Journal of the Operational Research Society (2010) 61, 191–201

www.palgrave-journals.com/jors/

First, do no harm? A framework for evaluating new *versus* reprocessed medical devices

T Sloan*

University of Massachusetts Lowell, Lowell, MA, USA

More than ever before, health care providers are under intense pressure to control costs. Medical devices represent a significant 'hard' cost, with worldwide spending exceeding USD 235 billion. A growing number of health care providers are engaging in the practice of *reprocessing*—sterilizing and reusing medical devices labelled only for a single use. The ethical and technical dimensions of this practice have received much attention, but its economic aspects remain largely unexamined. This paper presents a Markov decision process framework that a health care provider can use to decide whether to use new or reprocessed devices in a given context. Two cases are studied: completely observable device condition and partially observable device condition. After briefly discussing structural results for the two cases, several examples are presented to illustrate how the model can be applied in practice. Useful results can be computed quickly with minimal data. A key insight of the model is that perfect information regarding the device condition is often not required to make a sound decision.

Journal of the Operational Research Society (2010) **61,** 191–201. doi:10.1057/jors.2008.137
Published online 7 January 2009

Keywords: cost benefit; health service; hospitals; replacement policy; risk

1. Introduction

Managing health-sector supply chains presents unique challenges with respect to both quality of care and control of costs. Medical devices account for a significant portion of the 'hard goods' cost incurred by hospitals and other health care providers and therefore represent an important focus area for cost reduction. Spending on medical devices exceeds USD 235 billion worldwide, and that figure is expected to grow substantially in the coming years (Standard & Poor's, 2007). Over the past 30 years, changes in technology and increasing fears of disease transmission have prompted many device makers to shift from multi-use to single-use devices (SUDs). One way in which health care providers are seeking to reduce cost is through *reprocessing*: the practice of refurbishing, sterilizing, and repackaging single-use medical devices. These devices are typically sold back to hospitals at about half the cost of new devices. Recent surveys suggest that this is a significant and growing practice: 45% of medium and large hospitals in the United States (Kerber, 2005), 86% of medium and large hospitals in Canada (MEDEC, 2004), and more than 90% of hospitals in Japan (Koh and Kawahara, 2005) reuse some devices labelled for one use. Although banned in some parts of Europe, reuse of SUDs has been described as 'widespread' in other parts (Hope, 2006).

The practice of reprocessing has mainly been studied from an ethical standpoint (ie *should* it be done?) or from a technical standpoint (ie *how* should it be done?). However, there has been little or no research on the economic aspects of reprocessing. This is perhaps due to the fact that applying operational research models in the health services context is so challenging, as observed by Harper and Pitt (2004). Consider the case of coronary angioplasty, a procedure in which narrow 'balloons' are inflated in a patient's arteries to remove blockages. Often, a patient has several arteries cleared during a single procedure. Angioplasty is performed on millions of patients each year and accounts for billions of dollars worth of health care spending worldwide. Since a new balloon costs up to USD 1000 and reprocessed balloons typically cost less than half this amount, the economic benefits of reprocessing seem obvious. But what about patient safety? How is a health care provider to make a sound decision in this context? At one end of the spectrum, some health care providers and/or regulatory agencies require that a new balloon be used for *each artery*, clearly placing a premium on safety. At the other end of the spectrum, it is common practice in some areas to sterilize and reuse balloons on *other patients*, clearly placing a premium on cost.

The purpose of this paper is not to support a position for or against reprocessing. Rather, the purpose is to develop a model—built on well-established principles of equipment maintenance research—that a health care provider can use to examine the sourcing dimension of medical device reuse. This paper makes three contributions. First, necessary and

**Correspondence: T Sloan, College of Management, University of Massachusetts Lowell, Lowell, MA 02155, USA.*
E-mail: Thomas_Sloan@uml.edu

sufficient conditions that characterize the optimal policy are developed. These conditions go beyond previous operational research models by identifying the exact indifference point between the 'replace' and 'refurbish' decisions. Second, the model reveals situations where information about the device condition has little impact on the optimal decision, which greatly simplifies the health care provider's choice. Finally, and most importantly, the paper provides a framework with which health care providers can evaluate reprocessing decisions. Discussion of reprocessing is often cast in general terms and is often charged with emotion. This framework enables objective examination of specific devices, even without perfect and complete information.

The paper proceeds as follows. The next section presents an overview of the relevant literature. The model is presented in Section 3, and two cases are studied: completely observable device condition and partially observable device condition. Section 4 presents several example applications, demonstrating the use of the model for a wide variety of device types. Conclusions and directions of future research are discussed in Section 5.

2. Literature review

The literature related to this problem can be divided into three main categories: medical device reprocessing, supplier selection, and equipment maintenance. The research on device reprocessing is extensive; however, it tends to focus on very high-level issues or on very low-level issues. For example, there has been much attention on the ethical and legal implications of SUD reprocessing and reuse (eg Gottfried, 2000; Wang, 2001; Dunn, 2002). On the other end of the spectrum, there has also been much effort devoted to exploring the technical aspects of reprocessing, for example, testing different sterilization procedures (Tessarolo *et al*, 2006, 2007) and examining the effects of reprocessing on the device properties (Brown *et al*, 2002; Fedel *et al*, 2006). Both types of research are important, but they do not assist health care providers decide whether to use new or reprocessed devices in a particular context.

A final work on medical devices is that of Bennett *et al* (2005), who develop a model to assess the risk of transmitting disease via surgical instruments. Their focus is on variant *Creutzfeldt–Jakob disease*, which is known to be resistant to standard instrument decontamination procedures. Although their model mainly deals with instruments intended for multiple uses, several insights emerge from their study which are relevant here. First, the risk of person-to-person transmission of disease via medical instruments cannot be ignored. Second, it is extremely challenging to assess the risk of such disease transmission. Third, although patient well-being is of primary importance, economic considerations must be a part of the decision-making process with respect to medical devices.

The second category of relevant research is the work on supplier selection within a supply chain. The main question addressed is how a firm can balance the many criteria by which potential vendors can be judged. These criteria include cost, quality, delivery, flexibility, etc and may be both quantitative and qualitative in nature. This topic has received much attention in the last few years. Many of the papers make use of fuzzy decision models (Amid *et al*, 2006; Chen *et al*, 2006), while others use analytic hierarchy process or analytic network process approaches (Bayazit, 2006; Xia and Wu, 2007), and still others combine one or more of these approaches (Haq and Kannan, 2006; Yang and Chen, 2006). Another approach involves the use of a capability index, which measures the ability of a process to meet design specifications, in combination with cost measures (Chen *et al*, 2005; Linn *et al*, 2006). The existing research sheds much light on the general issues and trade-offs involved in supplier selection; however, sourcing decisions in the medical device context present unique challenges. Specifically, all products must go through extensive certifications, typically by government agencies or their designates. Thus, health care providers have a limited pool of vendors from which to choose. In addition, quality problems may be difficult or impossible to detect before a patient is harmed, so the stakes are much higher in health-care-related sourcing than in most environments.

