Data Uncertainty in Markov Chains: Application to Cost-effectiveness Analyses of Medical Innovations

Joel Goh[1], Mohsen Bayati[1], Stefanos A. Zenios[1], Sundeep Singh[2], David Moore[3]

[1]Stanford Graduate School of Business, CA 94305

[2]Stanford University Division of Gastroenterology, CA 94305

[3]Stanford Clinical Excellence Research Center, CA 94305

Abstract

Cost-effectiveness studies of medical innovations often suffer from data inadequacy. When Markov chains are used as a modeling framework for such studies, this data inadequacy can manifest itself as imprecise estimates for many elements of the transition matrix. In this paper, we study how to compute maximal and minimal values for the discounted value of the chain (with respect to a vector of state-wise costs or rewards) as these uncertain transition parameters jointly vary within a given uncertainty set. We show that these problems are computationally tractable if the uncertainty set has a row-wise structure. Conversely, we prove that if the row-wise structure is relaxed slightly, the problems become computationally intractable (NP-hard). We apply our model to assess the cost-effectiveness of fecal immunochemical testing (FIT), a new screening method for colorectal cancer. Our results show that despite the large uncertainty in FIT's performance, it is highly cost-effective relative to the prevailing screening method of colonoscopy.

1 Introduction

Cost-effectiveness analyses are essential to the process of innovation. They weigh the proposed benefits of an innovation against the costs of implementing or developing the innovation, and are often necessary to obtain buy-in from decision makers to adopt the innovation or to invest resources to conduct further studies. Models based on Markov chains (henceforth, *Markov models*) have been extensively employed for such analyses, especially in the medical domain (see, e.g., Sonnenberg and Beck 1993). For example, they have been used in evaluating screening methodologies for breast cancer (Duffy et al. 1995, Hillner and Smith 1991), colorectal cancer (Ladabaum et al. 2001, 2004), prostate cancer (Ross et al. 2000), and HIV (Brandeau et al. 1992). They have also been used to evaluate the cost-effectiveness of vaccination programs (Hutton et al. 2007, 2010) as well as medical devices and drugs (Jönsson et al. 1999, Kobelt et al. 2003, Sanders et al. 2005).

Data inadequacy can limit the accuracy of cost-effectiveness studies. Innovations, by their very nature, represent departures from established practice or operations, and therefore often lack extensive evidence bases for their cost-effectiveness, especially during their inception. This data inadequacy is exacerbated for *medical* innovations (e.g., a new medical device, drug, or a new method of healthcare delivery), which usually require costly clinical trials on human subjects as to definitively establish their efficacy. Therefore, decision-makers who make preliminary assessments of the efficacy and costs of a medical innovation often have to base their analyses on evidence that is scarce and indirect (e.g., different sample populations, different contexts).

When Markov models are used for cost-effectiveness studies of innovations, this data inadequacy can manifest itself as uncertainty in the transition matrix parameters of the Markov chain. Therefore, decision-makers often have to grapple with the following question: *How does uncertainty in the model's input transition parameters translate into uncertainty in the model's output metrics of interest*? In the applied literature, Markov models are usually analyzed using Monte Carlo simulation, and a common approach is to perform one-way or two-way sensitivity analyses over a few parameters (e.g., Briggs et al. 1994, Eddy 1990). Other authors (e.g., Ladabaum et al. 2001, Heitman et al. 2010) perform randomized multi-way sensitivity analyses by making distributional assumptions about the uncertain parameters (in particular, they almost always assume that these parameters are independent), and randomly draw the parameters from these distributions. However, both types of sensitivity analyses are *ad hoc* in the sense that they do not exhaustively explore the parameter space. Consequently, they under-estimate the true aggregate effect of the model uncertainty, and leave analysts with an incomplete picture of the effect of model uncertainty.

1.1 Contributions

In this paper, we study the problem of evaluating the best-case and worst-case performance of a Markov model with respect to the *aggregate uncertainty* in its transition matrix. We employ the modeling paradigm of robust optimization (Ben-Tal and Nemirovski 1998, Bertsimas and Sim 2004) where uncertain parameters are not described by probability distributions but instead, modeled as residing within an *uncertainty set*. We allow this uncertainty set to be constructed using multiple data sources, which is a common phenomenon in medical applications. We then use mathematical optimization to find values of the uncertain parameters that yield the best-case and worst-case performance of the model (see Bandi and Bertsimas (2012) for a recent overview of applying this approach to various problems, as well as an in-depth discussion of how typical uncertainty sets can

be constructed). Even though we adopt a non-probabilistic model of uncertainty, we demonstrate in §5.4 how our approach can be adapted to derive probabilistic conclusions of the type: Innovation A is more cost-effective than the status quo B with $p_{\text{value}} < 0.05$.

Our primary analysis assumes that the uncertainty set has a *row-wise structure*, which we formally define later. The row-wise structure is flexible enough to encompass a variety of practical modeling applications, and is a necessary structural assumption: We formally prove that a mild relaxation of this assumption causes the problem to be NP-hard. However, with the row-wise assumption, we show that the problem can be solved efficiently by an iterative approach, allowing modeling of large state spaces. Our approach can also be modified to exploit special problem structures to obtain further efficiency improvements.

1.2 Application to Colorectal Cancer Screening

Our research was motivated by the specific problem of evaluating the cost-effectiveness of fecal immunochemical testing (FIT), which is a new screening method for colorectal cancer (CRC). CRC is the third-largest cause of cancer-related deaths for men and women in the U.S. (U.S. Cancer Statistics Working Group 2013), but has been shown to be largely preventable by effective screening (e.g., Eddy 1990, Frazier et al. 2000). The two main screening modalities presently in use are endoscopic methods (e.g., colonoscopy, flexible sigmoidoscopy) and stool-based testing (e.g., FIT, guaiac-based fecal occult blood tests, fecal DNA tests). Presently, colonoscopy is the most prevalent screening modality by a large margin (Joseph et al. 2012).

The majority of papers that study the cost-effectiveness of various CRC screening modalities use Markov models (Eddy 1990, Ladabaum et al. 2001, 2004). However, a major challenge in applying Markov modeling to analyze newer screening methods, such as FIT, is that there is limited evidence for its test characteristics, which ultimately translates into substantial uncertainty in the transition matrix of the Markov chain, and a rigorous treatment of the aggregate effect of model uncertainty is necessary in order for practitioners and policymakers to have a complete picture of the cost-effectiveness of these newer methods relative to the status quo.

1.3 Organization of the Paper

After a review of related literature in §2, we introduce our model of a Markov chain with uncertain transition probabilities in §3, where we introduce the main problem studied in this paper, formally define the row-wise uncertainty structure, and provide a proof of computational intractability when

the row-wise structure is absent. We analyze the model in §4, and in §5, we apply our methods to perform a cost-effectiveness analysis of various CRC screening strategies that accounts for the aggregate effect of parameter uncertainty. Finally, §6 concludes. All proofs are provided in the Appendix.

