}\sum_{\ell'}\mathcal{H}\{h' \mid h\}\cdot\mathcal{K}\{k' \mid k\}\cdot\mathcal{L}\{\ell' \mid k'\}\cdot v(h',k',\ell')\}$ and AS2 implies $r_T(h,\ell) \geqslant r_T(h,\ell+1)$ for $\ell = 1,\ldots,L-1$, we find $v(h,\ell,k) = r_T(h,\ell)$. Therefore, $a^*(h,\ell,k) = T$ for any $h\in\Omega$ and $k\in\Psi$. □

Next, we introduce the concept of a Column-wise Concave with the maximum element of each column on the Diagonal (CCD) matrix, which facilitates the results associated with the rank component (k) of the MDP model. Several properties of CCD matrices are given in the appendix.

DEFINITION 1. An $n \times n$ stochastic matrix $\mathcal{P}$ is called CCD (Column-wise Concave with the maximum element of each column on the Diagonal) if, for $i = 1, \ldots, n-1$, it satisfies

$$\text{(i)}\ \ \mathcal{P}\{j \mid i\} \geqslant \mathcal{P}\{j \mid i+1\} \quad \text{for } j = 1, 2, \ldots, i, \tag{3}$$

$$\text{(ii)}\ \ \mathcal{P}\{j \mid i\} \leqslant \mathcal{P}\{j \mid i+1\} \quad \text{for } j = i+1, \ldots, n. \tag{4}$$

This definition implies that within each column of a CCD matrix, the values are nondecreasing up to (and including) the diagonal element, and nonincreasing after the diagonal element. As a consequence, a necessary (but not sufficient) condition for a stochastic matrix to be CCD is to have the maximum value within each column at the diagonal entry.

In the context of our MDP model, a CCD rank transition probability matrix implies that the likelihood that a patient with rank i moves to a better rank $j_1 < i$ in the next period is at least as large as the likelihood of moving to the same rank j_1 from a rank that is further down the list. On the other hand, the likelihood that a patient with rank i moves to a worse rank $j_2 > i$ in the next period is no more than the likelihood of moving to the same rank j_2 if the patient is further down the list than i. Proposition 2 presents an inequality for the row differences of a CCD matrix, which is used in proving Theorem 2. The proof is given in the appendix.

PROPOSITION 2. *Let $f\colon \mathbb{R} \to \mathbb{R}_+$ and $g\colon \mathbb{R} \to \mathbb{R}$ be two functions. If $g(\cdot)$ is nonincreasing and $\mathcal{P}$ is a CCD transition probability matrix, then the following hold*:

$$\text{(i)}\ \sum_{j \leqslant k} [\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}] f(j) g(j) \geqslant g(k) \cdot \sum_{j \leqslant k} [\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}] f(j), \tag{5}$$

$$\text{(ii)}\ \sum_{j > k} [\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}] f(j) g(j) \geqslant g(k) \cdot \sum_{j > k} [\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}] f(j). \tag{6}$$

An immediate result of Proposition 2 is Corollary 1.

COROLLARY 1. *If $\mathcal{H}$ is a CCD matrix and $v(h, \ell, k)$ is nonincreasing in k for any $h \in \Omega$ and $\ell \in \Phi$, then the following hold*:

$$\text{(i)}\ \sum_{\ell'} \sum_{k' \leqslant k} [\mathcal{H}\{k' \mid k\} - \mathcal{H}\{k' \mid k+1\}] \mathcal{L}\{\ell' \mid k'\} v(h', \ell', k') \geqslant \sum_{\ell'} v(h', \ell', k) \sum_{k' \leqslant k} [\mathcal{H}\{k' \mid k\} - \mathcal{H}\{k' \mid k+1\}] \mathcal{L}\{\ell' \mid k'\},$$

$$\text{(ii)}\ \sum_{\ell'} \sum_{k' > k} [\mathcal{H}\{k' \mid k\} - \mathcal{H}\{k' \mid k+1\}] \mathcal{L}\{\ell' \mid k'\} v(h', \ell', k') \geqslant \sum_{\ell'} v(h', \ell', k) \sum_{k' > k} [\mathcal{H}\{k' \mid k\} - \mathcal{H}\{k' \mid k+1\}] \mathcal{L}\{\ell' \mid k'\}.$$

Lemma 1 states that, for a given set of nonincreasing weights, the nonnegative linear combination of the values obtained by taking the difference of any two successive rows of a CCD matrix is always nonnegative. The proof of this lemma follows immediately from the fact that CCD matrices have the increasing failure rate (IFR) property (see Proposition 4 in the appendix and Lemma 4.7.2 of Puterman 1994).

LEMMA 1. *Let $\{z_j\}_{j=1}^{\infty}$ be a given sequence of nonnegative and nonincreasing numbers. If $\mathcal{P}$ is a CCD transition probability matrix, then for any given k,*

$$\sum_j [\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}] z_j \geqslant 0.$$

Theorem 2 states that a patient's maximum expected total discounted reward does not increase in her rank. The condition on $\mathcal{L}$ simply states that a patient's chance of receiving a liver offer does not increase in her rank.

THEOREM 2. *If $\mathcal{H}$ is CCD and $\mathcal{L}\{\ell \mid k\}$ is monotonically nonincreasing in $k \in \Psi$ for all $\ell \neq L+1$, then $v(h, \ell, k)$ is monotonically nonincreasing in $k \in \Psi$ for any $h \in \Omega$ and $\ell \in \Phi$.*

PROOF. By induction on the steps of the value iteration algorithm.

We prove the theorem for $\ell \neq L+1$ and note that the proof for $\ell = L+1$ follows similarly.

If we can show that the value functions at each iteration of the value iteration algorithm are monotonically nonincreasing in $k \in \Psi$ for given $h \in \Omega$ and $\ell \in \Phi$, then the result holds by the convergence of value iteration. Let $v^i(h, \ell, k)$ be the value associated with state $(h, \ell, k) \in \mathcal{S}$ at the ith iteration of the value iteration algorithm and assume, without loss of generality, that the algorithm starts with a value of zero for each state, i.e., $v^0(h, \ell, k) = 0$ for all $(h, \ell, k) \in \mathcal{S}$.

It is clear that $v^1(h, \ell, k)$ is constant and therefore nonincreasing in $k \in \Psi$ for all $h \in \Omega$ and $\ell \in \Phi$.

Next, assume, as the induction hypothesis, that for a given $h \in \Omega$ and $\ell \in \Phi$, $v^i(h, \ell, k) \geqslant v^i(h, \ell, k+1)$ for all $k \in \Psi$ for iterations $i = 2, \ldots, n$.

By Equation (2),

$$v^{n+1}(h, \ell, k) = \max\left\{ r_T(h, \ell), r_W(h) + \lambda \sum_{h'} \sum_{k'} \sum_{\ell'} \mathcal{H}\{h' \mid h\} \cdot \mathcal{H}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \right\} \tag{7}$$

and

$$v^{n+1}(h, \ell, k+1) = \max\left\{ r_T(h, \ell), r_W(h) + \lambda \sum_{h'} \sum_{k'} \sum_{\ell'} \mathcal{H}\{h' \mid h\} \cdot \mathcal{H}\{k' \mid k+1\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \right\}. \tag{8}$$

If $a^{n+1}(h, \ell, k+1) = T$, then $v^{n+1}(h, \ell, k+1) = r_T(h, \ell) \leqslant v^{n+1}(h, \ell, k)$.

If $a^{n+1}(h,\ell,k+1)=W$, then by Equations (7) and (8),

$$\begin{aligned}
&v^{n+1}(h,\ell,k)-v^{n+1}(h,\ell,k+1)\\
&\geqslant\lambda\sum_{h'}\mathcal{H}\{h'\,|\,h\}\Bigg\{\sum_{k'}\sum_{\ell'}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\\
&\qquad\cdot\mathcal{L}\{\ell'\,|\,k'\}v^n(h',\ell',k')\Bigg\}\\
&=\lambda\sum_{h'}\mathcal{H}\{h'\,|\,h\}\Bigg\{\sum_{\ell'}\sum_{k'\leqslant k}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\mathcal{L}\{\ell'\,|\,k'\}\\
&\qquad\cdot v^n(h',\ell',k')+\sum_{\ell'}\sum_{k'>k}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\\
&\qquad\cdot\mathcal{L}\{\ell'\,|\,k'\}v^n(h',\ell',k')\Bigg\}. \quad (9)
\end{aligned}$$

Because $v^n(h',\ell',k')$ is nonincreasing in $k'\in\Psi$ for all $h'\in\Omega$ and $\ell'\in\Phi$, by the induction hypothesis, and the fact that $\mathcal{K}$ is CCD, Corollary 1 implies that inequality (9) is preserved if we replace $v^n(h',\ell',k')$ by $v^n(h',\ell',k)$ for all (h',ℓ',k'). Making this substitution and rearranging yields

$$\begin{aligned}
&v^{n+1}(h,\ell,k)-v^{n+1}(h,\ell,k+1)\\
&\quad\geqslant\lambda\sum_{h'}\mathcal{H}\{h'\,|\,h\}\Bigg\{\sum_{\ell'\neq L+1}v^n(h',\ell',k)\sum_{k'}[\mathcal{K}\{k'\,|\,k\}\\
&\qquad-\mathcal{K}\{k'\,|\,k+1\}]\mathcal{L}\{\ell'\,|\,k'\}\\
&\qquad+v^n(h',L+1,k)\sum_{k'}[\mathcal{K}\{k'\,|\,k\}\\
&\qquad-\mathcal{K}\{k'\,|\,k+1\}]\mathcal{L}\{L+1\,|\,k'\}\Bigg\}.
\end{aligned}$$