The third relevant area of research is equipment maintenance models. Wang (2002) presents an extensive survey of this vast literature. The work which relates most closely to the sourcing problem discussed above is research on multi-action maintenance models. Hopp and Wu (1990) study a scenario where the condition of a machine deteriorates over time according to a Markov process, and after observing the state of the equipment, the decision maker can choose to perform one of several maintenance actions, to replace the equipment, or to do nothing. Su *et al* (2000) extend this model by including random failures and by incorporating the notion of imperfect maintenance, that is, performing maintenance may actually make the equipment condition worse.

Another relevant subcategory of equipment maintenance addresses scenarios where the state of the equipment is not completely observable. Monahan (1982) presents an overview of partially observable Markov decision processes (POMDPs) and related literature. Important structural results were developed by White (1979), and efforts to extend these results in various ways continues (Fernández-Gaucherand *et al*, 1991; Makis and Jiang, 2003; Ivy and Pollock, 2005; Hsu *et al*, 2006).

The paper by Hopp and Wu (1988) bridges the completely and partially observable subgroups by modelling a multi-action maintenance problem where the equipment state is only partially observable. The key distinction between Hopp and Wu (1988) and the model presented in the next section is that in their model the equipment state becomes known after maintenance is performed. In our model, however, the state is only known with certainty after the device has been replaced.

Otherwise, the state cannot be known with certainty, which more accurately reflects the medical device context.

3. Model

We model the decision about new *versus* reprocessed medical devices using a discrete-time, discrete-state, Markov decision process (MDP) framework. A medical device can be in one of several states, indexed by i. The states are strictly ordered such that state $i = 0$ indicates the best possible condition (also referred to as 'new' and 'like-new' condition), and state $i = J$ indicates the worst possible condition (also referred to as 'failed' and 'unusable'). At each decision epoch, the decision maker observes the condition of the device and decides which action to take. Initially, we examine the case where the true condition is known with certainty; later, the case of imperfect state information is studied. Possible actions, indexed by a, include using the device to perform the procedure ($a{=}r$), refurbishing the used device ($a{=}u$), and replacing the device with a new one ($a = n$). Performing the medical procedure earns a reward of R, while refurbishing a previously used device incurs a cost of U, and replacing the device with a new one costs N. Whether the device is new or used, there is a possibility that a failure occurs, for example, the device malfunctions or some other device-related harm comes to the patient. In the event of a failure, a cost of C is incurred.

The set of states is denoted as $\mathscr{S} = \{0, 1, \ldots, J\}$, and the set of actions available in state i is denoted as $\mathscr{A}_i \subseteq \{r, u, n\}$. It is assumed that the medical procedure is only performed when the process (device) is in like-new condition; in addition, it is unnecessary to refurbish or replace the device when it is in the best state. Therefore, performing the procedure is the only action available in state 0 (ie $\mathscr{A}_0 = \{r\}$). Performing the procedure causes the device to make a transition out of the like-new state into a worse state (ie $p^r_{00} = 0$, and $p^r_{0j} > 0$ for some $j \neq 0$). When the process is in the worst state, then replacement is the only action available, and the process returns to state 0 with probability one (ie $\mathscr{A}_J = \{n\}$, and $p^n_{J0} = 1$). In the intermediate states, performing the procedure is not allowed due to the potential harm to the patient; thus, the decision maker may choose either to replace or refurbish the device (ie $\mathscr{A}_j = \{u, n\}$ for $0 < j < J$). Reprocessing improves the condition of the device—or at least does not make it worse (ie $p^u_{ij} > 0$ for some $j < i$, and $p^u_{ij} = 0$ for all $j > i$, where $0 < i < J$)—while replacing the device always returns the process to state 0 with probability one (ie $p^n_{i0} = 1$ for all i).

Letting X_t denote the state of the process at decision epoch t and a_t denote the action taken at decision epoch t, we define the probability of making a transition from state i to state j when action a is taken as p^a_{ij}. The state at the next decision epoch depends only on the state and action at the current epoch, so $p^a_{ij} \equiv \Pr\{X_{t+1} = j | X_t = i, a_t = a\} = \{X_{t+1} = j | X_t = i, a_t = a; X_{t-1}, a_{t-1}; \ldots; X_1, a_1\}$. Figure 1 shows a state transition diagram for a four-state problem using

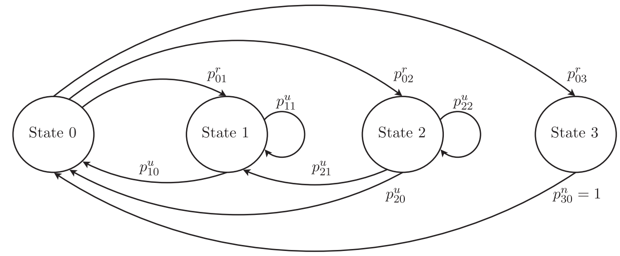

Figure 1 State transition diagram for policy $[r, u, u, n]$.

policy $[r, u, u, n]$, that is, performing the procedure in state 0, choosing to refurbish the device if the state is 1 or 2, and choosing to replace the device if the state is 3.

Note that several implicit assumptions have been made in order to formulate a tractable model. For example, we have defined the process state so that it is not tied to a particular device. Typically, third-party reprocessing firms receive used devices of many different types (and levels of wear) from many different hospitals, and they also supply refurbished devices to many different customers. Thus, hospitals are unlikely to get back the same devices which they sent out to be reprocessed. Treating the state as something that is not connected to a particular device enables the model to account for the randomness of different 'streams' through the use and reuse cycles.

In addition, the state transition probabilities are fixed and depend only on the current action and state. In reality, the transition from state to state may be difficult to predict due to variables such as the type of procedure being performed, the patient, and the clinician using the device. Additionally, the cost of failure may not be known with certainty as we have assumed here. However, similar assumptions regarding costs and state transitions are needed for equipment maintenance models (such as those discussed in the last section), and valuable lessons have been gleaned from these models. Our goal is to develop similar insights in the context of medical device reprocessing.

3.1. *Completely observable states*

To begin our analysis, we assume that the condition of the device is completely observable; that is, the true state of the process is known at each decision epoch. If the process is in state i and action a is taken at the current decision epoch, then the process makes a transition to state j with probability p^a_{ij}. Since the expected rewards/costs and state transitions depend only on the current state and action taken, the process can be modelled as an MDP.