2 Related Literature

Blanc and Den Hertog (2008) study a Markov chain with row-wise uncertainty, and compute bounds on its stationary distributions, n-step state probabilities, expected hitting times and absorption probabilities of a set of states. Their method uses a single convex optimization step for each bound. Although we share a common modeling framework, our present work is novel in three important respects. First, we use a complexity argument to motivate why a row-wise uncertainty structure is needed. Second, our solution approach is iterative, which is better-suited for problems with large state-spaces, which, as we will see in §5, is the case for our target application of CRC screening. Third, our row-wise uncertainty structure is more general, allowing for multiple data sources to contribute estimates for the same row of the transition matrix.

Several mathematical papers, such as those by Škulj (2007), Škulj and Hable (2009), De Cooman et al. (2009), Hermans and De Cooman (2012), investigate the theoretical aspects of uncertain Markov chains. Generally, these papers aim to derive conceptual analogs of elements of classical Markov chain theory for uncertain Markov chains (e.g., n-step distribution of states, limiting behavior, convergence rates). These papers focus on advancing the theory of uncertain Markov chains, while our present work focuses on developing computational methods that can be applied to bound the performance of an uncertain Markov chain. Compared to these papers, we focus more narrowly on a specific problem class, but delve deeper into its computational aspects and design algorithms to study the problem.

Our paper shares some similarities with the literature on robust Markov Decision Processes (MDP) (Iyengar 2005, Nilim and El Ghaoui 2005, Xu and Mannor 2012), but differs in its setup. In robust MDP models, a controller and an adversary (i.e., nature) alternately make decisions, and the goal is to find an operating policy for the controller that is robust (in a minimax sense) to the sequential decisions of the adversary (i.e., the worst-case outcomes by nature). The papers in this stream of research also assume row-wise uncertainty. In contrast, in our model, nature makes a single decision (at time 0) that fixes the uncertain parameters for the entire time horizon of the

problem, and our goal is to evaluate the best-case and worst-case value of the chain, which represent the most beneficent and most adversarial choices of nature respectively. In the case of row-wise uncertainty, we will show that our model can be interpreted as an MDP where a (single) dynamic controller is optimizing the value of the chain.

Notation

Let Δ_N represent the standard $N-1$ simplex, i.e., $\Delta_N := \left\{\mathbf{x} \in \mathbb{R}_+^N : \sum_{i=1}^N x_i = 1\right\}$. We denote matrices and vectors with a bold typeface, and by convention, all vectors are column vectors unless stated otherwise. Inequalities applied to vectors and matrices are interpreted as being applied component-wise. We denote matrix transposition with $'$. Also, define $\mathbf{e} := (1, \ldots, 1)' \in \mathbb{R}^N$ as the vector of ones, $\mathbf{e_k} = (0, \ldots, 1, \ldots, 0)'$ as the kth standard basis vector, and $\mathbf{I}$ as the identity matrix. For optimization problems that are infeasible, we adopt the convention that their optimal value is $+\infty(-\infty)$ if the problem is stated as a minimization (maximization) problem.

3 Model

Consider a *discrete time Markov chain* (DTMC) $\{X_t, t = 0, 1, 2, \ldots\}$ on a finite set of states, $\mathcal{X} := \{1, \ldots, N\}$, with transition matrix $\mathbf{P} \in \mathbb{R}^{N \times N}$. In particular, for all $i, j \in \mathcal{X}$, $\mathbf{P}_{ij}$ is the probability of transition from state i to state j. We associate a reward $\mathbf{r} \in \mathbb{R}^N$ to each state (negative rewards are interpreted as costs), and let $\boldsymbol{\pi_0} \in \Delta_N$ represent the (known) initial distribution over the states. Future periods are discounted at a rate of $\alpha \in (0, 1)$, which is assumed to be rational for a technical reason (see Remark 2 in §3.2).

The expected net present value (NPV) of the chain over an infinite horizon, expressed as a function of its transition matrix, $\mathbf{P}$, may be written as

$$f(\mathbf{P}) := \boldsymbol{\pi_0}' \sum_{t=0}^{\infty} \alpha^t \mathbf{P}^t \mathbf{r} = \boldsymbol{\pi_0}' \left(\mathbf{I} - \alpha \mathbf{P}\right)^{-1} \mathbf{r}. \tag{1}$$

We assume that the true transition matrix $\mathbf{P}$ is unknown to the modeler (i.e., is *uncertain*) , but instead, the modeler only knows that $\mathbf{P}$ resides within a certain *uncertainty set*, $\mathcal{P}$, which we assume throughout this paper to be convex and compact.

This uncertainty set captures the imprecision in estimates of the transition matrix. If confidence intervals for each entry of the transition matrix are available (at a certain confidence level), a simple and data-driven way of constructing $\mathcal{P}$ is to treat these confidence intervals as deterministic

quantities, and define $\mathcal{P}$ as the set of all matrices $\mathbf{P} \in \mathbb{R}^{N\times N}$ where each entry of $\mathbf{P}$ belongs to its prescribed interval and $\mathbf{P}$ satisfies row-stochastic constraint on its elements. The uncertainty set $\mathcal{P}$ constructed in this manner has the interpretation that with a certain probability, the true transition matrix lies within $\mathcal{P}$, regardless of how the original confidence intervals are jointly distributed. We describe more concrete examples of constructing $\mathcal{P}$ in the next subsection.

Since $\mathbf{P}$ is uncertain, the true NPV of the chain cannot be precisely known. However, we can find a pair of values, $\underline{f}$ and $\overline{f}$, that capture the range of its NPV as its true transition matrix varies across the uncertainty set $\mathcal{P}$. We will call the interval $[\underline{f}, \overline{f}]$ the *certainty interval* for the NPV of the chain. The lower end-point $\underline{f}$ is defined as the optimal value of the minimization problem,

$$\underline{f} := \min_{\mathbf{P}\in\mathcal{P}} f(\mathbf{P}), \tag{2}$$

and the upper end-point $\overline{f}$ is defined analogously as $\overline{f} := \max_{\mathbf{P}\in\mathcal{P}} f(\mathbf{P})$. Our analysis will focus on how to evaluate these end-points.