Employing the identity $\mathcal{L}\{L+1\,|\,k'\}=1-\sum_{\ell'\neq L+1}\mathcal{L}\{\ell'\,|\,k'\}$ yields

$$\begin{aligned}
&v^{n+1}(h,\ell,k)-v^{n+1}(h,\ell,k+1)\\
&\quad\geqslant\lambda\sum_{h'}\mathcal{H}\{h'\,|\,h\}\Bigg\{\sum_{\ell'\neq L+1}v^n(h',\ell',k)\\
&\qquad\cdot\sum_{k'}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\\
&\qquad\cdot\mathcal{L}\{\ell'\,|\,k'\}-v^n(h',L+1,k)\\
&\qquad\cdot\sum_{k'}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\\
&\qquad\cdot\sum_{\ell'\neq L+1}\mathcal{L}\{\ell'\,|\,k'\}+v^n(h',L+1,k)\\
&\qquad\cdot\sum_{k'}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\Bigg\}.
\end{aligned}$$

Eliminating the final term in the right-hand side of the last inequality because it is always zero and rearranging yields

$$\begin{aligned}
&v^{n+1}(h,\ell,k)-v^{n+1}(h,\ell,k+1)\\
&\quad\geqslant\lambda\sum_{h'}\mathcal{H}\{h'\,|\,h\}\Bigg\{\sum_{\ell'\neq L+1}[v^n(h',\ell',k)-v^n(h',L+1,k)]\\
&\qquad\cdot\Bigg[\sum_{k'}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\mathcal{L}\{\ell'\,|\,k'\}\Bigg]\Bigg\}.
\end{aligned}$$

Now $\lambda\geqslant 0$, $\mathcal{H}\{h'\mid h\}\geqslant 0$ for all $h,h'\in\Omega$, and $v^n(h',\ell',k)-v^n(h',L+1,k)\geqslant 0$ for all $h'\in\Omega$, $\ell'\in\Phi$, and $k\in\Psi$, by Proposition 1. Furthermore, because $\mathcal{K}$ is CCD and $\mathcal{L}$ is nonincreasing in $k'\in\Psi$ for fixed $\ell'\neq L+1$, Lemma 1 implies that $\sum_{k'}[\mathcal{K}\{k'\,|\,k\}-\mathcal{K}\{k'\,|\,k+1\}]\mathcal{L}\{\ell'\,|\,k'\}\geqslant 0$ for all $\ell'\neq L+1$. Therefore, $v^{n+1}(h,\ell,k)-v^{n+1}(h,\ell,k+1)\geqslant 0$, which completes the proof. □

The main result on the structure of the optimal policy is given in Theorem 3, which establishes conditions under which there exists a rank-based control-limit optimal policy. Analogous to a liver-based control-limit optimal policy, a *rank-based control-limit optimal policy* is an optimal policy that prescribes, for a given health state h, and a liver quality ℓ, "wait" if the rank of the patient is below some threshold rank k^* and "transplant" for all ranks greater than k^*.

THEOREM 3. *If $\mathcal{K}$ is CCD and $\mathcal{L}\{\ell\,|\,k\}$ is monotonically nonincreasing in $k\in\Psi$ for all $\ell\neq L+1$, then there exists a rank-based control-limit optimal policy.*

PROOF. By contradiction.

To prove the claim, we need to show that, given $h\in\Omega$ and $\ell\in\Phi$, $a(h,\ell,k)=T$ for any $k\in\Psi$ implies that $a(h,\ell,k')=T$ for all $k'\geqslant k$. In other words, given $h\in\Omega$ and $\ell\in\Phi$, the optimal policy is of the following form:

$$\begin{aligned}
&a(h,\ell,1)=\cdots=a(h,\ell,k-1)=W,\\
&a(h,\ell,k)=a(h,\ell,k+1)=\cdots=T,
\end{aligned}$$

for some $k\in\Psi$.

Consider any rank $k\in\Psi$ and assume otherwise, i.e., assume that $a(h,\ell,k)=T$ but $a(h,\ell,k+1)=W$ uniquely. This assumption, respectively, implies that $v(h,\ell,k)=r_T(h,\ell)$ and $v(h,\ell,k+1)>r_T(h,\ell)$. Therefore, we find

$$v(h,\ell,k)-v(h,\ell,k+1)<0,$$

which contradicts the result of Theorem 2. □

Before we proceed with the results associated with the health component (h) of the MDP model, we provide a technical lemma that is used in the proof of Theorem 4. The proof of the lemma and the definition of an IFR matrix is given in the appendix. For notational convenience, we denote the dead state, Δ, as $H+1$ in the remainder of the paper. Let the $(H+1)\times(H+1)$ *augmented health transition probability matrix* $\widehat{\mathcal{H}}$, where the first H states of this matrix represent the health states and the last state represents death, be given by

$$\widehat{\mathcal{H}}=\begin{bmatrix}\mathcal{H} & (I-\mathcal{H})e\\ 0 & 1\end{bmatrix},$$

where I is the $H\times H$ identity matrix and e is an $H\times 1$ vector of ones.

LEMMA 2. *Let $\widehat{\mathcal{H}}$ be an IFR transition probability matrix and $v(h,\ell,k)$ be a nonincreasing function of $h \in \Omega$ for any $k \in \Psi$ and $\ell \in \Phi$. Then, the following hold for all $h \in \Omega$:*

$$\text{(i)} \sum_{h' \leqslant h} \Bigg\{ [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \Bigg[\sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v(h', \ell', k') \Bigg] \Bigg\} \geqslant \Bigg[\sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v(h, \ell', k') \Bigg] \cdot \sum_{h' \leqslant h} [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}], \tag{10}$$

$$\text{(ii)} \sum_{h' > h} \Bigg\{ [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \Bigg[\sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v(h', \ell', k') \Bigg] \Bigg\} \geqslant \Bigg[\sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v(h+1, \ell', k') \Bigg] \cdot \sum_{h' > h} [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}]. \tag{11}$$

Theorem 4 presents the conditions under which the value function is monotone in patient health for fixed rank and liver quality.

THEOREM 4. *If $\widehat{\mathcal{H}}$ is IFR, then $v(h,\ell,k)$ is monotonically nonincreasing in h for any $k \in \Psi$ and $\ell \in \Phi$.*

PROOF. By induction on the steps of the value iteration algorithm.

We prove the theorem for $\ell \neq L+1$ and note that the proof for $\ell = L+1$ follows similarly.

Let $v^i(h,\ell,k)$ be the value associated with state $(h,\ell,k) \in \mathcal{S}$ at the ith iteration of the value iteration algorithm and assume, without loss of generality, that the algorithm starts with a value of zero for each state, i.e., $v^0(h,\ell,k) = 0$ for all $(h,\ell,k) \in \mathcal{S}$. It is clear, by Assumptions AS1 and AS2, that the result holds for iteration 1. Given that $v^n(h,\ell,k) \geqslant v^n(h+1,\ell,k)$, we must show that $v^{n+1}(h,\ell,k) \geqslant v^{n+1}(h+1,\ell,k)$. Because $v(H+1) = 0$, by Equation (2),

$$v^{n+1}(h,\ell,k) = \max \Bigg\{ r_T(h,\ell), r_W(h) + \lambda \sum_{h'} \sum_{k'} \sum_{\ell'} \widehat{\mathcal{H}}\{h' \mid h\} \cdot \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \Bigg\} \tag{12}$$

and

$$v^{n+1}(h+1,\ell,k) = \max \Bigg\{ r_T(h+1,\ell), r_W(h+1) + \lambda \sum_{h'} \sum_{k'} \sum_{\ell'} \widehat{\mathcal{H}}\{h' \mid h+1\} \cdot \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \Bigg\}. \tag{13}$$

If $a^{n+1}(h+1,\ell,k) = T$, then $v^{n+1}(h+1,\ell,k) = r_T \cdot (h+1,\ell) \leqslant r_T(h,\ell) \leqslant v^{n+1}(h,\ell,k)$.