The objective is to find a decision rule that minimizes the long-run expected average cost per unit time. Since all expected costs are bounded and the number of states is finite, an optimal stationary policy exists; furthermore, the optimal policy is stationary, that is, it is time-independent (Heyman and Sobel, 1984). Define $\mathbf{A} = [a_0, a_1, \ldots, a_J]$ as a stationary policy, that is a decision rule which specifies that action a_i be taken when the process is in state i. Any stationary policy (based on the permissable actions identified above) will induce a Markov chain with a unique set of steady-state probabilities that are independent of the initial state of the process. Define $\pi_i(\mathbf{A})$ as the steady-state probability of being in state i when policy $\mathbf{A}$ is used. Let $C(i, a)$ denote the expected cost of performing action a while the process is in state i. The minimal expected average cost per unit time while using policy $\mathbf{A}$ can then be expressed as

$$g(\mathbf{A}) = \sum_{i=0}^{J} C(i, a_i)\pi_i(\mathbf{A}). \tag{1}$$

Note that if action r is taken, then an expected reward of R is earned; in this cost-minimization framework, the value of R would be less than zero, indicating a negative 'cost'. A policy $\mathbf{A}^*$ is average cost optimal if $g(\mathbf{A}^*) \leqslant g(\mathbf{A})$ for each stationary policy $\mathbf{A}$. The optimal action in state i is defined as a_i^*.

Previous work has focussed on the structure of the optimal policy rather than on the solution. For example, one might be interested in knowing conditions that ensure monotonicity of the actions, that is, where increasingly effective (and costly) maintenance actions are chosen as the process state deteriorates. However, in the context of medical devices, where the stakes are so high, it is useful to know the exact indifference point between new and reprocessed devices.

The key question for a health care provider in this context is: Under what conditions is it preferable to use a refurbished device rather than a new device? The key variable that differentiates the two device types is the probability of returning to like-new condition. Choosing a new device returns the process to state 0 with probability one, while choosing a reprocessed device may not. Thus, we seek to determine the exact indifference point with respect to p_{i0}^u, the probability that reprocessing returns the device to like-new condition. This is accomplished by employing an approach similar to that of Kazaz and Sloan (2008): First, express the objective function (1) for two different policies in terms of the steady-state probabilities; next, compare two policies that differ by only one action in one state, identifying the value of p_{i0}^u for which one is indifferent between the two.

For ease of exposition, let us examine the case with four possible states—labelled 0 through 3—where the states are ordered such that state 0 is like-new and state 3 is failed. To begin our analysis, we compare two policies which differ only by the device type used in state 1: $\mathbf{A}_1 = [r, n, n, n]$ and $\mathbf{A}_2 = [r, u, n, n]$. Define $\rho_i^*(\mathbf{A}_k, \mathbf{A}_l)$ as the p_{i0}^u value at which the decision maker is indifferent between policies $\mathbf{A}_k$ and $\mathbf{A}_l$. The following expression identifies the p_{10}^u value at which the decision maker is indifferent between policies $\mathbf{A}_1$ and $\mathbf{A}_2$:

$$\rho_1^*(\mathbf{A}_1, \mathbf{A}_2) = \frac{2U - N - R - Cp_{03}^r}{N - R - Cp_{03}^r}. \tag{2}$$

When $p_{10}^u > \rho_1^*(\mathbf{A}_1, \mathbf{A}_2)$, using a refurbished device is optimal in state 1; when $p_{10}^u < \rho_1^*(\mathbf{A}_1, \mathbf{A}_2)$, using a new device is optimal in state 1 (formal statements of this and other technical results appear in Appendix A). Comparing policies $\mathbf{A}_1 = [r, n, n, n]$ and $\mathbf{A}_4 = [r, n, u, n]$ yields the same result for state 2.

The indifference point with respect to p_{20}^u is determined by comparing two policies which differ only by the device type used in state 2: $\mathbf{A}_2 = [r, u, n, n]$ and $\mathbf{A}_3 = [r, u, u, n]$. This indifference point is written as

$$\rho_2^*(\mathbf{A}_2, \mathbf{A}_3) = \frac{(p_{10}^u + p_{21}^u)(2U - N - R - Cp_{03}^r) - p_{10}^u p_{21}^u (N - R - Cp_{03}^r) + p_{01}^r p_{10}^u (N - U)}{p_{10}^u (N - R - Cp_{03}^r) + p_{01}^r (N - U)}. \tag{3}$$

When $p_{20}^u > \rho_2^*(\mathbf{A}_2, \mathbf{A}_3)$, using a refurbished device is optimal in state 2; when $p_{20}^u < \rho_2^*(\mathbf{A}_2, \mathbf{A}_3)$, using a new device is optimal in state 2. Comparing policies $\mathbf{A}_3 = [r, u, u, n]$ and $\mathbf{A}_4 = [r, n, u, n]$ yields a similar result.

Equipped with these expressions, the decision maker can easily determine the optimal policy for a given set of parameter values for the completely observable case. Next, we examine the situation in which perfect information about the device condition is not available.

3.2. Partially observable states

No health care provider would endanger a patient by knowingly using a damaged or contaminated device. However, the true state of the device may not observable to the naked eye or without disassembling the device, as documented by cases reported in the popular press (Kerber, 2005; Klein, 2005) and medical literature (Roth *et al*, 2002; Tessarolo *et al*, 2007). To account for the possibility of imperfect information, we reframe the decision problem as a POMDP. Readers are referred to Monahan (1982) for an overview of POMDPs. The key difference between this model and the completely observable Markov decision process (COMDP) in the last section is the addition of an observation process. As before, the state of the process at epoch t is denoted by X_t. However, since it is not directly observed, we also define an observation process, the state of which is denoted as Y_t at epoch t. The state of the core process continues to make transitions according to the probabilities denoted by p_{ij}^a. We adopt the standard convention that after an action is selected, the state of the core process makes a transition, and then an observation is made. The probability of observing state k when the state of the core process is j and action a was taken last is defined as $q_{jk}^a \equiv \Pr\{Y_t = k | X_t = j, a_{t-1} = a\}$, where $k \in \mathscr{S}$. Thus, the decision maker does not know the true state of the process but does know the probabilistic relationship between X_t and Y_t.