Before we conduct our analysis, we present two reductions to simplify our exposition. Both reductions are without loss of generality. First, we can limit our analysis to the minimization problem (2) since we can transform the maximization problem into a minimization as $\max_{\mathbf{P}\in\mathcal{P}} f(\mathbf{P}) = -\min_{\mathbf{P}\in\mathcal{P}} \boldsymbol{\pi_0}' \left(\sum_{t=0}^{\infty} \alpha^t \mathbf{P}^t\right)(-\mathbf{r})$. Second, we may further limit the problem to nonnegative $\mathbf{r}$. This is because defining $r_* := \min_{n=1}^{N} r_n$, and noting that $\left(\sum_{t=0}^{\infty} \alpha^t \mathbf{P}^t\right)\mathbf{e} = \sum_{t=0}^{\infty} \alpha^t \mathbf{e} = \frac{1}{1-\alpha}\mathbf{e}$, we have $f(\mathbf{P}) = \boldsymbol{\pi_0}' \left(\sum_{t=0}^{\infty} \alpha^t \mathbf{P}^t\right)(\mathbf{r} - r_*\mathbf{e}) + \frac{r_*}{1-\alpha}$. Since $\mathbf{r} - r_*\mathbf{e} \geq \mathbf{0}$, it is sufficient to consider $\mathbf{r} \geq \mathbf{0}$.

Note that our model assumes that the vector $\mathbf{r}$ is known precisely, even though uncertainty in $\mathbf{r}$ might be important in applications. In the (very natural) case that each element of $\mathbf{r}$ is described by an interval (i.e., the uncertainty set of $\mathbf{r}$ is a hyper-rectangle), the best and worst cases can be easily evaluated. This is because the NPV is increasing in each component of $\mathbf{r}$. Hence, the best and worst cases of the NPV can be obtained by first substituting the extreme values of $\mathbf{r}$, and then performing the optimization in (2). For the remainder of this paper, we will treat $\mathbf{r}$ as a fixed parameter.

3.1 Structure of Uncertainty Sets

For the main analysis in this paper (§4), we assume that $\mathcal{P}$ has a *row-wise* structure. The row-wise structure is formally defined in Assumption 1 below.

Assumption 1 *Assume that for some $K \in \mathbb{N}$, there exist matrices $\mathbf{U} := [\mathbf{u_1}, \ldots, \mathbf{u_K}], \mathbf{Z} := [\mathbf{z_1}, \ldots, \mathbf{z_K}] \in \mathbb{R}^{N\times K}$, a set $\mathcal{Z} := \mathcal{Z}_1 \times \ldots \times \mathcal{Z}_K$, where $\mathcal{Z}_k \subseteq \Delta_N$, $k = 1, \ldots, K$ are tractable*

convex sets (i.e., they have explicit semidefinite representations) such that

1. $\mathbf{z_k} \in \mathcal{Z}_k$, $k = 1, \ldots, K$,
2. $\mathbf{u_k} \geq \mathbf{0}$, $k = 1, \ldots, K$,
3. $\sum_{k=1}^{K} \mathbf{u_k} = \mathbf{e}$, *and*
4. $\mathbf{P} \in \mathcal{P}$ *iff* $\mathbf{P} = \mathbf{UZ}' = \sum_{k=1}^{K} \mathbf{u_k}\mathbf{z_k}'$.

An intuitive interpretation of Assumption 1 is that there are K uncertain vectors $\mathbf{z_k}, k = 1, \ldots, K$, which drive the uncertainty in $\mathcal{P}$. Each transition matrix $\mathbf{P} \in \mathcal{P}$ is constructed as the product $\mathbf{P} = \mathbf{UZ}'$, which means that each row of $\mathbf{P}$ is some convex combination of the K uncertain vectors, with the weights for the convex combinations specified by the entries of matrix $\mathbf{U}$. Note that Assumption 1 encompasses, but is more general than, the row-wise assumption used in the existing literature of uncertain DTMCs (Blanc and Den Hertog 2008) and uncertain MDPs (Iyengar 2005, Nilim and El Ghaoui 2005, Xu and Mannor 2012).

More importantly, Assumption 1 is sufficiently general to many practically-relevant situations. The technical requirement for $\mathcal{Z}_k$ to be "computationally tractable" is a mild one that is standard in robust optimization (see, e.g., Ben-Tal et al. 2009), and furthermore encompasses most of the sets one would use to model uncertainties in practice (e.g, polytopes, ellipsoids). Below, we describe several example constructions of convex sets that fit Assumption 1.

Example 1 (Element-wise Interval Uncertainty). For each $k = 1, \ldots, K$, construct $\mathcal{Z}_k$ by a cartesian product of uncertainty intervals. Specifically, for some a_{ik}, b_{ik}, $i = 1, \ldots, N$, such that $0 \leq a_{ik} \leq b_{ik} \leq 1$, construct $\mathcal{Z}_k = ([a_{1k}b_{1k}] \times \ldots \times [a_{Nk}, b_{Nk}]) \cap \Delta_N$. In this case, $\mathcal{Z}_k$ is a polytope. This setting represents the case that $\mathcal{Z}_k$ represents an interval of uncertainty on each element of the unknown transition matrix $\mathbf{P}$. In practice, the quantities a_{ik}, b_{ik} can be obtained from the confidence intervals of the transition probabilities (at a certain confidence level), which are estimated from data.

Example 2: (Weighted Distance from an Empirical Distribution). For each $k = 1, \ldots, K$, let $\mathbf{z_k}^\circ \in \mathcal{Z}_k$ represent a point estimate for the unknown transition probability vector $\mathbf{z_k}$. Then, for some rational $q > 1$, real $\epsilon > 0$, and weight vector $\mathbf{w} \in \mathbb{R}_+^N$, construct $\mathcal{Z}_k = \left\{\mathbf{z} \in \Delta_N : \sum_{i=1}^{N} w_i \left|z_i - z_{ki}^\circ\right|^q \leq \epsilon\right\}$. We note that $\mathbf{Z}_k$ is tractable because it can expressed in terms of a finite number of second-order conic constraints using the methods outlined in Alizadeh and Goldfarb (2003). This construction captures transition probabilities that are a certain weighted distance (in the q-norm sense) away from the point estimate. Two special cases of this construction

are sets constructed by bounding the (a) χ *divergence of order* q or (b) *modified* χ^2 *distance*. We refer readers to Ben-Tal et al. (2013) for details.

Example 3: (Multiple Data Sources). Suppose that, for each row of $\mathbf{P}$, there are $m > 1$ distinct data sources that each contribute a noisy estimate of that row (i.e., there are m uncertainty sets for that row of $\mathbf{P}$, which can, e.g., be constructed as in Examples 1 or 2). If we wanted $\mathbf{P}$ to be an equally-weighted estimate of these m data sources, we could choose $K = mN$, and $\mathbf{u_k} = (1/m)\mathbf{e_k}$ for $k = j(N-1)+1, \ldots, jN$ and $j = 1, \ldots, m$. Using a similar construction, we can also model the case where each row of $\mathbf{P}$ is estimated by a different number of data sources, and the case where the weights need not be equal.