If $a^{n+1}(h+1,\ell,k) = W$, then by Equations (12) and (13),

$$\begin{aligned} v^{n+1}(h,\ell,k) - v^{n+1}(h+1,\ell,k) &\geqslant r_W(h) - r_W(h+1) + \lambda \sum_{h'} \widehat{\mathcal{H}}\{h' \mid h\} \cdot \Bigg\{ \sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \Bigg\} \\ &\quad - \lambda \sum_{h'} \widehat{\mathcal{H}}\{h' \mid h+1\} \Bigg\{ \sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \Bigg\} \\ &\geqslant \lambda \sum_{h'} \Bigg\{ [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \Bigg[\sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h', \ell', k') \Bigg] \Bigg\}. \end{aligned} \tag{14}$$

Because $\widehat{\mathcal{H}}$ is IFR by assumption, $\mathcal{K}\{\cdot\}$ and $\mathcal{L}\{\cdot\}$ are nonnegative by definition, and $v^n(h',\ell',k')$ is nonincreasing in $h' \in \Omega$ for all $k' \in \Psi$ and $\ell' \in \Phi$, Lemma 2 implies that inequality (14) is preserved if we replace $v^n(h',\ell',k')$ by $v^n(h,\ell',k')$ for all (h',ℓ',k'). Therefore,

$$\begin{aligned} v^{n+1}(h,\ell,k) - v^{n+1}(h+1,\ell,k) &\geqslant \lambda \Bigg[\sum_{k'} \sum_{\ell'} \mathcal{K}\{k' \mid k\} \mathcal{L}\{\ell' \mid k'\} v^n(h, \ell', k') \Bigg] \cdot \sum_{h'} [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}]. \end{aligned}$$

The result follows because $\sum_{h'} [\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] = 0$. □

Analogous to a liver-based or a rank-based control-limit optimal policy, a *health-based control-limit optimal policy* is defined as an optimal policy that prescribes, for a given rank state k, and a liver quality ℓ, "wait" in all health states up to (and including) a threshold health state h^* and "transplant" in all health states greater than h^*. Given the result of Theorem 4 and similar conditions to the conditions of Theorem 3 in Alagoz et al. (2004), it can easily be shown that there exists a health-based control-limit optimal policy.

4. Numerical Results

In this section, we present numerical results driven by clinical data. Section 4.1 discusses the parameter estimation process for the MDP model formulated in §2 and presents a numerical example for a patient with Hepatitis B. Section 4.2 discusses the concept of price of privacy in greater detail and presents the results of a numerical study for 200 ESLD patients.

4.1. Parameter Estimation and an Example

In our computational experiments, we define each period to be one day and consider the objective of maximizing the patient's total expected remaining lifetime. Therefore, we set $r_W(h) = 1$ for all $h \in \Omega$ and estimate the patient-specific

total expected posttransplant life days, $r_T(h, \ell)$ for all $h \in \Omega$ and $\ell \in \Phi$, using the posttransplant survival model of Roberts et al. (2004).

Adult ESLD patients are classified by disease severity into Status 1 patients (i.e., patients with a life expectancy of less than seven days without a liver transplant) and MELD (Model for End-Stage Liver Disease) patients. We only consider MELD patients in this study because there are typically fewer than a dozen Status 1 patients nationwide at a given time. In the allocation mechanism, each MELD patient has an integer-valued MELD score based on several lab values between 6 (healthiest) and 40 (sickest). However, due to sparsity of the available data, we represent patient health (h) by MELD scores aggregated in groups of two. Because the natural history of ESLD depends on the type of diagnosis, we estimate different $\mathcal{H}$ matrices for different disease groups using the Natural History Model of Alagoz et al. (2005).

We also follow the liver quality classification scheme of Alagoz et al. (2007a), which considers 14 liver qualities as determined by the age, race, and gender of the donor (Roberts et al. 2004). Detailed descriptions of the liver quality assignment scheme and the estimation of $r_T(h, \ell)$ are provided in Alagoz et al. (2007a).

To estimate $\mathcal{K}$ and $\mathcal{L}$, we use the national liver allocation model of Shechter et al. (2005), which simulates the evolution of the waiting list under various liver allocation policies for the United States based on clinical data. We resort to this simulation model to collect rank information, which is not available in any form from any of the clinical resources. Tracking the rank of a patient in the nationwide waiting list would result in enormously large $\mathcal{K}$ and $\mathcal{L}$ matrices given that the waiting list contains nearly 20,000 patients. Moreover, the vast majority of offers are made to patients in the same geographic area as the donated liver because the allocation mechanism exhausts the local geographic area before considering other areas. Therefore, we simulate the national waiting list, but track the rank of the patients within the geographic area where they are registered. Because OPOs represent different populations, for each OPO we estimate different $\mathcal{K}$ and $\mathcal{L}$ matrices of varying sizes depending on the size of the geographic area served by the OPO. For notational convenience, we drop the dependency on OPOs in the following discussion.

We use 30 independent replications of the simulation to estimate $\mathcal{K}$ and $\mathcal{L}$. In their original paper, Shechter et al. (2005) also used 30 replications for their simulation model and found that the simulation output closely matches UNOS data for several important statistics such as number of new patients listed, number of cadaveric donors, number of transplants, median waiting time for a transplant, and one-year survival rates for patients and organs after receiving the transplant. To estimate $\mathcal{L}$, for each OPO, we count the number of offers each rank receives during the simulation, provided a liver is donated, and transform these counts into probabilities. Let $f(k)$ represent the probability that a rank k patient in an arbitrary OPO receives an offer given that a liver is offered in her OPO. When a liver is offered in an OPO, the patient having the highest priority in this OPO receives this offer, and depending on her decision, the liver may be offered to the patient with the second-highest priority and so on. Therefore, we assume that $f(1) = 1$ and $f(k)$ is monotonically nonincreasing in k. The second step in estimating each $\mathcal{L}$ matrix involves augmenting $f(k)$ to distinguish between liver qualities. To do so, we fit an exponential function of the form $\exp(-\theta k)$ to $f(k)$ using ordinary least squares. Then, to ensure that higher-quality livers are accepted earlier, we perturb θ to obtain $\tilde{f}(k, \ell)$, the probability that a rank k patient receives a donated liver of quality ℓ given that this liver is offered in her OPO. In doing so, we assume that the probability that a rank k patient receives an average-quality liver is at most twice that of the highest-quality liver (i.e., $\ell = 1$) and is at most $2/3$ of the probability of receiving the lowest-quality liver (i.e., $\ell = 14$). Specifically, we obtain the perturbed θ values, θ_ℓ, by

$$\theta_\ell = \begin{cases} \theta(1 + 1/\ell) & \text{if } \ell = 1, \ldots, 7, \\ \theta(1 - 0.05(\ell - 7)) & \text{if } \ell = 8, \ldots, 14. \end{cases}$$

For a typical OPO, Figure 2 shows the original $f(k)$ as estimated from the simulation and the best exponential fit to $f(k)$ as discussed above. Figure 3 shows the resulting $\tilde{f}(k, \ell)$ functions for $\ell = 1, \ldots, 14$. Next, we assume that liver arrivals in each OPO follow a Poisson process such that the probability that there is a liver offer in this OPO during an arbitrary day is given by α, where α is estimated using data obtained from UNOS (UNOS 2006a). Lastly, let $\gamma(\cdot)$ be the organ quality distribution of an arriving liver, which we obtain from the simulation by counting the number of liver offers of each quality. Then,

$$\mathcal{L}\{\ell \mid k\} = \alpha \cdot \gamma(\ell) \cdot \tilde{f}(k, \ell) \quad \text{for all } \ell \neq L + 1,\ k \in \Psi,$$

and

$$\mathcal{L}\{L + 1 \mid k\} = 1 - \sum_{\ell \in \Phi} \mathcal{L}\{\ell \mid k\} \quad \text{for all } k \in \Psi.$$

Figure 2. Liver offer probabilities as a function of rank given there is an offer for OPO "A."

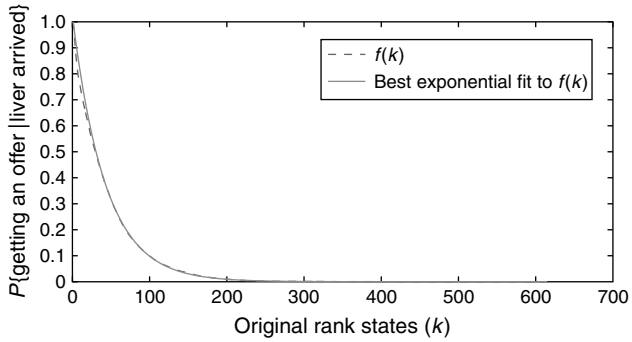

Note. $r^2 = 0.9991$ for the exponential fit to $f(k)$.

Figure 3. Liver offer probabilities as a function of rank given there is an offer of some quality for OPO "A."