Building on the notation from the last section, let $C(i, j, k, a)$ denote the immediate cost incurred at epoch $t+1$ when at epoch t the core state was i, observation k was made, action a was taken, and the process then made a transition to state j. The expected cost for taking action a when the core state is i at the current epoch is simply $C(i, j, k, a)$ weighted by the various transition and observation probabilities and is defined as follows:

$$\bar{C}(i, a) = \sum_{j \in \mathscr{S}} \sum_{k \in \mathscr{S}} C(i, j, k, a) p_{ij}^{a} q_{jk}^{a}. \tag{4}$$

With these new definitions in hand, we may now write the POMDP average cost optimality equation using the standard recursive form:

$$g(\boldsymbol{\alpha}) + h(\boldsymbol{\alpha}) = \min_{a} \left\{ \sum_{i=0}^{J} \alpha_i \bar{C}(i, a) + \sum_{k=0}^{J} \left(\sum_{l \in \mathscr{S}} q_{lk}^{a} \sum_{i \in \mathscr{S}} \alpha_i p_{il}^{a} \right) h[T(\boldsymbol{\alpha}|a, k)] \right\}, \tag{5}$$

where $g(\cdot)$ is the expected average cost per unit time, $h(\cdot)$ is a difference function, $\boldsymbol{\alpha}$ is a belief vector, and $T(\cdot)$ is a transformation function (see Sloan (2008) for details). The theory regarding average cost POMDPs is not nearly as well developed as that of COMDPs, particularly with respect to the structural properties of optimal solutions. Nevertheless, it can be shown that a stationary optimal policy exists and that the average expected minimal cost per unit time is constant (Fernández-Gaucherand *et al*, 1991). Despite the paucity of structural results for the partially observed case, we can still gain significant insights into the decision process, as shown in the examples in the next section.

4. Example applications

How can a health care provider use these models to shed light on the question of new *versus* refurbished devices? This section presents several example applications to illustrate the kinds of analyses which can be performed using the models from the last section. Ideally, the example problems would be based on empirical data regarding the devices studied. Unfortunately, such data are not readily available for all model parameters. Indeed, device makers and reprocessors alike have an incentive to keep some information confidential. One motivation for the development of a model is to help explore this decision in the absence of high-quality, detailed information about the condition of the devices. As we shall demonstrate shortly, decision makers need not have complete information to make sound decisions regarding medical device reuse. The model enables health care providers to identify broad classes of devices that probably should not be reprocessed—that is, for which the risk of adverse effect greatly outweighs the potential cost savings. In this way, the debate about reprocessing can be re-framed from one of ‘yes or no’ to one of ‘which devices’. Shifting the debate from the general to the specific can create incentives for device makers, reprocessors, regulators, and health care providers to engage in a more open discussion of the practice. This kind of discussion, in turn, can pave the way to collecting the kind of detailed information necessary to study these issues more fully.

The example problems below cover a wide range of device types (similar to the examples in Sloan, 2007); however, some parameters are fixed for all problems to conserve space. All problems have four device states, ordered from 0 to 3, where state 0 is new (or like-new) and state 3 is failed (or unusable). The use of a discrete state space is an approximation—in reality the device condition is probably a continuum. However, this approximation makes the solution much easier, while capturing sufficient detail about the real process. The number of states can easily be changed, depending on the level of detail desired in a given situation. Preliminary analyses revealed that using more than four states in the examples below slows down the solution process without producing additional insight. Using fewer than four states oversimplifies the problem.

The exact definition of a state will depend on the specific device being modelled and may be composed of multiple dimensions. For example, the state of an angioplasty balloon could refer to its sterility as well as its structural integrity. Tessarolo *et al* (2006) use four measures of structural properties: crossing profile, slipperiness, compliance, and burst pressure. A detailed discussion of these measures is beyond the scope of this paper, but they are well known and easily evaluated. Using these or similar measures, an angioplasty balloon that is within the tolerances allowed for a new device would be classified as like-new (state 0). States 1, 2, and 3 would refer to increasing levels of variation beyond these tolerances (where the precise levels would be defined by experts). Thus, reaching the ‘failed’ state may not mean a device malfunction but merely a level of wear or contamination which makes the device unusable.

For all problems, replacement returns the process to state 0, that is, $p_{i0}^{n} = 1$ for all i. With respect to observation probabilities, perfect information quality for a given action a means that $q_{jj}^{a} = 1$, and very high quality means that $q_{jj}^{a} = 0.99$ for all j. Unless otherwise noted, new and reprocessed device cost estimates come from Klein (2005) and Landro (2008). All cost figures are reported in US dollars. Complete data for each example problem are reported in Appendix B.

4.1. *Angioplasty balloon*

Percutaneous transluminal coronary angioplasty is a procedure used to clear a blocked artery by inserting a catheter into the artery and then inflating a ‘balloon’ that is mounted on the tip of the catheter. New angioplasty balloons range in cost, depending on the type. Suppose that a new balloon costs 515, while a reprocessed balloon costs only about 250—clearly a very substantial cost reduction. On the other hand, studies have shown that reusing balloons is associated with longer procedure times, which in turn is associated with an increase in the rates of adverse effects such as death, myocardial infarction, and the need for further revascularization (Mak *et al*, 1996). Therefore, reusing the balloons also comes with a cost.

Suppose that the probability of an adverse effect is one in a million and that the cost of such an effect is 1 000 000. Note that the 'adverse effect' may not be outright failure but rather any of the unfavourable outcomes listed above, which could be attributed to the device. The charges for this type of procedure are substantial—an average of 16 000 in the US (Mak *et al*, 1996).

The most conservative policy calls for replacement of the device after each procedure, which is policy $\mathbf{A}_1 = [r, n, n, n]$. Perhaps the health care provider could reduce costs by performing the procedure when the device state is 0 (like-new) and refurbishing the device when it reaches state 1. Beyond state 1, the device is replaced. This policy—still rather conservative—is expressed as $\mathbf{A}_2 = [r, u, n, n]$. How effective does refurbishing need to be in order to make a used device preferable to a new device? In other words, what value of p^u_{10} is sufficient to make policy $\mathbf{A}_2 = [r, u, n, n]$ optimal? The parameter values given above are sufficient to answer this question. Plugging the numbers into Equation (2) reveals that $\rho^*_1(\mathbf{A}_1, \mathbf{A}_2) \approx 0.9679$. This result tells us that there must be nearly a 97% chance of returning the device to like-new condition for reprocessing to be optimal in state 1. Otherwise, it will be optimal to replace the device with a new one before each procedure.

Suppose that reprocessing in state 1 has a 98% success rate (ie $p^u_{10} = 0.98$). For the completely observable case, policy $\mathbf{A}_2 = [r, u, n, n]$ is optimal since $p^u_{10} = 0.98 > 0.9679 = \rho^*_1(\mathbf{A}_1, \mathbf{A}_2)$. However, when the device state is only partially observable after performing a procedure, then the POMDP *never* calls for reprocessing—even with a very high probability of knowing the true state of the device after reprocessing. That is, policy $\mathbf{A}_1 = [r, n, n, n]$ is optimal for all cases tested such that the information quality is less than perfect. Table 1 reports the complete results.

This example supports the intuitive idea that when the cost of failure is very high, then one must have great confidence in the outcome of reprocessing. In the absence of such confidence, the health care provider is probably better off simply using new devices for this procedure. Resources can be better allocated to assessing risk and improving processes for other device types.