Example 4: (Positive Correlation Between Rows). We can model positive correlations between rows of $\mathbf{P}$, using common factors $\mathbf{z_k}$ to construct distinct rows of $\mathbf{P}$. This is achieved by letting some $\mathbf{u_k}$ have multiple strictly positive entries. For example, to model that rows 1 and 2 of $\mathbf{P}$ are perfectly positively correlated, we can choose $\mathbf{u_k} = \mathbf{e_1} + \mathbf{e_2}$ for some k.

3.2 Complexity Analysis of (2) for General Uncertainty Sets

The following analysis further motivates why the row-wise structure of Assumption 1 is necessary. We show that if this assumption is even slightly relaxed, solving (2) becomes intractable. In particular, a natural relaxation of Assumption 1 is to simply assume that $\mathcal{P}$ is itself a tractable convex set. For example, this would allow $\mathcal{P}$ to be specified by constraints that involve elements from distinct rows of $\mathbf{P}$. However, as we proceed to show, it is generally NP-hard to calculate the value of $\underline{f}$ under this setting, even if $\mathcal{P}$ is as simple as a polytope.

Our proof strategy is as follows. Given a directed graph, we construct a special case of problem (2) that violates Assumption 1 using the graph structure as an input. We will then show that we can determine whether this graph has a Hamiltonian cycle[1] by checking the optimal value of that problem. The former is known to be an NP-complete problem (Karp 1972). The special case is constructed as follows. Let $G(V, E)$ be a given directed graph with vertex set $V = \{1, \ldots, N\}$ and edge set E with no self loops. Consider the problem

$$\begin{aligned} \underline{f} = \min_{\mathbf{P} \in \mathbb{R}^{N \times N}} \quad & \boldsymbol{\pi_0}'(\mathbf{I} - \alpha \mathbf{P})^{-1}\mathbf{r} \\ \text{s.t.} \quad \mathbf{Pe} &= \mathbf{e},\ \mathbf{P}'\mathbf{e} = \mathbf{e},\ \mathbf{P} \geq \mathbf{0}, \\ P_{ij} &= 0 \qquad\qquad (i,j) \notin E. \end{aligned} \tag{3}$$

[1] A Hamiltonian cycle is a directed cycle that visits each vertex of a graph exactly once.

where $\boldsymbol{\pi_0} = \mathbf{r} = [1, 0, \ldots, 0]$. Note that (3) violates Assumption 1 because of the column stochastic constraint on $\mathbf{P}$. Problem (3) can be interpreted as finding the transition probabilities for a random walk on G to minimize the expected discounted cost, if a unit cost is incurred in state 1 and nothing elsewhere, with the additional requirement that the transition matrix be doubly-stochastic. We now can state our complexity result in the following proposition.

Proposition 1 *G contains a Hamiltonian cycle iff $\underline{f} = \frac{1}{1-\alpha^N}$ in (3). Therefore, problem (2) without Assumption 1 is NP-hard.*

Remark 1 *One corollary of Proposition 1 is that G contains a Hamiltonian cycle if and only if there is a doubly-stochastic matrix $\mathbf{P}$, which is feasible in (3), and represents a DTMC whose first return time to state 1 is equal to N with probability 1. This particular corollary has been shown before by Borkar et al. (2009), but their proof is quite complicated: It focuses on Markov chains with a single ergodic class (i.e., unichains) and applies perturbations (on doubly-stochastic matrices) in order to carry the result through. In contrast, our proof is much simpler, only relying on elementary probabilistic arguments, and may be of independent interest.*

Remark 2 *If α is irrational, the proof of Proposition 1 goes through, but it is not meaningful. This is because testing whether $\underline{f}$ is equal to $\frac{1}{1-\alpha^N}$ cannot be done under the standard model of computation when α is irrational. This technical issue was why we assumed that α is rational.*

4 Analysis

In this section, we proceed to analyze the optimization problem (2) when the uncertainty set $\mathcal{P}$ has a row-wise structure in the form of Assumption 1. In §4.1 we prove an equivalent reformulation of our problem as a semi-infinite optimization problem and will interpret it as a discounted MDP. We follow this with a description of iterative methods for the problem in §4.2, and present a particularly efficient method in the special case that the uncertainty set $\mathcal{P}$ is constructed by intervals.

4.1 Formulation as a Semi-infinite Optimization Problem

Our approach to solving (2) begins by considering the following semi-infinite optimization problem over the decision variable $\mathbf{v} \in \mathbb{R}^N$:

$$\begin{aligned} f_* := \max_{\mathbf{v} \in \mathbb{R}^N} \quad & \boldsymbol{\pi_0}' \mathbf{v} \\ \text{s.t.} \quad & v_i \leq r_i + \alpha\, \mathbf{e_i}' \mathbf{P} \mathbf{v} \qquad \forall \mathbf{P} \in \mathcal{P}, i = 1, \ldots, N. \end{aligned} \tag{4}$$

In (4), for *each* of the $i = 1, \ldots, N$ constraints, the inequality $v_i \leq r_i + \alpha \mathbf{e_i}'\mathbf{P}\mathbf{v}$ must hold for all $\mathbf{P} \in \mathcal{P}$. Problem (4) is labelled as "semi-infinite" because it has an infinite number of constraints as $\mathcal{P}$ is not a finite set.

Problem (4) has several appealing properties. First, it is intimately connected to our original problem of interest (2). It is equivalent to (2) when the row-wise assumption holds, and bounds the value of (2) otherwise.

Theorem 1 *If $\mathcal{P}$ satisfies Assumption 1, then $f_* = \underline{f}$. For general uncertainty sets $\mathcal{P}$, $f_* \leq \underline{f}$.*

The second appealing property of problem (4) is that it is usually computationally tractable. It can be re-formulated into an equivalent convex optimization problem by well-established techniques from robust optimization (Ben-Tal and Nemirovski 1998, Bertsimas and Sim 2004), or can even be modeled directly using various algebraic modeling toolboxes (Löfberg 2004, Goh and Sim 2011). Therefore, this approach provides a direct method for solving (2) when $\mathcal{P}$ has a row-wise structure. However, when the state space is very large, it faces a computational limitation similar to Blanc and Den Hertog's (2008) approach: The equivalent convex optimization problem has a size that is at least on the order of $\mathcal{O}(KN)$, and may be too large to be practically solvable.