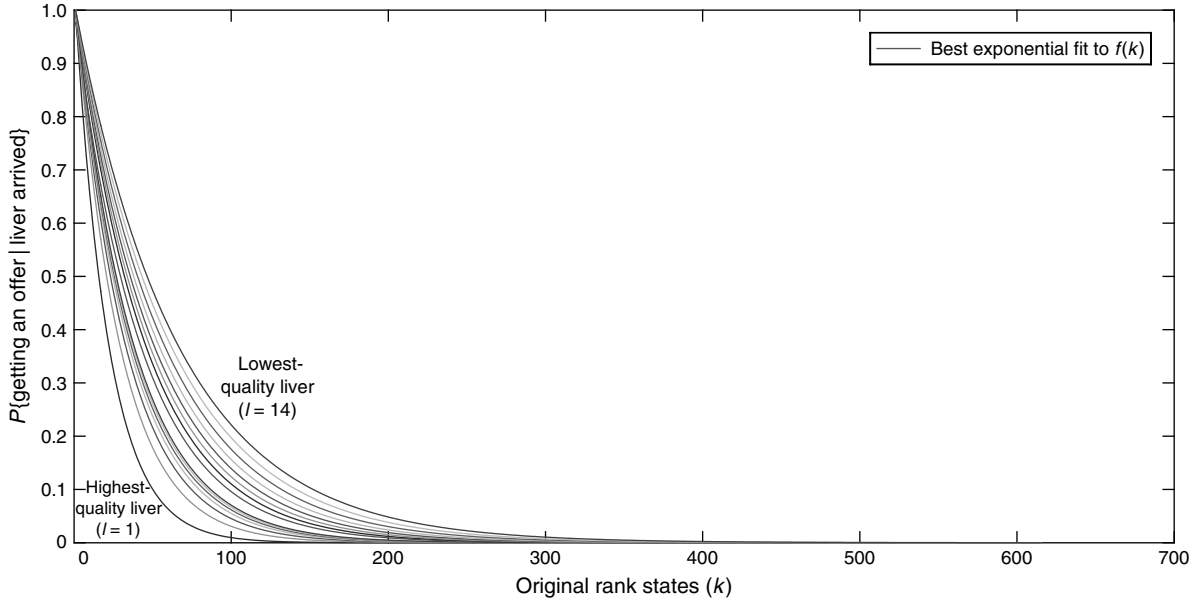

Notes. Given $f(1) = 1$ and $f(k)$ is nonincreasing in k, we construct $\tilde{f}(1, \ell) = 1$ and $\tilde{f}(k, \ell)$ is monotonically nonincreasing in k for all $\ell = 1, \ldots, 14$. Furthermore, $\partial\tilde{f}(k, \ell)/\partial k < \partial\tilde{f}(k, \ell + 1)/\partial k$ for $\ell = 1, \ldots, 13$, indicating that higher-quality livers are accepted sooner.

Even within an OPO, the number of rank states can be as large as several thousand for OPOs serving large populations, which may yield computationally intractable problems. For this reason, for each OPO, we aggregate the rank states into 30 new rank states, which is found to be computationally tractable for the numerical study presented in §4.2, using the original $f(k)$ estimates in the following manner. Given a number of original rank states, we start from rank 1 and consolidate the first j ranks, where j is the maximum number of ranks such that the difference between the average function value over these j states and the function value at the $(j+1)$st state (e.g., $(1/j)\sum_{i=1}^{j} f(i) - f(j+1)$) is larger than some predetermined threshold. The threshold is found by line search so as to guarantee 30 final rank states. We then repeat this process starting from the $(j+1)$st rank and so on. Figure 4 depicts the approximation generated by this aggregation scheme for a typical OPO. Table 2 displays detailed information about the aggregated rank states for this OPO.

To estimate the corresponding 30×30 $\mathcal{K}$ matrix for each OPO, we count the number of transitions during the simulation from each rank to every rank. We then transform

Figure 4. Effect of state aggregation on $f(k)$ for OPO "A."

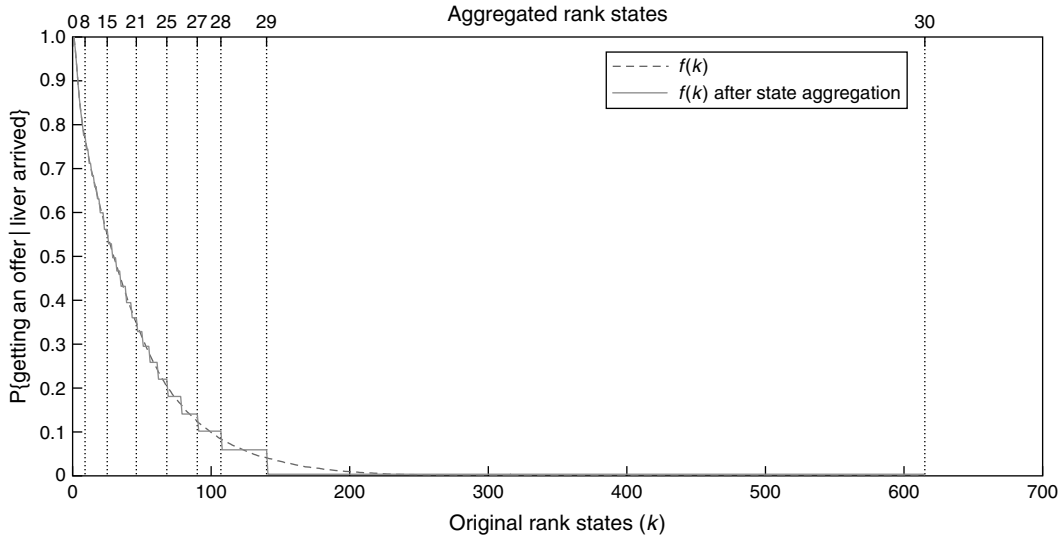

Notes. This OPO originally had 615 rank states. The first seven original rank states remain unaggregated, the next two form rank 8, and so on.

Table 2. Aggregated rank information for OPO "A."

Aggregated rank	Original rank	Aggregated rank	Original rank	Aggregated rank	Original rank
1	1	11	14–15	21	43–46
2	2	12	16–17	22	47–50
3	3	13	18–19	23	51–55
4	4	14	20–22	24	56–61
5	5	15	23–25	25	62–68
6	6	16	26–28	26	69–78
7	7	17	29–31	27	79–90
8	8–9	18	32–34	28	91–107
9	10–11	19	35–38	29	108–140
10	12–13	20	39–42	30	141–615

these counts into transition probabilities by dividing each count by its row sum. Finally, we average the resulting transition probabilities across replications. Across all estimates of $\mathcal{K}$, the maximum standard error of the point estimates varies between 0.1245% and 0.3682% with an average of 0.1844% and a standard deviation of 0.0448%. Similarly, the maximum standard error of the point estimates in $\mathcal{L}$ varies between 0.0684% and 0.3616% with an average of 0.1594% and a standard deviation of 0.0709%.

Finally, we assume an annual discount rate of 0.97, which translates into a daily discount rate (λ) of 0.999917.

Given these parameter estimates, Figure 5 depicts the optimal policy, which exhibits control-limit structure in all components, for a 50-year-old female patient in OPO "B" with Hepatitis B. In Figure 5, liver quality 15 represents the "no liver offer" case. As shown in Figure 5, the optimal action varies across liver quality and health as measured by MELD score. If the patient is at the top of the list and has a MELD score of 20, then the optimal policy prescribes "transplant" for a liver of quality 1 and "wait" otherwise. This is an example of a liver-based control-limit optimal policy. Similarly, if the patient is at the top of the list and receives a liver offer of quality 2, then the optimal policy prescribes the "wait" action if her MELD score is below 25 and the "transplant" action otherwise. This is an example of a health-based control-limit optimal policy. Figure 5 further shows that the optimal action varies significantly by rank as well. The "Transplant" region is smallest when the patient is at the top of the list and gradually grows as her rank deteriorates. In other words, the patient is more selective if she is at the top of the list and becomes less selective if her rank decreases. For example, when the patient's MELD score is 20 and she is at the top of the list, she rejects all liver offers of quality $\ell \geqslant 2$ and accepts only the highest-quality liver. However, at the same MELD score, if she is at the bottom of the list, she rejects only liver offers of quality $\ell \geqslant 12$ and accepts all liver offers of quality $\ell < 12$, which is an example of a rank-based control-limit optimal policy.

In all of our 200 test problems that are randomly generated from the simulation model of Shechter et al. (2005), (AS1) and (AS2) are satisfied by the reward estimates $r_W(h)$ and $r_T(h, l)$, and the conditions on $\mathcal{L}$ in Theorems 2 and 3 are satisfied by the parameter estimates. However, the CCD condition on the $\mathcal{K}$ matrix is not always satisfied,

Figure 5. Control-limit optimal policy in all parameters for a 50-year-old patient with Hepatitis B.