Since the numbers used in this example are approximate, performing some sensitivity analysis is wise. Figure 2 illustrates how the indifference point (ρ^*) changes as a function of the device failure probability (p^r_{03}). The figure suggests that greater device reliability is associated with higher reprocessing standards; that is, as the p^r_{03} value decreases, the ρ^* value increases, indicating a need for a higher repair probability. Similar analyses can easily be performed with respect to other parameters.

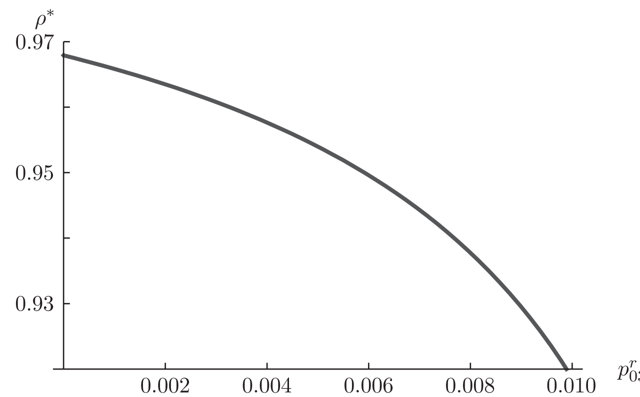

Figure 2 ρ^* as a function of p^r_{03}.

4.2. Pulse oximetry sensor

Pulse oximetry is a technique used to measure the amount of oxygen in a patient's blood. This is accomplished by placing a small, plastic sensor on the patient's fingertip, which is attached to a piece of equipment that measures waves of infrared light. The sensors themselves are quite inexpensive: 10 for a new one and six for a reprocessed one. Suppose that the probability of failure is $1e-09$, the cost of failure is 10 000, and the reward for the procedure is 15.

Suppose that success rate for reprocessing is 99% in state 1 and 85.65% in state 1 (ie $p^u_{10} = 0.99$ and $p^u_{20} = 0.8565$). Entering the above values into Equation (2) yields $\rho^*_1(\mathbf{A}_1, \mathbf{A}_2) = 0.68$. Thus, only a 68% success rate is required in this instance to make reprocessing the optimal choice in state 1. Similarly, Equation (3) yields $\rho^*_2(\mathbf{A}_2, \mathbf{A}_3) \approx 0.68$. Since $p^u_{10} > \rho^*_1(\mathbf{A}_1, \mathbf{A}_2)$ and $p^u_{20} > \rho^*_2(\mathbf{A}_2, \mathbf{A}_3)$, policy $\mathbf{A}_3 = [r, u, u, n]$ is optimal for the completely observable case.

Not surprisingly, reprocessing is also the optimal choice in state 1 in the partially observable case even when the information quality after performing the procedure is very low. The optimal choice in state 2 varies as the observation probabilities change. Complete results are reported in Table 1. Clearly, two factors are at work here: the high probability of returning the device to an acceptable state and the relatively low cost of adverse effect. Sensitivity analyses with respect to these and other parameters could easily be performed. Even without such analyses, however, the results clearly suggest that reprocessing is a viable choice for this scenario, which makes sense given the nature of the device.

4.3. Orthopaedic blade

Blades are commonly used to cut bone and tissue during orthopaedic surgery. A new orthopaedic blade costs about 30, while a refurbished blade costs about 15. Suppose that the chance of adverse effect (attributable to the device itself)

Table 1 Example Problems—optimal policies for different information quality levels

Information quality after action		*Example problem*			
Procedure	*Reprocessing*	1	2	3	4
Perfect	Perfect	$[r, u, n, n]$	$[r, u, u, n]$	$[r, u, n, n]$	$[r, u, u, n]$
Perfect	Very high	$[r, n, n, n]$	$[r, u, u, n]$	$[r, u, n, n]$	$[r, u, u, n]$
Very high	Very high	$[r, n, n, n]$	$[r, u, u, n]$	$[r, n, n, n]$	$[r, u, u, n]$
High	Very high	$[r, n, n, n]$	$[r, u, u, n]$	$[r, n, n, n]$	$[r, n, n, n]$
Medium	Very high	$[r, n, n, n]$	$[r, n, n, n]$	$[r, n, n, n]$	$[r, n, n, n]$
Low	Very high	$[r, n, n, n]$	$[r, n, n, n]$	$[r, n, n, n]$	$[r, n, n, n]$

Note: All example problems have four device states, ordered from 0 to 3, where 0 is like-new and 3 is failed. The policies indicate the optimal action for each state, $[a_0, a_1, a_2, a_3]$, where a_i is the action specified for state i. A policy for the completely observable case (perfect information) specifies the action taken for the actual state. A policy for a partially observable case specifies the action taken for the *observed* state. The shaded cell for each example indicates the information quality level at which reprocessing is no longer optimal for any state.

is only one in 100 million and that the cost of a failure is 100 000. In addition, suppose that a modest reward of 650 is earned for performing the procedure.

Using the above values in Equation (2), we can determine that $\rho_1^*(\mathbf{A}_1, \mathbf{A}_2) \approx 0.9559$. This result means that reprocessing must have approximately a 96% chance of returning the device to state 0 for that to be the optimal choice in state 1. Otherwise, it will be optimal to replace the device with a new one before each procedure. Suppose that $p_{10}^u = 0.98$; since $p_{10}^u > \rho_1^*(\mathbf{A}_1, \mathbf{A}_2)$, reprocessing in state 1 is optimal for the completely observable case. In the partially observable case, reprocessing is also optimal in state 1 when there is perfect information about the device state after the procedure and very high-quality information about the device state after reprocessing (ie when $q_{jj}^u = 0.99$). However, if there is even slight uncertainty about the device condition after the procedure, then reprocessing is no longer optimal. Table 1 reports the results for all information quality levels tested.

This example illustrates that the decision about reprocessing cannot be driven by device cost nor by failure cost alone. It is difficult to determine the best choice without understanding the connections between various parameters.

4.4. Laparoscopic surgery

Laparoscopic surgery is a minimally invasive technique used for a wide range of procedures. Here we focus on cholecystectomy (removal of the gall bladder). Surgery is performed by first making several small incisions in a patient's abdomen and then inserting special instruments—including graspers, scissors, and a telescopic lens (laparoscope) which is connected to a video camera. After insufflating the patient's abdomen with carbon dioxide, the surgeon is able to use the video to guide the surgery. New instruments for laparoscopic cholecystectomy cost about 1200, while reprocessed devices cost only about 250 (Jacobs and Noorani, 2008)—obviously a major cost savings.

Suppose that the device failure probability is one in 10 million, the cost of adverse effect is 1 000 000, and the reward for performing the procedure is 8000. The indifference point for state 1 is $\rho_1^*(\mathbf{A}_1, \mathbf{A}_2) \approx 0.7935$, which is computed using Equation (2). Similarly, Equation (3) tells us that $\rho_2^*(\mathbf{A}_2, \mathbf{A}_3) \approx 0.792$. Thus, reprocessing must have a success probability greater than approximately 79% for that to be the optimal choice in states 1 and 2; otherwise, replacement is optimal.