Third, problem (4) has a natural interpretation: It represents a discounted MDP on the finite state space $\{1, \ldots, N\}$, but with a continuous (but compact) action space $\mathcal{Z}$ (see, e.g., Manne 1960). The "action" in each state may be interpreted as choosing the transition probabilities into the next state. In the case that $\mathcal{P}$ satisfies Assumption 1, the choice of actions in each state is the collection $(\mathbf{z_1}, \ldots, \mathbf{z_K})$, and the transition probability of going from state i to j is given by $\sum_{k=1}^{K} u_{ki} z_{kj}$. The Bellman equation for this MDP is

$$v_i = \min_{\mathbf{z_k} \in \mathcal{Z}_k, k=1,\ldots,K} \left\{ r_i + \alpha \sum_{k=1}^{K} u_{ki} \mathbf{z_k}' \mathbf{v} \right\}, \tag{5}$$

where $\mathbf{v}$ represents the value function. The quantity f_* can be evaluated as $f_* = \boldsymbol{\pi_0}'\mathbf{v}$. Therefore, Theorem 1 implies that one strategy to evaluate $\underline{f}$ under Assumption 1 is to work with the MDP (5), which is what we will pursue below.

4.2 Policy Iteration

We now describe our policy iteration approach to solve the MDP (5). A parallel approach based on value iteration is provided in our full paper. Policy iteration generates a sequence of solutions

$(\mathbf{z_1}^{(0)}, \ldots, \mathbf{z_K}^{(0)}), (\mathbf{z_1}^{(1)}, \ldots, \mathbf{z_K}^{(1)}), \ldots$ that successively improve the value function until no further improvements can be made. These updates occur in two steps:

$$\begin{aligned} \text{Policy Improvement Step:} \quad \mathbf{z_k^{(t)}} &\in \underset{\mathbf{z_k} \in \mathcal{Z}_k}{\arg\min} \left\{ \mathbf{z_k}' \mathbf{v}^{(t)} \right\} \qquad k = 1, \ldots, K \\ \text{Policy Evaluation Step:} \quad \mathbf{v}^{(t+1)} &= \left(\mathbf{I} - \alpha \sum_{k=1}^{K} \mathbf{u_k} \mathbf{z_k^{(t)}}' \right)^{-1} \mathbf{r}. \end{aligned} \tag{6}$$

Each update requires solving K optimization problems separately, each of which only has $\mathcal{O}(N)$ decision variables. Therefore, when K and N are large, these iterative approaches are computationally appealing.

In the simple, but practically important, special case that all the sets $\mathcal{Z}_k$ have interval uncertainty structures (as described in §3.1) the optimizations over $\mathcal{Z}_k$ in each update step can be computed very efficiently by exploiting this special structure. Specifically, for every k, there exists an optimal $\mathbf{z_k^*}$ that has a "threshold" structure, which is described in Proposition 2 below.

Proposition 2 *For each $k = 1, \ldots, K$, suppose that $\mathcal{Z}_k$ has an interval uncertainty structure, that is, $\mathcal{Z}_k = \{\mathbf{z} \in \Delta_N : \underline{\mathbf{z}}_\mathbf{k} \leq \mathbf{z} \leq \overline{\mathbf{z}}_\mathbf{k}\}$ for parameters $\underline{\mathbf{z}}_\mathbf{k}, \overline{\mathbf{z}}_\mathbf{k}$. For a fixed t, let $\{\sigma(1), \ldots, \sigma(N)\}$ be a permutation of $\{1, \ldots, N\}$ such that $v^{(t)}_{\sigma(1)} \leq v^{(t)}_{\sigma(2)} \leq \ldots \leq v^{(t)}_{\sigma(N)}$ and for each k, define the indices $i_k := \inf \left\{ i \in \{1, \ldots, N\} : \sum_{j=1}^{i} \overline{z}_{k,\sigma(j)} + \sum_{j=i+1}^{N} \underline{z}_{k,\sigma(j)} \geq 1 \right\}$. Then, the vectors $\mathbf{z_k^{(t)}}$, $k = 1, \ldots, K$ constructed as follows,*

$$z^{(t)}_{k,\sigma(i)} = \begin{cases} \overline{z}_{k,\sigma(i)} & \text{if } i < i_k, \\ \underline{z}_{k,\sigma(i)} & \text{if } i > i_k, \\ 1 - \sum_{i=1}^{i_k - 1} \overline{z}_{k,\sigma(i)} - \sum_{i=i_k+1}^{N} \underline{z}_{k,\sigma(i)} & \text{if } i = i_k. \end{cases} \tag{7}$$

satisfy $\mathbf{z_k^{(t)}} \in \arg\min_{\mathbf{z_k} \in \mathcal{Z}_k} \left\{ \mathbf{z_k}' \mathbf{v}^{(t)} \right\}$.

The policy iteration algorithm for this special case of interval uncertainty is described in Algorithm 1. We note that each update step of the policy iteration requires a single sort of the current value function $\mathbf{v}^{(t)}$ (requiring $\mathcal{O}(N \log N)$ operations), which is followed by K evaluations of the optimal $\mathbf{z_k^{(t)}}$. In practice, we expect many degenerate indices (i.e., indices i such that $\overline{z}_{ki} = \underline{z}_{ki}$), and if we let N_s represent the maximum number of nondegenerate indices of $\mathbf{z_k}$ for every k, evaluation of $\mathbf{z_k^{(t)}}$ takes at most $\mathcal{O}(N_s)$ operations. Therefore, the total number of operations taken for each update is on the order of $\mathcal{O}(N \log N + K N_s)$.

Algorithm 1 Policy Iteration Method for $\mathcal{Z}_k$ with Interval Uncertainty Structure

Initialize $t \leftarrow 0$ and any feasible $\mathbf{z_k}, k = 1, \ldots, K$
Evaluate $\mathbf{v}^{(0)} = \left(\mathbf{I} - \alpha \sum_{k=1}^{K} \mathbf{u_k z_k}'\right)^{-1} \mathbf{r}$.
repeat
    $\widetilde{\mathbf{v}}^{(t)} = \text{SORT}(\mathbf{v}^{(t)})$.
    for $k = 1$ to K **do**
        Construct $\mathbf{z_k^{(t)}}$ using (7).
    end for
    Evaluate $\mathbf{v}^{(t+1)} = \left(\mathbf{I} - \alpha \sum_{k=1}^{K} \mathbf{u_k z_k^{(t)}}'\right)^{-1} \mathbf{r}$.
    Update $t \leftarrow t + 1$.
until $\left\|\mathbf{v}^{(t)} - \mathbf{v}^{(t-1)}\right\| = 0$
return $f_* = \boldsymbol{\pi_0}' \mathbf{v}^{(t)}$.

4.3 Generalizations

Even though our method was designed to compute a certainty interval around the discounted NPV of an uncertain DTMC, the same method can be used (with small modifications) to obtain certainty intervals for other quantities of interest for the DTMC such as the expected hitting time and the hitting probability to a subset of states. Details are omitted for brevity.