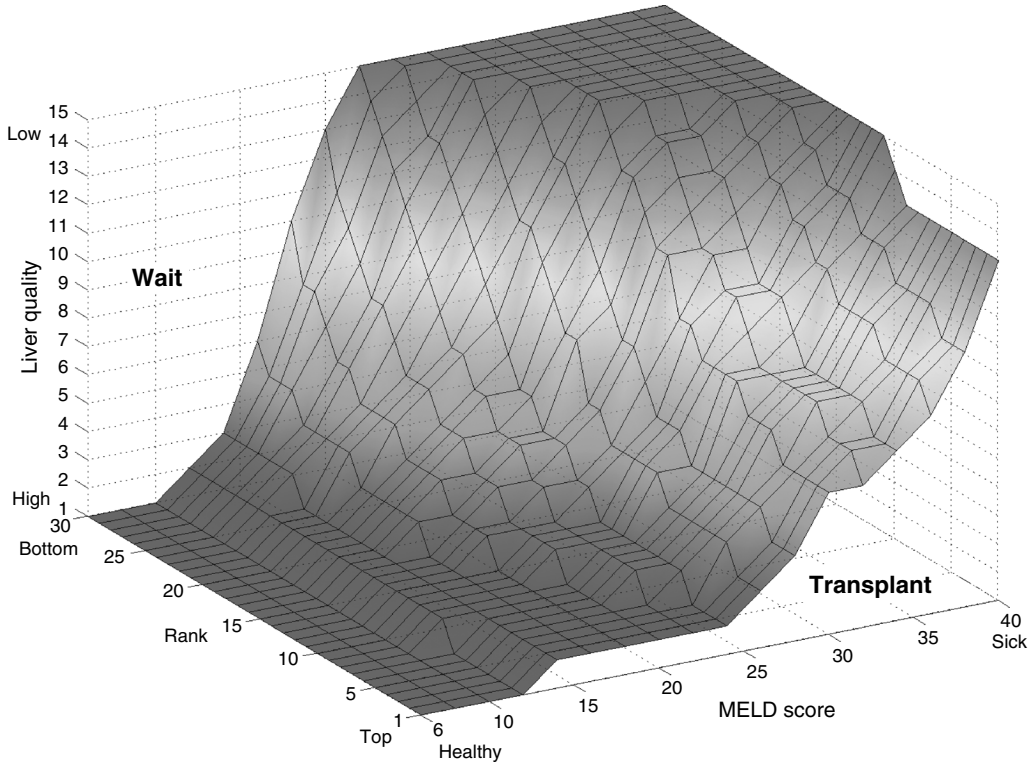

although violations are not large. To quantify the magnitude of the violation of the CCD condition, we define the following metric:

$$\epsilon = \sum_{k'} \epsilon_{k'},$$

where

$$\epsilon_{k'} = \begin{cases} \sum_{k'} \max\{0, \mathcal{K}\{k' \mid k\} - \mathcal{K}\{k' \mid k+1\}\} \\ \quad \text{for } k' = 1, \ldots, k-1, \\ \sum_{k'} \max\{0, \mathcal{K}\{k' \mid k+1\} - \mathcal{K}\{k' \mid k\}\} \\ \quad \text{for } k' = k, k+1, \ldots. \end{cases}$$

Note that ϵ depends on the geographic area. Across all the estimates of $\mathcal{K}$ for each of the OPOs, ϵ varies between 0.001838 and 0.005425 with an average of 0.003332 and a standard deviation of 0.000806. Given the monotonicity of the value function in k and the rank-based control-limit optimal policy in all of the 200 test instances, we conclude that the results of Theorems 2 and 3 are fairly robust to small violations of the CCD requirement for the $\mathcal{K}$ matrix. We refer the reader to Alagoz et al. (2004) for a discussion of violations on the $\widehat{\mathcal{H}}$ matrix.

4.2. Estimating the Price of Privacy

The societal price of privacy is the aggregate benefit that society would accrue if the waiting list were made transparent. An estimate for the true price of privacy would be obtained by comparing a system (in which every patient has partial rank information as in the current allocation system and behaves optimally with this information) to a benchmark system (in which every patient has full rank information and behaves optimally). Our current model is unable to provide an exact value for the societal price of privacy because if the waiting list were to become transparent, the organ offer probabilities would change substantially as the allocation system moved to a new equilibrium, thus making precise parameter estimation using existing data impossible. Rather, due to difficulties in identifying an equilibrium in either of these systems, we focus on a special case where only one patient, who is provided the waiting list information, is considered. As a result, the quantities we provide can be viewed as estimates for the true values. We define the *patient's price of privacy* (PPoP) as the amount of life days gained when she acts optimally based on full knowledge of the waiting list, as opposed to her optimal actions under the current allocation rules.

We provide an estimate of the PPoP by comparing the model of §2, denoted the explicit waiting list model (EWLM), to the implicit waiting list model (IWLM) of Alagoz et al. (2007a). More specifically, let π_a be an optimal IWLM policy and π_E be an optimal EWLM policy for the same patient. For any given $h \in \Omega$ and $\ell \in \Phi$, define

$$\pi_I(h, \ell, k) = \pi_a(h, \ell) \quad \text{for all } k \in \Psi.$$

This policy π_I may be viewed as the projection of the optimal IWLM policy onto the EWLM state space. Intuitively, if the patient does not have any rank information and solves IWLM, then her optimal actions as prescribed by this model should be same for each (h, ℓ) pair regardless of her rank, k. Let the benefit of using policy π_E over policy π_I in state $(h, \ell, k) \in \mathcal{S}'$ be given by

$$b(h, \ell, k) = v^{\pi_E}(h, \ell, k) - v^{\pi_I}(h, \ell, k),$$

where $v^{\pi_E}(h, \ell, k)$ and $v^{\pi_I}(h, \ell, k)$, respectively, is the maximum total expected discounted reward for state (h, ℓ, k) associated with policy π_E and π_I.

Proposition 3 establishes that the benefit of using policy π_E over policy π_I is nonnegative in every state, which implies that an optimal policy recommended by IWLM may not provide the true optimal policy for EWLM. The proof is obvious and is therefore omitted.

PROPOSITION 3. *$b(h, \ell, k) \geqslant 0$ for all $h \in \Omega$, $\ell \in \Phi$, and $k \in \Psi$.*

We provide an estimate of a PPoP ratio (i.e., the ratio of the patient's price of privacy to her optimal reward under the current allocation rules) using the following formula:

$$\rho = \frac{b(\tilde{h}, L+1, K)}{v^{\pi_I}(\tilde{h}, L+1, K)}, \tag{15}$$

where $\tilde{h}$ is the patient's health at the time of her registration to the waiting list and $K = \sup_{k \in \Psi}\{k\}$. This metric measures the improvement associated with using the optimal EWLM policy over the optimal IWLM policy as a fraction of the optimal value of being in state $(\tilde{h}, L+1, K)$. We choose state $(\tilde{h}, L+1, K)$ because the rank of a new patient is usually very low and a patient rarely receives a liver offer on the day she joins the list. The quantity given by Equation (15) provides an estimate of the true PPoP ratio partly because of the fact that IWLM does not model the partial information availability in the current liver allocation system.

Consider, for example, the patient whose optimal policy is depicted in Figure 5. The estimate of her PPoP ratio, as computed using Equation (15), is 4.51%, which corresponds to 103.01 additional expected life days. Figure 6 adds the projected optimal IWLM policy to Figure 5 for the same patient. First note that, by construction, the optimal action of this projected policy does not vary across rank states. Furthermore, although the trends are same, the control limits h^* and ℓ^* of the projected policy do not always coincide with those of the optimal EWLM policy. In other words, ℓ^* is nondecreasing in h for fixed k in both policies; however, they do not always coincide. Similar observations apply for h^* as well.

We compute the estimate ρ for the true PPoP ratio for each of the 200 patients generated by the national liver allocation model of Shechter et al. (2005). Figure 7 presents a histogram of the ρ values for all 200 patients. The ρ values range between 0.31% and 15.57%, with a median value of

Figure 6. Comparison of EWLM and IWLM optimal policies.

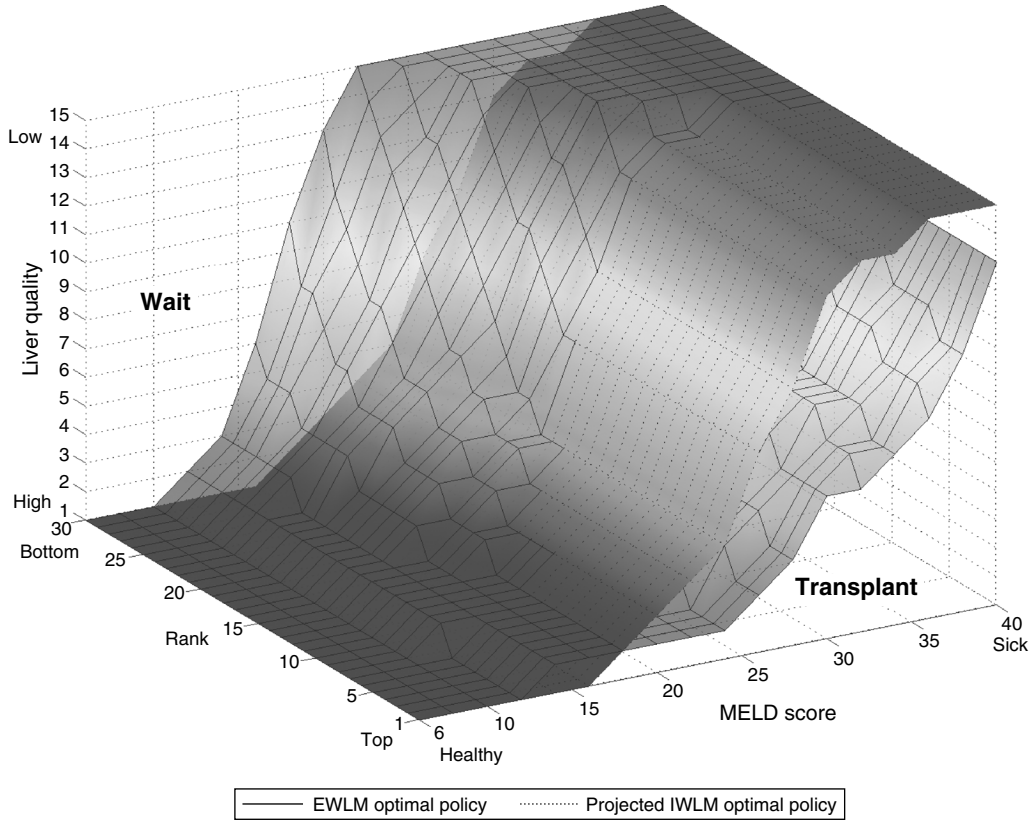

4.59%, an average value of 5.22%, and a standard deviation of 3.82%.