Suppose that reprocessing has the following success rates: $p_{10}^u = 0.8609$ and $p_{20}^u = 0.8354$. Both p_{10}^u and p_{20}^u are high enough to warrant reprocessing in the completely observable case, and therefore policy $\mathbf{A}_3 = [r, u, u, n]$ is optimal. In the partially observable case, however, the policy depends heavily on the information quality. When the information quality level after the procedure is very high, then $[r, u, u, n]$ remains optimal. However, when the information quality level is less than very high, then $[r, n, n, n]$ becomes the optimal policy—a dramatic shift from the completely observable case. The results for all information quality levels tested are reported in Table 1.

4.5. Managerial insights

These examples illustrate how the model works for a wide variety of device types. Two main conclusions can be drawn. First, the quality of information about devices has an important impact on the optimal choice. Thus, the results from the COMDP model are an excellent indicator of when reprocessing should *not* be pursued. In short, if it is not optimal to reprocess in the completely observable case, then it will not be optimal in the partially observable case either. The completely observable problem requires minimal data and is easily solved using MS Excel or other readily available software. Therefore, a health care provider could examine a number of devices/cases with minimal effort.

The second main conclusion that can be drawn is that no obvious patterns exist to tell when reprocessing is warranted. For example, one might expect common-sense rules of thumb to emerge, such as 'If the device failure probability is low, then reprocessing is advisable'. However, the blade (Example 3) has a lower probability of adverse effect than the laparoscope (Example 4)—yet the results show that reprocessing is optimal in more cases for the laparoscope than for the blade. One might expect the difference between new and reprocessed device costs to determine the desirability of reprocessing. However, the percentage discount for the reprocessed angioplasty balloon (Example 1) is greater than the discount for the oximetry sensor (Example 2), showing that this rule of thumb does not hold either. These results illustrate the challenges of making broad generalizations and demonstrate the need for rigorous study of different device types and scenarios.

Up to now, health care providers have not approached reprocessing decisions in a systematic way. Decisions have generally been made by an external government agency or on an *ad hoc* basis by individuals within the health care system, such as hospital administrators or physicians. In some cases health care providers have made a ‘global’ choice, accepting all reprocessed devices or accepting none. As one might expect, economic factors point one way and safety concerns point the other. As our model and examples illustrate, however, no single factor can adequately capture the different dimensions of this complex problem. Equipped with this model, a decision maker can undertake more sophisticated analyses.

The lessons of the model, however, extend beyond the boundaries of a single health care provider. The insights revealed by the model can be used by the provider to negotiate prices with the device makers and/or reprocessors. The device makers can use the model to guide improvements in current devices, perhaps changing designs to improve a device’s reusability.

The model can be used to help design a supplier qualification programme, requiring reprocessors to provide more information about their processes and results. The example problems illustrate that the post-procedure information quality can have a significant impact on the results. This finding can serve as the impetus for the health care provider and reprocessor to improve the collection of used devices. Clearly there is a difference between a device that has only been removed from its package and one that has actually been used on a patient. Currently, however, many providers do not sort ‘used’ devices, meaning that valuable information is lost.

Regulators can also benefit from the model’s insights. Currently, there are few distinctions between device types with respect to reprocessing regulations. In some countries, reprocessing is banned outright; in others, the decision is left up to the health care providers and reprocessors. As the examples above suggest, perhaps some device classes should not be reprocessed; however, banning all reprocessing may not be the best choice either. A government agency may be the only entity with the necessary perspective and authority to make such judgements. This type of high-level analysis can have implications for research and development, both in terms of the devices and their sterilization/reprocessing protocols.

In summary, the model provides a flexible, easy-to-use framework for examining reprocessing decisions. Its main utility is not in identifying the optimal policy for a particular device; rather, it is best suited to provide rough-cut analysis for different device classes. Ruling out reprocessing as a viable alternative for some device types can free up resources for the design, regulation, and reprocessing of other device types.

5. Conclusions

Medical devices represent a significant, and growing, ‘hard good’ expenditure for health care providers. These providers are under increasing pressure to contain costs in an industry that faces unique financial challenges. One approach to cutting costs has been to sterilize and refurbish devices that are labelled for a single use. This practice, known as *reprocessing*, has been studied from ethical and technical perspectives, but relatively little effort has been expended to examine the economic and operational aspects of this important issue.

An MDP model was formulated to apply the time-tested principles of equipment maintenance research to the decision about whether a health care provider should use new or refurbished devices. The optimal policy was characterized by finding the exact point of indifference between ‘replace’ and ‘refurbish’ decisions. The model was then extended to a POMDP, accounting for the fact that the true state of the device may not be easily observed.

Example applications covering a wide range of device types were presented. These examples explain how device states can be defined, illustrate the kinds of analyses possible with the model, and demonstrate how the model could be put into practice. The examples show that if the MDP model indicates that reprocessing is not optimal, then it is not warranted in the POMDP scenario either. Not surprisingly, even when very high-quality information about the true state is available, the optimal POMDP policy is quite conservative. Thus, health care providers need not worry about obtaining optimal solutions to the POMDP for all devices—the high-level results provided by the completely observable MDP can be used to identify classes or families of products for which reprocessing is not viable. This finding represents the main managerial insight of the model: Perfect information about the device state and/or infection risk is generally not necessary to make a sound decision regarding reprocessed devices.

Assessment of infection risk is challenging and costly (Bennett *et al*, 2005), so eliminating some device types from consideration can save substantial resources. Creating a flexible framework with which health care providers can perform these kinds of assessments for specific device types is the main contribution of the paper. Until now, most discussion of this important health care issue has been in very general terms and has primarily been based on appeals to emotion.

Future research in several areas would be useful. First, a more extensive experimental study—perhaps focusing on a particular class of devices—would be of value. Second, examining more than two 'maintenance' actions would be of interest. Finally, it would be of great interest to collect data on the failure and observation probabilities for actual devices. These data are not readily available at present, and collecting such data would undoubtedly require the cooperation of original equipment manufacturers and device reprocessors—two groups that are not inclined to work together. Nevertheless, it is hoped that the insights revealed by this paper will provide an incentive for health care providers to find ways to ascertain the relevant data and use the framework to make better decisions.

References

Amid A, Ghodsypour SH and OBrien C (2006). Fuzzy multiobjective linear model for supplier selection in a supply chain. *Int J Prod Econ* **104**: 394–407.

Bayazit O (2006). Use of analytic network process in vendor selection decisions. *Benchmarking: Int J* **13**: 566–579.