We can also apply our method to calculate certainty intervals around cost-effectiveness ratios. An important example of this ratio in health applications is the total cost per expected life year. Mathematically, minimizing this ratio is can be represented by

$$\min_{\mathbf{P} \in \mathcal{P}} \frac{\boldsymbol{\pi_0}' \left(\mathbf{I} - \alpha \mathbf{P}\right)^{-1} \mathbf{r_1}}{\boldsymbol{\pi_0}' \left(\mathbf{I} - \alpha \mathbf{P}\right)^{-1} \mathbf{r_2}}, \tag{8}$$

for appropriate choices of $\mathbf{r_1}, \mathbf{r_2} \geq \mathbf{0}$ (an analogous formulation holds for maximization). While this does not fall within the scope of our model *per se*, our model can be adapted to perform this computation. We note that problem (8) is equivalent to finding a minimal $\beta^* \geq 0$ such that the set $\left\{\mathbf{P} \in \mathcal{P} : \boldsymbol{\pi_0}' \left(\mathbf{I} - \alpha \mathbf{P}\right)^{-1} \left(\mathbf{r_1} - \beta^* \mathbf{r_2}\right) \leq 0\right\}$ is non-empty. To find β^*, consider the following optimization problem for a given *fixed* β:

$$R_\beta := \min_{\mathbf{P} \in \mathcal{P}} \boldsymbol{\pi_0}' \left(\mathbf{I} - \alpha \mathbf{P}\right)^{-1} \left(\mathbf{r_1} - \beta \mathbf{r_2}\right), \tag{9}$$

which is within the scope of our model. It is easy to see that R_β decreases in β, and that $R_{\beta^*} = 0$. Therefore, we can estimate β^* using bisection by iteratively solving (9).

5 Application: Cost Effectiveness of FIT for CRC Screening

We applied of our model to analyze the cost-effectiveness of FIT, a recently-developed screening modality for CRC. There are existing medical cost-effectiveness studies that use Markov models to study the cost-effectiveness of FIT (e.g., Parekh et al. 2008, Telford et al. 2010, Heitman et al. 2010), but none of these studies account for the aggregate effect of model uncertainty as we do.

For this study, we compare the cost-effectiveness of three screening modalities: No screening (NOSCREEN), annual FIT screening (FIT), and screening with colonoscopy at 10-year intervals (COLO). The Markov model used here is based on the colonoscopy screening model from the medical literature by Parekh et al. (2008), which has been well-tested and validated.

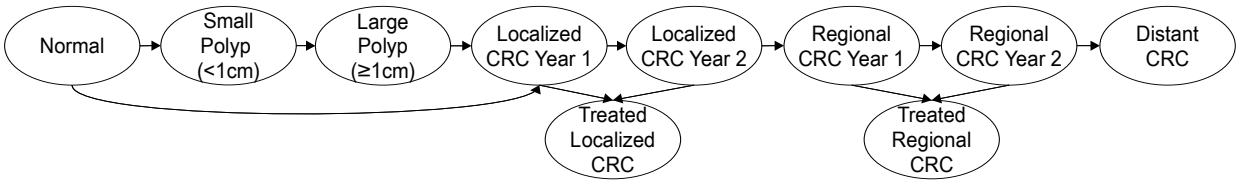

Figure 1: Health states in colonoscopy screening model. Arrows represent non-zero transition probabilities. Omitted from figure: Dead state and all self-transitions.

A period in the model represents a year, and at time 0, all subjects are assumed to be at age 50, and were followed to a maximum age of 100. In each period, subjects may reside in one of 11 principal health states depicted in Figure 1, which also illustrates the possible transitions between states. Age-specific non-CRC death probabilities were obtained from U.S. life tables (Arias 2011). FIT test characteristics were obtained from studies listed in Table 1. All other transition probability data, as well as the initial distribution over health states were obtained from Parekh et al. (2008). We used interval uncertainty sets for the uncertain parameters that were constructed as the reported 95% confidence intervals, and applied equal weights when there were multiple sources of estimates (see Examples 1 and 3 of §3.1). These confidence intervals, along with all other clinical and cost parameters for the model are listed in Table 1.

We superimposed the various screening modalities on top of the health states. In NOSCREEN, we assumed that CRC in each stage is detected upon symptomatic presentation (we used rates of symptomatic presentation reported by Parekh et al. (2008)). In FIT, testing is conducted every year, and in COLO, colonoscopy is conducted every 10 years. In all screening strategies,

- If polyps are detected, they are removed by colonoscopy,

Parameter Description	Values	References
Clinical		
Annual transition rate from Normal to Small Polyp (%)	1.1 – 1.9[a]	Parekh et al. (2008)
Annual transition rate from Small Polyp to Large Polyp (%)	1.5	Parekh et al. (2008)
Annual transition rate from Normal to Localized CRC (%)	0.006 – 0.086[a]	Parekh et al. (2008)
Annual transition rate from Large Polyp to Localized CRC (%)	5	Parekh et al. (2008)
Symptomatic presentation of Localized CRC (%)	22/year	Parekh et al. (2008)
Symptomatic presentation of Regional CRC (%)	40/year	Parekh et al. (2008)
Mortality rate from treated localized cancer (%)	1.74/year	Parekh et al. (2008)
Mortality rate from treated regional cancer (%)	8.6/year	Parekh et al. (2008)
Mean survival from distant cancer (years)	1.9	Parekh et al. (2008)
Colonoscopy sensitivity for cancer (%)	95 (90 – 97)	Parekh et al. (2008)
Colonoscopy sensitivity for large polyp (%)	90 (85 – 95)	Parekh et al. (2008)
Colonoscopy sensitivity for small polyp (%)	85 (80 – 90)	Parekh et al. (2008)
FIT sensitivity for cancer (%)	69 (51 — 86)	Allison et al. (1996)
	83 (62 — 100)	Greenberg et al. (2000)
	88 (71 — 100)	Greenberg et al. (2000)
	89 (51 — 99)	Wong et al. (2003)
	66 (55 — 76)	Morikawa et al. (2005)
FIT sensitivity for large polyp (%)	67 (57 — 76)	Allison et al. (1996)
	37 (21 — 53)	Greenberg et al. (2000)
	36 (21 — 51)	Greenberg et al. (2000)
	20 (17 — 23)	Morikawa et al. (2005)
FIT specificity (%)	95 (95 — 96)	Allison et al. (1996)
	98 (97 — 100)	Rozen et al. (1997)
	97 (95 — 100)	Rozen et al. (1997)
	90 (87 — 93)	Greenberg et al. (2000)
	88 (85 — 91)	Greenberg et al. (2000)
	99 (98 — 99)	Rozen et al. (2000)
Cost Parameters ($)		
Colonoscopy	920	Parekh et al. (2008)
FIT testing	22	Parekh et al. (2008)
Cost of Localized CRC Care	51 000	Parekh et al. (2008)
Cost of Regional CRC Care	98 000	Parekh et al. (2008)
Cost of Distant CRC Care	200 000	Parekh et al. (2008)

[a] Age-specific.