Table 3 presents the descriptive statistics associated with the estimated PPoP ratios for the sampled patients in different disease groups; Table 4 presents these statistics by age range; and Table 5 presents these statistics by geographic area. We observe that the mean ρ values for patients in disease groups 1 and 4 and for those in disease groups 2 and 3 are close to each other. However, the mean ρ for patients in disease groups 1 and 4 is more than 25% higher compared to that for patients in disease groups 2 and 3. Patients in disease groups 1 and 4 have the best post-transplant survival (e.g., 10-year survival is approximately 75%), whereas patients in disease groups 2 and 3 have much poorer posttransplant survival (e.g., 10-year survival is approximately 60%). Therefore, it appears that the PPoP ratio declines as the benefit of transplantation declines.

Figure 7. Histogram of the estimate for PPoP ratio for 200 patients generated from simulation.

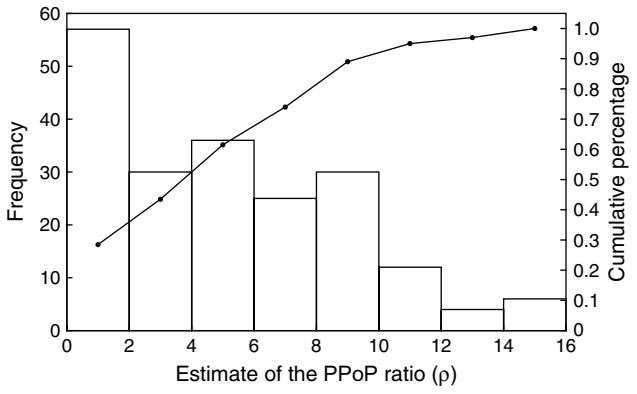

The mean ρ for patients younger than 20 and older than 70 is about 40% less than that for patients in the remaining age groups. Although we do not observe any major difference in the mean ρ values for patients in the remaining age groups, the results indicate a general decrease in the PPoP ratio as age increases. Because elderly patients have shorter posttransplant survival, this observation further supports the hypothesis that a PPoP ratio decreases with decreased posttransplant survival.

As seen in Table 5, Regions 6–11 serve significantly larger populations than those served by regions 1–5. The mean ρ for patients registered in OPOs serving larger populations is observed to be about 3.5 times the mean ρ for patients registered in OPOs serving smaller populations. This result is intuitive because as the size of population served in a geographic area increases, the liver offer probability differs significantly across ranks.

5. Summary and Future Research

We develop an MDP model to optimize the accept/reject decision faced by ESLD patients. This model explicitly

Table 3. Descriptive statistics for estimated PPoP ratio (ρ) by disease group.

Disease group	Number of patients	Min	Max	Median	Mean	Standard deviation
1	81	0.31%	15.01%	4.82%	5.61%	4.34%
2	71	0.47	11.53	4.35	4.79	2.96
3	22	0.75	10.36	2.88	3.85	2.82
4	26	0.60	15.57	6.34	6.36	4.55
ALL	200	0.31	15.57	4.59	5.22	3.82

Notes. Disease group 1 includes primary biliary cirrhosis, primary sclerosing cholangitis, alcoholic liver disease, and autoimmune disorders; disease group 2 includes Hepatitis B and C viruses; disease group 3 patients have acute liver failure; and disease group 4 patients have metabolic disorders (e.g., glycogen storage disease types I and II, and Gaucher's disease).

considers waiting list effects by augmenting the state space of the implicit waiting list model studied by Alagoz et al. (2007a). We derive conditions under which the optimal value function is monotone in each dimension of the state space, namely, health, liver quality, and rank, and conditions under which control-limit optimal policies exist for each dimension. In establishing these results, we define a new class of stochastic matrices, termed CCD matrices, and explore its relationship to well-known classes of matrices (i.e., IFR and TP_2 matrices). Computational experiments parameterized by clinical data reveal that complete knowledge of the composition of the waiting list significantly affects the optimal policy. In particular, a patient is much more selective if she knows that she is near the top of the waiting list and becomes gradually less selective as her position deteriorates.

We solve our explicit waiting list model for 200 randomly generated patients. Although the conditions of Theorems 2 (monotonicity in rank) and 4 (monotonicity in health) are not always satisfied, in all 200 cases the value function is monotone in each component of the model (i.e., h, ℓ, and k). This result suggests that the monotonicity of the value function is robust to small violations of the CCD condition on the rank transition probability matrix, as well as the IFR condition on the augmented health transition probability matrix. In all of our computational experiments, the conditions of Theorem 1 hold, and therefore we find the optimal policy to be of liver-based control-limit type. Although the CCD condition of Theorem 3 is not always satisfied, the magnitudes of the violations are not significant, and, as a result, the optimal policy always has a rank-based control limit in our experiments. However, for some patients, the optimal policy does not have a health-based control limit. As noted by Alagoz et al. (2007a), as the patient deteriorates, the rate of increase in the probability of receiving higher-quality liver offers may be sufficiently high so that she rejects low-quality livers that she would have accepted in better health in anticipation of higher-quality liver offers.

We use our explicit waiting list model to estimate a patient's price of privacy, which is incurred due to suboptimal decision making from a lack of complete waiting list information. By comparing the results of our model to those of the implicit waiting list model of Alagoz et al. (2007a), we provide a quantitative estimate for a patient's price of privacy. Our computational experiments reveal that this quantity varies significantly by the particular etiology of ESLD. Indeed, although the majority of the patients would realize a less than 6% increase in total expected remaining lifetime by having complete information about the waiting list, there are patients in our study who realize improvements as high as 15%. In particular, patients diagnosed with diseases that yield shorter posttransplant survivals tend to have a smaller PPoP ratio. Similarly, the PPoP ratio tends to be smaller for older patients, who typically have shorter posttransplant survivals. Furthermore, our results indicate that patients registered in OPOs serving larger populations experience a higher PPoP ratio.

The numerical estimates presented in this paper are immediately useful to the policy makers if the value of maintaining a concealed waiting list exceeds these quantities. Furthermore, we are unaware of other research that attempts to quantify a cost of privacy in health care; we are aiding policy makers by demonstrating that such a cost may indeed exist. However, our findings are several steps away from providing precise estimates of the price of privacy. Furthermore, even if our estimates were exact, the benefits may not be shared equally among all patients, and it is likely that some patients will be worse off under a transparent waiting list. Our results need further support before any changes are implemented. Due to the difficulties inherent in describing the new equilibrium in response to a

Table 4. Descriptive statistics for estimated PPoP ratio (ρ) by age group.

Age group	Number of patients	Min	Max	Median	Mean	Standard deviation
Age $<$ 20	5	1.27%	5.38%	2.05%	2.66%	1.73%
20 $\leqslant$ Age $<$ 30	8	1.09	11.78	4.46	5.27	4.06
30 $\leqslant$ Age $<$ 40	34	0.31	15.57	4.95	5.78	4.09
40 $\leqslant$ Age $<$ 50	54	0.73	14.56	6.02	5.96	3.32
50 $\leqslant$ Age $<$ 60	56	0.35	14.50	3.92	4.76	4.09
60 $\leqslant$ Age $<$ 70	39	0.40	15.01	4.18	4.84	4.01
Age $\geqslant$ 70	4	1.43	6.04	3.83	3.78	2.08
ALL	200	0.31	15.57	4.59	5.22	3.82

Table 5. Descriptive statistics for estimated PPoP ratio (ρ) by geographic area.

Region ID	Number of patients	Min	Max	Median	Mean	Standard deviation
1	2	0.47%	0.60%	0.53%	0.53%	0.10%
2	6	0.64	2.91	0.88	1.18	0.86
3	8	0.40	2.34	1.33	1.28	0.70
4	7	0.44	3.70	1.47	1.84	1.18
5	6	0.31	5.40	2.25	2.55	2.15
6	20	0.72	7.68	3.09	3.11	1.96
7	25	1.41	9.18	4.18	4.66	2.26
8	28	0.55	15.01	5.77	5.18	3.76
9	34	0.54	15.57	5.01	6.33	5.30
10	15	1.76	10.66	5.51	6.71	2.89
11	49	0.59	12.14	8.35	7.32	3.12
ALL	200	0.31	15.57	4.59	5.22	3.82

Note. We group the OPOs by region; however, the regions are renumbered due to a data use agreement with UNOS.

change in allocation policy, we suggest that policy makers initiate limited pilot studies over small populations (e.g., in a single OPO) to gauge the accuracy of our estimates if they believe that increasing the transparency of the waiting list may be beneficial.

One potential limitation of the current numerical study is that the parameter estimates used in this study are point estimates. As more data become available, future work should include extensive sensitivity analysis on these point estimates. Also, this limitation can potentially be overcome by incorporating robust dynamic programming techniques (Iyengar 2005, Nilim and Ghaoui 2005).