Bennett P, Hare A and Townshend J (2005). Assessing the risk of vCJD transmission via surgery: Models for uncertainty and complexity. *J Opl Res Soc* **56**: 202–213.

Brown SA, Merritt K, Woods TO, McNamee SG and Hitchins VM (2002). Effects of different disinfection and sterilization methods on tensile strength of materials used for single-use devices. *Biomed Instrument Technol* **36**: 23–27.

Chen CT, Lin CT and Huang SF (2006). A fuzzy approach for supplier evaluation and selection in supply chain management. *Int J Prod Econ* **102**: 289–301.

Chen KL, Chen KS and Li RK (2005). Suppliers capability and price analysis chart. *Int J Prod Econ* **98**: 315–327.

Dunn D (2002). Reprocessing single-use devices—The ethical dilemma. *AORN J* **75**: 989–999.

Fedel M, Tessarolo F, Ferrari P, Losche C, Ghassemieh N, Guarrera GM and Nollo G (2006). Functional properties and performance of new and reprocessed coronary angioplasty balloon catheters. *J Biomed Mater Res Part B, Appl Biomater* **78B**: 364–372.

Fernández-Gaucherand E, Arapostathis A and Marcus SI (1991). On the average cost optimality equation and the structure of optimal policies for partially observable Markov decision processes. *Ann Opns Res* **AC-29**: 439–470.

Gottfried KL (2000). Are reprocessed medical devices really as good as new, and does it matter? *HEC Forum* **12**: 311–316.

Haq AN and Kannan G (2006). Design of an integrated supplier selection and multi-echelon distribution inventory model in a built-to-order supply chain environment. *Int J Prod Res* **44**: 1963–1985.

Harper PR and Pitt MA (2004). On the challenges of healthcare modelling and a proposed project life cycle for successful implementation. *J Opl Res Soc* **55**: 657–661.

Heyman DP and Sobel MJ (1984). *Stochastic Models in Operations Research, Volume II: Stochastic Optimization*. McGraw-Hill: New York.

Hope J (2006). Cleared, surgeon who reused his tools to save cash. *Daily Mail (London)* 4 July, p 12.

Hopp WJ and Wu SC (1988). Multiaction maintenance under Markovian deterioration and incomplete state information. *Nav Res Log* **35**: 447–462.

Hopp WJ and Wu SC (1990). Machine maintenance with multiple maintenance actions. *IIE Trans* **22**: 226–232.

Hsu SP, Chuang DM and Arapostathis A (2006). On the existence of stationary optimal policies for partially observed MDPs under the long-run average cost criterion. *Syst Control Lett* **55**: 165–173.

Ivy JS and Pollock SM (2005). Marginally monotonic maintenance policies for a multi-state deteriorating machine with probabilistic monitoring and silent failures. *IEEE Trans Reliab* **54**: 489–497.

Jacobs P and Noorani H (2008). Re-use of single-use medical devices in Canada: Using evidence to inform policy. *Proceedings of the 2008 CADTH Invitational Symposium*. Canadian Agency for Drugs and Technologies in Health: Carling Avenue, Ottawa, Ontario, pp 600–865. http://cadth.ca/index.php/en/events/sympos-2008, accessed November 2008.

Kazaz B and Sloan T (2008). Production policies under deteriorating process conditions. *IIE Trans* **40**: 187–205.

Kerber R (2005). Device makers fight reuse of surgical tools. *Boston Globe*, 19 October, p D.1.

Klein A (2005). Hospitals save money, but safety is questioned. *Washington Post*, 11 December, p A1.

Koh A and Kawahara K (2005). Current practices and problems in the reuse of single-use devices in Japan. *J Med Dental Sci* **52**: 81–89.

Landro L (2008). The informed patient: Hospitals reuse medical devices to lower costs. *Wall Street Journal*, 19 March, p D.1.

Linn RJ, Tsung F and Ellis LWC (2006). Supplier selection based on process capability and price analysis. *Qual Eng* **18**: 123–129.

Mak KH, Eisenberg MJ, Eccleston DS, Brown KJ, Ellis SG and Topol EJ (1996). Cost-efficacy modeling of catheter reuse for percutaneous transluminal coronary angioplasty. *J Am Coll Cardiol* **28**: 106–111.

Makis V and Jiang X (2003). Optimal replacement under partial observations. *Math Opns Res* **28**: 382–394.

MEDEC (2004). *The reuse of single use medical devices*. Technical Report. MEDEC: Canada's Medical Device Technology Companies, Toronto, Canada.

Monahan G (1982). A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Mngt Sci* **28**: 1–16.

Roth K, Heeg P and RR (2002) Specific hygiene issues relating to reprocessing and reuse of single-use devices for laparoscopic surgery. *Surg Endoscop* **16**: 1091–1097.

Sloan T (2007). Safety-cost trade-offs in medical device reuse: A Markov decision process model. *Health Care Mngt Sci* **10**: 81–93.

Sloan T (2008). *First do no harm? A framework for evaluating new vs. reprocessed medical devices*. Working Paper. University of Massachusetts Lowell: Lowell, MA.

Standard & Poor's (2007). Healthcare: Products and supplies. *Standard & Poor's Ind Surv* **175**: 14–15.

Su CT, Wu SC and Chang CC (2000). Multiaction maintenance subject to action-dependent risk and stochastic failure. *Eur J Opl Res* **125**: 133–148.

Tessarolo F, Caola I, Nollo G, Antolini R, Guarrera GM and Caciagli P (2006). Efficiency in endotoxin removal by a reprocessing protocol for electrophysiology catheters based on hydrogen peroxide plasma sterilization. *Int J Hygiene Environment Health* **209**: 557–565.

Tessarolo F, Caola I, Fedel M, Stacchiotti A, Caciagli P, Guarrera GM, Motta A and Nollo G (2007). Different experimental protocols for decontamination affect the cleaning of medical devices: Preliminary electron microscopy analysis. *J Hosp Infect* **65**: 326–333.

Wang EP (2001). Regulatory and legal implications of reprocessing and reuse of single-use medical devices. *Food Drug Law J* **56**: 77–98.

Wang H (2002). A survey of maintenance policies of deteriorating systems. *Eur J Opl Res* **139**: 469–489.

White CC (1979). Optimal control-limit strategies for a partially observed replacement problem. *Int J Syst Sci* **10**: 321–331.

Xia W and Wu Z (2007). Supplier selection with multiple criteria in volume discount environments. *Omega* **35**: 494–504.

Yang CC and Chen BS (2006). Supplier selection using combined analytical hierarchy process and grey relational analysis. *J Manuf Technol Mngt* **17**: 926–941.

Appendix A. Main technical results

The main technical results of the model are summarized here. For more details and for proofs, please refer to Sloan (2008).