Table 1: Table of parameters for Markov modeling of colonoscopy screening. Uncertain parameters are listed as a nominal value with an accompanying uncertainty range provided in parentheses.

- If CRC is detected, colonoscopy is performed to confirm its presence and the patient enters treatment,
- If CRC is undetected, Localized (Regional) CRC progresses to Regional (Distant) CRC deterministically in 2 years,
- If polyps or CRC is detected, the patient enters surveillance, and colonoscopy is performed every 5 years,
- All screening is done up to age 80. After age 80, only colonoscopy is performed to evaluate symptoms.

These correspond to clinical best practices guidelines (e.g., Winawer et al. 2003). Using this model logic and parameter values from Table 1, we built a nominal transition matrix for the model, and constructed an uncertainty set around it. We applied a discount rate of 3% per year on all computations, as is standard in the medical literature (see, e.g., Parekh et al. 2008).

It is important to note that while there are only a few health states, the overall system state is a *product state* (comprising health states, surveillance strategies, various time intervals, etc.) and is very high-dimensional (> 15000 states). Because of the large size of the state space, we could not use the single-step optimization method by Blanc and Den Hertog (2008) or the direct robust optimization method in §4.1 because we were not able to fit the resulting optimization problem(s) into memory (even on a high-end Linux server using efficient sparse matrix representations).

5.1 Life Expectancy and Total Expected Cost of Each CRC Screening Modality

We first present a simple application of our model to compute separate certainty intervals on each subject's life expectancy and expected total cost. For the latter, we did some preprocessing by annualizing the cost of cancer care into state-wise costs (the raw data are reported as one-time costs).

We used policy iteration (Algorithm 1) to compute the certainty intervals, and the results are reported in Table 2. Convergence was quick and took less than 10 iterates in each computation (i.e., $t < 10$ upon termination in Algorithm 1).

Screening Method	Life Expectancy (years)	Total Cost ($)
NOSCREEN	19.069 (19.069 – 19.069)	2,975 (2,975 – 2,975)
COLO	19.143 (19.140 – 19.145)	3,875 (3,833 – 3,926)
FIT	19.150 (19.142 – 19.153)	2,324 (2,043 – 2,621)

Table 2: Nominal values and certainty intervals (in parentheses) of life expectancy and expected costs for no screening (NOSCREEN), colonoscopy at 10-year intervals (COLO), and annual FIT screening (FIT), discounted at a rate of 3% per year.

Our results complement the findings of Parekh et al. (2008): Both COLO and FIT screening methodologies yield longer life expectancies than NOSCREEN, but COLO costs more on average. FIT has the lowest total cost of all the options. FIT's life expectancies are nominally higher than COLO, but because their certainty intervals overlap, our analysis does not conclude that FIT is

definitively better along this dimension.

The central managerial insight that we obtain from these results is that given the test characteristics and costs of FIT and colonoscopy from the medical literature, the higher frequency (and substantially lower cost) of FIT outweighs its lower per-screening accuracy, both in terms of life expectancy and total costs.

5.2 Performance of FIT with Correlations

A contributor to FIT's good performance relative to COLO is our implicit Markovian assumption that consecutive FIT tests are statistically independent. This is a standard assumption used in the literature (e.g., Heitman et al. 2010, Telford et al. 2010). Nevertheless, a natural question to ask is what happens if the independence assumption is relaxed, and consecutive tests using FIT are positively correlated in some way? Specifically, we are interested in a specific type of correlation where certain types of polyps/CRC cannot be detected even with repeated FIT, e.g., due to their position, size, or biochemical properties.

To investigate this question, we suppose for simplicity that there are only two types of polyps/CRCs: The first type (existing with proportion $1-p$) is detectable by repeated FIT, while the second type (with proportion p) is undetectable by repeated FIT. Our base model corresponds to the case that $p = 0$, whereas the case of $p = 1$ means that if the first FIT fails to detect a polyp/CRC, every subsequent FIT test will also fail. There is very limited evidence on what p might be in practice: We are only aware of a single study (Zorzi et al. 2007) that suggests $p \approx 0.74$ for large polyps (but not CRCs). We investigate how the metrics of life expectancy and total cost vary as p is varied. Results are plotted in Figures 2.

From Figure 2, we observe that the certainty intervals for life expectancy from FIT and COLO overlap for all values of p. On the other hand, the certainty intervals for total cost from FIT and COLO are disjoint, with FIT costing much less than COLO. The insight gained from the combined result of both figures is that even in the extreme case where FIT are perfectly correlated (i.e., if FIT fails to detect a true polyp/CRC once, it will fail henceforth), the life expectancy obtained by FIT is comparable to COLO, and delivers this life expectancy with a much lower average cost.

5.3 Maximal and Minimal Cost-Effectiveness

Previously, we evaluated the certainty intervals around each metric of interest, life expectancy and expected total cost, separately. Because there may be different transition matrices that at-

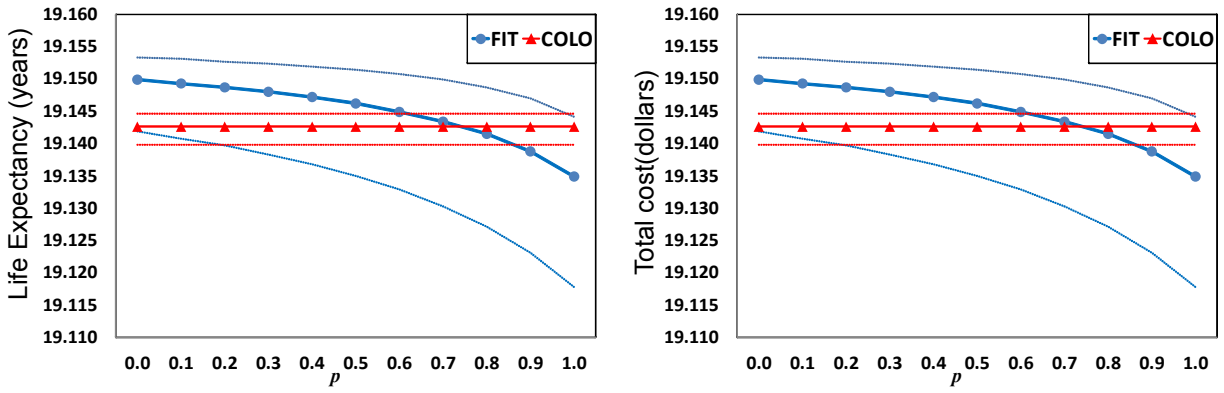

Figure 2: Nominal values and certainty intervals on life expectancy in years (left) and expected total cost in dollars (right) from FIT and COLO as a function of the proportion p of polyps/CRC that are undetectable by repeated FIT.

tain the optimal values for each metric, this is still not the ideal way to compare both screening modalities. A more rigorous comparison involves directly computing a certainty interval around the cost-effectiveness ratio (i.e., the total cost per expected year of life) of each modality, which we can do using the generalization described in §4.3.