We emphasize that the estimate we provide for a patient's price of privacy can be further refined by incorporating the partially observable nature of the waiting list under the current liver allocation system. Modeling this partial information availability, and hence obtaining a better estimate of the price of privacy, is left for future research. Moreover, if a single transplant center dominates an OPO, then the price of privacy for the patients listed in such an OPO may be significantly smaller than what is suggested by our numerical study. This result is expected because physicians in such OPOs may currently be able to discern a patient's rank fairly precisely. Our model also provides a more realistic representation of the decision-making process in these situations.

Furthermore, we assume that the current equilibrium of the liver allocation system is not affected by the completely observable waiting list assumption. In other words, we assume that all other patients continue to make decisions according to their current policies, and therefore the rank transition probability matrix and the liver offer probability matrix can be estimated using existing data. If completely publishing the waiting list causes the current equilibrium of the liver allocation system to shift, then estimating the price of privacy becomes dramatically more difficult, and would require modeling techniques such as competitive Markov decision processes (Filar and Vrieze 1996, Neyman and Sorin 2003). A complete equilibrium analysis would require analyzing an asymmetric stochastic game with thousands of players each with nonzero-sum rewards. Analysis of such a large-scale game-theoretic model and hence answering the question of how the match (e.g., median waiting time before a transplant, average quality-adjusted life years gained) between patients and organs changes under a new equilibrium that emerges when every ESLD patient has access to the liver transplant waiting list is left for future research.

Another interesting research direction is to study how to reduce the price of privacy. Individual patients might choose to reveal their relevant information to other such patients by joining an information-sharing consortium. This practice would give members of the consortium a more complete view of the waiting list. Several interesting questions arise from such a data-sharing agreement: How large must such a consortium be before benefits accrue to its members? How would the composition of a consortium affect the benefits/costs to the members? How would patients who refuse to join be affected? Finally, if UNOS does not manage such a consortium, can data-sharing agreements be made so that each patient has an incentive to represent her information truthfully? We leave these and other such questions for future research.

Appendix. Properties of CCD Matrices

This appendix gives the proofs of the technical materials presented in §3 and collects together several properties of CCD matrices whose definition was introduced in §3. We start with the proofs.

Proof of Proposition 2. (i) If $\mathcal{P}$ is a CCD matrix, then the term in the square brackets on the left-hand side of inequality (5) is nonnegative by definition (see Equation (3)). Thus, the coefficient of $g(j)$ is nonnegative for all j on the left-hand side of (5) because we are given nonnegative $f(j)$. Furthermore, because $g(j) \geqslant g(k)$ for any $j \leqslant k$, we can write

$$\sum_{j\leqslant k}[\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}]f(j)g(j) \geqslant g(k)\sum_{j\leqslant k}[\mathcal{P}\{j \mid k\} - \mathcal{P}\{j \mid k+1\}]f(j),$$

which establishes the result.

(ii) Similar to the proof in case (i). □

Proof of Lemma 2. (i) Starting with the left-hand side of inequality (10), we can write

$$\sum_{h'\leqslant h}\left\{[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'}\mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h', \ell', k')\right]\right\}$$

$$= [\widehat{\mathcal{H}}\{1 \mid h\} - \widehat{\mathcal{H}}\{1 \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(1,\ell',k')\right] + \sum_{h'=2}^{h}\left\{[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h',\ell',k')\right]\right\}$$

$$\geqslant \sum_{h'=1}^{2}[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(2,\ell',k')\right] + \sum_{h'=3}^{h}\left\{[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h',\ell',k')\right]\right\},$$

where the last inequality follows from the fact that $\widehat{\mathcal{H}}$ IFR implies $\widehat{\mathcal{H}}\{1 \mid h\} - \widehat{\mathcal{H}}\{1 \mid h+1\} \geqslant 0$, $\mathcal{K}\{\cdot\}$ and $\mathcal{L}\{\cdot\}$ are nonnegative, and $v(1,\ell',k') \geqslant v(2,\ell',k')$. Repeating this argument (i.e., $\widehat{\mathcal{H}}$ IFR implies $\sum_{h'=1}^{2}[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\} \geqslant 0$, $\mathcal{K}\{\cdot\}$ and $\mathcal{L}\{\cdot\}$ are nonnegative by definition, and $v(2,\ell',k') \geqslant v(3,\ell',k')$) yields

$$\sum_{h' \leqslant h}\left\{[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h',\ell',k')\right]\right\} \geqslant \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(3,\ell',k')\right] \cdot \sum_{h'=1}^{3}[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] + \sum_{h'=4}^{h}\left\{[\widehat{\mathcal{H}}\{h' \mid h\} - \widehat{\mathcal{H}}\{h' \mid h+1\}] \cdot \left[\sum_{k'}\sum_{\ell'} \mathcal{K}\{k' \mid k\}\mathcal{L}\{\ell' \mid k'\}v(h',\ell',k')\right]\right\}.$$

Continuing in this manner until $v(h-1,\ell',k')$ is replaced by $v(h,\ell',k')$ establishes the result.

(ii) Similar to the proof in case (i). □

Next, we show the relationship between CCD matrices and two other important classes of stochastic matrices in the literature, namely, IFR matrices and TP_2 matrices. We start by recalling the definitions of IFR matrices and TP_2 matrices.

Definition 2. An $n \times n$ stochastic matrix P is called IFR (Increasing Failure Rate) if, for $i = 1, \ldots, n-1$,

$$\sum_{j=k}^{n} P\{j \mid i\} \leqslant \sum_{j=k}^{n} P\{j \mid i+1\} \quad \text{for } k = 1, \ldots, n.$$

Definition 3. An $n \times n$ stochastic matrix P is called TP_2 (Totally Positive of order 2) if the determinants of all 2×2 submatrices of P are nonnegative.

IFR and totally positive matrices have been extensively studied. We know that a TP_2 matrix is also IFR (Karlin 1968) and are interested in finding any relationship between the CCD class and these two classes. Proposition 4 states that every CCD matrix is also an IFR matrix.

Proposition 4. *If P is CCD, then it is also IFR.*

Proof. Let P be an $n \times n$ stochastic CCD matrix and pick an arbitrary index $i \in \{1, \ldots, n\}$.

Initially consider the first set of inequalities, (3), for the given i that must hold for a CCD matrix. Summing these inequalities for $j = 1$ to i_2, where $i_2 = 1, \ldots, i$, we obtain

$$\sum_{j=1}^{i_2} P\{j \mid i\} \geqslant \sum_{j=1}^{i_2} P\{j \mid i+1\} \quad \text{for } i_2 = 1, \ldots, i,$$

$$\Rightarrow 1 - \sum_{j=i_2+1}^{n} P\{j \mid i\} \geqslant 1 - \sum_{j=i_2+1}^{n} P\{j \mid i+1\} \quad \text{for } i_2 = 1, \ldots, i,$$

$$\Rightarrow \sum_{j=i_2+1}^{n} P\{j \mid i\} \leqslant \sum_{j=i_2+1}^{n} P\{j \mid i+1\} \quad \text{for } i_2 = 1, \ldots, i,$$

$$\Rightarrow \sum_{j=k}^{n} P\{j \mid i\} \leqslant \sum_{j=k}^{n} P\{j \mid i+1\} \quad \text{for } k = 2, \ldots, i+1. \quad (16)$$

Now consider the second set of inequalities, (4), for the given i that must hold for a CCD matrix. Summing these inequalities for i_1 through $j = n$, where $i_1 = i+2, \ldots, n$, we obtain

$$\sum_{j=i_1}^{n} P\{j \mid i\} \leqslant \sum_{j=i_1}^{n} P\{j \mid i+1\} \quad \text{for } i_1 = i+2, \ldots, n. \quad (17)$$

Finally, combining results (16) and (17), we obtain

$$\sum_{j=k}^{n} P\{j \mid i\} \leqslant \sum_{j=k}^{n} P\{j \mid i+1\} \quad \text{for } k = 2, \ldots, n. \quad (18)$$

In addition, because both sides of the inequality in (18) are equal to one when $k = 1$, we conclude that the inequality holds for $k = 1, \ldots, n$, which is the definition of an IFR matrix. □

Proposition 5 states that not every IFR matrix is CCD.

Proposition 5. *If P is IFR, then it is not necessarily CCD.*

Proof. Consider the following stochastic matrix:

$$P = \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0.7 & 0.1 & 0.2 \\ 0.5 & 0.2 & 0.3 \end{bmatrix}.$$

P is IFR, but it is not CCD because the maximum element of the second column does not appear on the diagonal entry. □

Having proved that CCD is a stronger condition than IFR, we now investigate how CCD is related to TP_2. Proposition 6 shows that *CCD* and TP_2 conditions are equivalent for 2×2 matrices. However, the answer remains ambiguous for matrices of larger size, as indicated in Remark 1.