A.1. COMDP: indifference points

Proposition A.1 *With respect to the repair probability of a used device* (p_{10}^u), *comparing policies* $\mathbf{A}_1 = [r, n, n, n]$ *and* $\mathbf{A}_2 = [r, u, n, n]$ *yields the following indifference point*:

$$\rho_1^*(\mathbf{A}_1, \mathbf{A}_2) = \frac{2U - N - R - Cp_{03}^r}{N - R - Cp_{03}^r}. \tag{A.1}$$

When $p_{10}^u > \rho_1^*(\mathbf{A}_1, \mathbf{A}_2)$, *using a refurbished device is optimal in state 1; when* $p_{10}^u < \rho_1^*(\mathbf{A}_1, \mathbf{A}_2)$, *using a new device is optimal in state 1.*

Proposition A.2 *With respect to the repair probability of a used device* (p_{20}^u), *comparing policies* $\mathbf{A}_2 = [r, u, n, n]$ *and* $\mathbf{A}_3 = [r, u, u, n]$ *yields the following indifference point*:

$$\rho_2^*(\mathbf{A}_2, \mathbf{A}_3) = \frac{(p_{10}^u + p_{21}^u)(2U - N - R - Cp_{03}^r) - p_{10}^u p_{21}^u (N - R - Cp_{03}^r) + p_{01}^r p_{10}^u (N - U)}{p_{10}^u (N - R - Cp_{03}^r) + p_{01}^r (N - U)}. \tag{A.2}$$

When $p_{20}^u > \rho_2^*(\mathbf{A}_2, \mathbf{A}_3)$, *using a refurbished device is optimal in state 2; when* $p_{20}^u < \rho_2^*(\mathbf{A}_2, \mathbf{A}_3)$, *using a new device is optimal in state 2.*

A.2. COMDP: monotonicity conditions

In this context, one condition is sufficient to ensure the monotonicity of the optimal policy: as the condition of a used device gets worse, the likelihood of returning it to like-new condition is non-increasing. In other words, as long as the repair probability does not increase as the state gets worse, the optimal policy will call for increasingly effective (and costly) maintenance actions. The next proposition formalizes this intuitive result.

Proposition A.3 *If* $p_{i0}^u \geqslant p_{j0}^u$ *for each* $j > i$, *then the optimal policy will be monotone with respect to the process state.*

A.3. POMDP: existence of an optimal policy

Proposition A.4 *A bounded solution to Equation* (5) *exists, and a unique, stationary policy exists that minimizes the average expected cost. Furthermore, the average cost per unit time is constant*; *that is,* $g(\boldsymbol{\alpha}) = g^*$, *independent of the initial belief vector.*

The result follows direction from Theorem 4.2 of Fernández-Gaucherand *et al* (1991).

Appendix B. Data for example problems

Most of the data used for the example problems are reported in Section 4. This section reports all of the remaining data necessary to solve the problem variants discussed. Cost data are based on numbers from Klein (2005) and Landro (2008) unless otherwise noted.

B.1. Post-procedure observation probabilities

Table B1 reports the post-procedure observation probabilities, q_{jk}^r, used for in the example problems. By definition, a device cannot be in state 0 (perfect condition) after a procedure is performed, so corresponding observation probabilities are not relevant. In addition, the probability of observing state 0 after a procedure is performed is 0.

B.2. Post-reprocessing observation probabilities

For all example problems, a *Very High* information quality level refers to the following observation probabilities:

$$[q_{jk}^u] = \begin{bmatrix} 0.99 & 0.01 & 0 & 0 \\ 0.005 & 0.99 & 0.005 & 0 \\ 0 & 0.005 & 0.99 & 0.005 \\ 0 & 0 & 0.01 & 0.99 \end{bmatrix},$$

where j (row) is the observed state and k (column) is the actual state.

B.3. Parameters for example 1: angioplasty balloon

Costs and rewards: $N = 515$, $U = 250$, $C = 1\,000\,000$, $R = -16\,000$.

Transition probabilities:

$$[p_{ij}^r] = \begin{bmatrix} 0 & 0.4999995 & 0.4999995 & 1e-6 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix},$$

$$[p_{ij}^u] = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.98 & 0.02 & 0 & 0 \\ 0.5 & 0.3 & 0.2 & 0 \\ 0.05 & 0.5 & 0.3 & 0.15 \end{bmatrix}.$$

Table B1 Information matrices for example problems

Information quality level	*Observation probability matrix*
Very high	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.99 & 0.01 & 0 \\ 0 & 0.005 & 0.99 & 0.005 \\ 0 & 0 & 0.01 & 0.99 \end{bmatrix}$
High	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.9 & 0.1 & 0 \\ 0 & 0.05 & 0.9 & 0.05 \\ 0 & 0 & 0.1 & 0.9 \end{bmatrix}$
Medium	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.55 & 0.45 & 0 \\ 0 & 0.225 & 0.55 & 0.225 \\ 0 & 0 & 0.45 & 0.55 \end{bmatrix}$
Low	$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.333 & 0.333 & 0.334 \\ 0 & 0.333 & 0.333 & 0.334 \\ 0 & 0.333 & 0.333 & 0.334 \end{bmatrix}$

B.4. Parameters for example 2: pulse oximetry sensor

Costs and rewards: $N=10$, $U=6$, $C=10\,000$, $R=-15$.
Transition probabilities:

$$[p^r_{ij}]=\begin{bmatrix} 0 & 0.6 & 0.399999999 & 1e-9 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0.25 & 0.75 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

$$[p^u_{ij}]=\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.99 & 0.01 & 0 & 0 \\ 0.8565 & 0.0935 & 0.05 & 0 \\ 0.05 & 0.75 & 0.15 & 0.05 \end{bmatrix}.$$

B.5. Parameters for example 3: orthopaedic blade

Costs and rewards: $N=30$, $U=15$, $C=100\,000$, $R=-650$.
Transition probabilities:

$$[p^r_{ij}]=\begin{bmatrix} 0 & 0.499999995 & 0.499999995 & 1e-8 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

$$[p^u_{ij}]=\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.98 & 0.02 & 0 & 0 \\ 0.75 & 0.2 & 0.05 & 0 \\ 0.05 & 0.75 & 0.15 & 0.05 \end{bmatrix}.$$

B.6. Parameters for example 4: laparoscopic surgery

Costs and rewards: $N=1200$, $U=250$, $C=1\,000\,000$, $R=-8000$.
Transition probabilities:

$$[p^r_{ij}]=\begin{bmatrix} 0 & 0.6 & 0.3999999 & 1e-7 \\ 0 & 0 & 0.6 & 0.4 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

$$[p^u_{ij}]=\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0.8609 & 0.1391 & 0 & 0 \\ 0.8354 & 0.0823 & 0.0823 & 0 \\ 0.05 & 0.75 & 0.1 & 0.1 \end{bmatrix}.$$

Received October 2007;
accepted September 2008 after two revisions