PROPOSITION 6. *Consider an $n \times n$ stochastic matrix P. For $n = 2$, P is CCD $\Leftrightarrow$ P is TP_2 $\Leftrightarrow$ P is IFR.*

PROOF. (i) Let P be given as

$$P = \begin{bmatrix} a & 1-a \\ b & 1-b \end{bmatrix},$$

where $0 \leqslant a, b \leqslant 1$. First, assume that P is CCD. Then, $a \geqslant b$ must hold. Therefore, $|P| = a(1-b) - b(1-a) = a - b \geqslant 0$, which implies P is TP_2. Next, assume that P is TP_2. Then, $|P| = a(1-b) - b(1-a) = a - b \geqslant 0$ must hold. Therefore, $a \geqslant b$, which implies P is *CCD*. This completes the proof for the equivalence of the CCD and TP_2 conditions.

Finally, assume that P is IFR. Then, $1 - a \leqslant 1 - b$ must hold. Therefore, $a \geqslant b$, which implies P is CCD. Combining this result with that of Proposition 4 completes the proof for the equivalence of the CCD and IFR conditions. □

REMARK 1. For $n > 2$, if P is CCD (TP_2), then it is not necessarily TP_2 (CCD) as indicated by the following examples.

First, it is easily verified that the following matrix is CCD:

$$P = \begin{bmatrix} 0.9 & 0.1 & 0.0 \\ 0.6 & 0.2 & 0.1 \\ 0.4 & 0.1 & 0.5 \end{bmatrix};$$

however, it is not TP_2 because the lower left 2×2 submatrix has a negative determinant.

Second, the following TP_2 matrix is not CCD because the maximum element in the second column is not on the diagonal:

$$P = \begin{bmatrix} 0.4 & 0.4 & 0.2 \\ 0.3 & 0.3 & 0.4 \\ 0.2 & 0.3 & 0.5 \end{bmatrix}.$$

Acknowledgments

This research was supported by National Science Foundation grants CMMI-0546960 and CMMI-0700094 and by the National Library of Medicine grant R21-LM008273. The authors thank Edwin Romejin and three anonymous referees for providing constructive comments that helped improve the paper.

The data and analyses reported in the 2005 Annual Report of the U.S. Organ Procurement and Transplantation Network and the Scientific Registry of Transplant Recipients have been supplied by UNOS and URREA under contract with HHS. The authors alone are responsible for reporting and interpreting these data; the views expressed herein are those of the authors and not necessarily those of the United States Government.

References

Ahn, J.-H., J. C. Hornberger. 1996. Involving patients in the cadaveric kidney transplant allocation process: A decision-theoretic perspective. *Management Sci.* **42**(5) 629–641.

Alagoz, O. 2004. Optimal policies for the acceptance of living- and cadaveric-donor livers. Ph.D. dissertation, University of Pittsburgh, Pittsburgh.

Alagoz, O., L. M. Maillart, A. J. Schaefer, M. S. Roberts. 2004. The optimal timing of living-donor liver transplantation. *Management Sci.* **50**(10) 1420–1430.

Alagoz, O., L. M. Maillart, A. J. Schaefer, M. S. Roberts. 2007a. Determining the acceptance of cadaveric livers using an implicit model of the waiting list. *Oper. Res.* **55**(1) 24–36.

Alagoz, O., L. M. Maillart, A. J. Schaefer, M. S. Roberts. 2007b. Choosing among living-donor and cadaveric livers. *Management Sci.* **53**(11) 1702–1715.

Alagoz, O., C. L. Bryce, A. J. Schaefer, C. H. Chang, D. C. Angus, M. S. Roberts. 2005. Incorporating biological natural history in simulation models: Empiric estimates of the progression of end-stage liver disease. *Medical Decision Making* **25**(6) 620–632.

Altman, E. 2005. Applications of dynamic games in queues. A. S. Nowak, K. Szajowski, eds. *Advances in Dynamic Games: Applications to Economics, Finance, Optimization, and Stochastic Control.* Birkhäuser, Boston, 309–342.

Cuende, N., B. Miranda, J. F. Canon, G. Garrido, R. Matesant. 2005. Donor characteristics associated with liver graft survival. *Transplantation* **79**(10) 1445–1452.

David, I., U. Yechiali. 1985. A time-dependent stopping problem with application to live organ transplants. *Oper. Res.* **33**(3) 491–504.

David, I., U. Yechiali. 1990. Sequential assignment match processes with arrivals of candidates and offers. *Probab. Engrg. Informational Sci.* **4**(4) 413–430.

David, I., U. Yechiali. 1995. One-attribute sequential assignment match processes in discrete time. *Oper. Res.* **43**(5) 879–884.

Feng, S., N. P. Goodrich, J. L. Bragg-Gresham, D. M. Dykstra, J. D. Punch, M. A. Debroy, S. M. Greenstein, R. M. Merion. 2006. Characteristics associated with liver graft failure: The concept of a donor risk index. *Amer. J. Transplantation* **6** 783–790.

Filar, J., K. Vrieze. 1996. *Competitive Markov Decision Processes*. Springer, New York.

Hassin, R., M. Haviv. 2003. *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Kluwer Academic Publishers, Boston.

HHS. 2005. Annual report of the U.S. Organ Procurement and Transplantation Network and the Scientific Registry of Transplant Recipients: Transplant Data 1995–2004. Tables 1.1, 1.2, 1.7 retrieved December 1, 2006, http://www.ustransplant.org/annual_reports/.

Hornberger, J. C., J.-H. Ahn. 1997. Deciding eligibility for transplantation when a donor kidney becomes available. *Medical Decision Making* **17**(2) 160–170.

Howard, D. H. 2001. Dynamic analysis of liver allocation policies. *Medical Decision Making* **21** 257–266.

Howard, D. H. 2002. Why do transplant surgeons turn down organs? A model of the accept/reject decision. *J. Health Econom.* **21** 957–969.

Iyengar, G. N. 2005. Robust dynamic programming. *Math. Oper. Res.* **30**(2) 257–280.

Karlin, S. 1968. *Total Positivity*, Vol. I. Stanford University Press, Stanford, CA.

NCHS. 2006. Deaths: Final data for 2004. Retrieved December 1, 2006, http://www.cdc.gov/nchs.

Neyman, A., S. Sorin. 2003. *Stochastic Games and Applications*. Kluwer Academic Publishers, Dordrecht, The Netherlands.

Nilim, A., L. E. Ghaoui. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Oper. Res.* **53**(5) 780–798.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York.

Righter, R. 1989. A resource allocation problem in a random environment. *Oper. Res.* **37**(2) 329–338.

Roberts, M. S., D. C. Angus, C. L. Bryce, Z. Valenta, L. Weissfeld. 2004. Survival after liver transplantation in the United States: A disease-specific analysis of the UNOS database. *Liver Transplantation* **10**(7) 886–897.

Roth, A. E., T. Sönmez, M. U. Ünver. 2004. Kidney exchange. *Quart. J. Econom.* **119**(2) 457–488.

Salizzoni, M., A. Franchello, F. Zamboni, A. Ricchiuti, D. Cocchis, F. Fop, A. Brunati, E. Cerutti. 2003. Marginal grafts: Finding the correct treatment for fatty livers. *Transplant Internat.* **16** (7) 486–493.

Shechter, S. M., C. L. Bryce, O. Alagoz, J. E. Kreke, J. E. Stahl, A. J. Schaefer, D. C. Angus, M. S. Roberts. 2005. A clinically based discrete-event simulation of end-stage liver disease and the organ allocation process. *Medical Decision Making* **25**(2) 199–209.

Su, X., S. A. Zenios. 2004. Patient choice in kidney allocation: The role of the queueing discipline. *Manufacturing Service Oper. Management* **6**(4) 280–301.

Su, X., S. A. Zenios. 2005. Patient choice in kidney allocation: A sequential stochastic assignment model. *Oper. Res.* **53**(3) 443–455.

Su, X., S. A. Zenios. 2006. Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Sci.* **52**(11) 1647–1660.

Swani, S., M. L. Puterman, C. B. Weinberg. 2001. Play it again, Sam? Optimal replacement policies for a motion picture exhibitor. *Manufacturing Service Oper. Management* **3**(4) 369–386.

UNOS. 2006a. View data reports. Retrieved December 1, 2006, http://www.unos.org/data/.

UNOS. 2006b. Policy 3.6: Allocation of livers. Retrieved December 1, 2006, http://www.unos.org/resources/.

Vieille, N. 2002. Stochastic games: Recent results. R. J. Aumann, S. Hart, eds. *Handbook of Game Theory with Economic Applications*. Elsevier, Amsterdam, 1833–1850.

Zenios, S. A. 1999. Modeling the transplant waiting list: A queuing model with reneging. *Queuing Systems* **31** 239–251.

Zenios, S. A. 2002. Optimal control of a paired-kidney exchange program. *Management Sci.* **48**(3) 328–342.

Zenios, S. A. 2004. Models for kidney allocation. M. L. Brandeau, F. Sainfort, W. P. Pierskalla, eds. *Operations Research and Health Care: A Handbook of Methods and Applications*. Kluwer Academic Publishers, Boston, 537–554.

Zenios, S. A., G. M. Chertow, L. M. Wein. 2000. Dynamic allocation of kidneys to candidates on the transplant waiting list. *Oper. Res.* **48**(4) 549–569.

Zenios, S. A., L. M. Wein, G. M. Chertow. 1999. Evidence-based organ allocation. *Amer. J. Medicine* **107**(1) 52–61.