The left panel of Figure 3 plots the total cost per expected year of life for FIT as we vary p, the fraction of polyps/CRCs undetectable by repeated FIT. In our base model that assumes no correlations (i.e., $p = 0$), COLO cost 202 (200 – 205) dollars per expected year of life, whereas FIT cost 122 (107 – 137) dollars per expected year of life.

This result sharpens our previous finding: For all values of p, FIT costs less per expected life-year gained than COLO. Even in the extreme case that $p = 1$, and even if we assume that the model uncertainty resulted in the worst possible performance for FIT, it costs about 172 dollars per expected year of life, which is approximately 85% of the most optimistic cost for COLO (200 dollars per expected year of life).

5.4 Probabilistic Uncertainty Sets and Hypothesis Testing

In our base model, we used the 95% confidence intervals of the uncertain model parameters as their uncertainty sets. We adopted this approach because of its simplicity and because it already provides a more robust and principled characterization of model uncertainty than typical sensitivity analyses used in cost-effectiveness studies. One limitation of this approach is that the computed certainty intervals do not have probabilistic interpretations.

We now present a simple, but conservative, method that will endow the certainty intervals with

a probabilistic interpretation. We do this by enlarging the intervals used to define our uncertainty set and repeat the calculations in the previous subsection. Specifically, we construct the uncertainty set using $1-\epsilon$ confidence intervals. Because there are no more than $30,000$ uncertain parameters in the base models for FIT or COLO, by choosing $\epsilon = 0.05/30,000$, we guarantee (using a union bound argument) that *all* model parameters will fall within the uncertainty set with at least 95% probability. Using this construction, the certainty intervals will contain the 95% confidence intervals on the metrics of interest. In particular, if the certainty intervals for COLO and FIT are disjoint, we may still conclude that FIT is more cost-effective than COLO with $p_{\text{value}} < 0.05$.

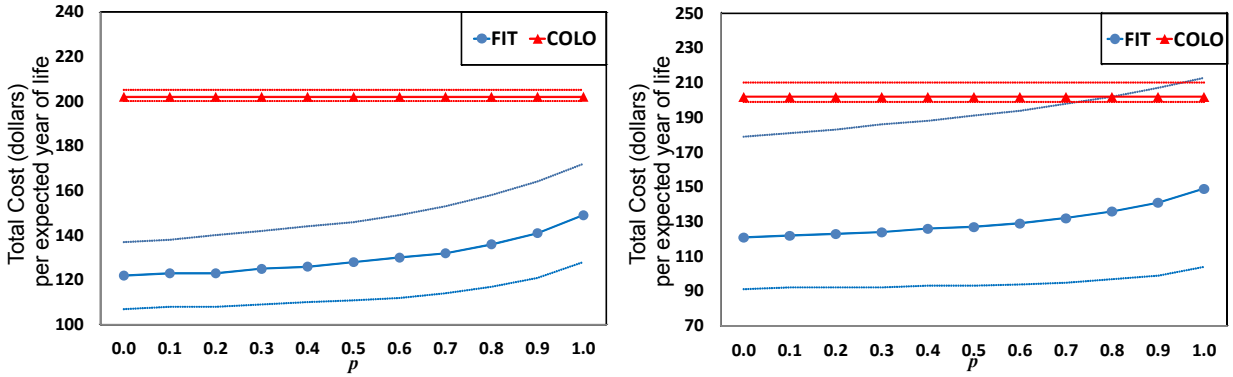

Figure 3: Nominal values and certainty intervals on the total cost in dollars per expected year of life from FIT and COLO as a function of the proportion p of polyps/CRC that are undetectable by repeated FIT. Left: base model. Right: using enlarged uncertainty sets.

The right panel of Figure 3 plots the costs per expected year of life and corresponding certainty intervals, as a function of p, using the enlarged intervals for the model parameters. The results strengthen our previous insight: For a large range of p (between 0.0 and 0.7), FIT is more cost-effective than COLO and now we havethe added guarantee that this holds with $p_{\text{value}} < 0.05$.

5.5 Probability of CRC Incidence

Even though life expectancy is the most common measure of health for CRC screening, one downside of solely relying on it is that only a small fraction about 5% of people ever get CRC throughout their remaining lifetime. We were interested to study how each screening modality performed along other metrics, specifically, the probability of ever being diagnosed with each stage of CRC. These can be viewed as asorption probability to certain states of the DTMC, and we applied the generalization of our model in §4.3 to compute certainty intervals around the quantities. Results are tabulated in

Table 3. From these results, NOSCREEN yields the poorest health outcomes overall as expected, but the comparison between COLO and FIT is not as clear as before. Nominally, FIT had fewer people on average being diagnosed with the more advanced stages of CRC (regional and distant) than COLO. However, the large widths of the certainty intervals around these quantities show that it does not dominate COLO along these metrics. In contrast, COLO nominally only had about 2/3 the number of people being diagnosed with the early-stage (i.e., localized) CRC, and is strictly better than FIT on this metric (since their certainty intervals do not overlap).

Screening Method	Localized CRC	Regional CRC	Distant CRC
NOSCREEN	2,393 (2,393 – 2,393)	2,169 (2,169 – 2,169)	1,197 (1,197 – 1,197)
COLO	1,002 (938 – 1,062)	651 (606 – 701)	266 (247 – 294)
FIT	1,559 (1,270 – 1,896)	510 (418 – 709)	211 (186 – 274)

Table 3: Nominal values and certainty intervals (in parentheses) of the average number of people (per 100,000) that are diagnosed with localized, regional, or distant CRC.

These results can be explained by the fact that COLO has much better test properties for detecting pre-cancerous polyps (both large and small) than FIT, which is the primary means through which localized CRC develops (see Table 1). Moreover, since the proportion of people who get diagnosed with the more advanced CRC stages are about the same, the primary cost savings of FIT relative to COLO (see Table 2) come from its much lower per-screening cost than colonoscopy.

6 Conclusion

For cost-effectiveness studies of innovations that employ Markov models, data insufficiency can result in uncertain estimates of many entries of the transition matrix. In this paper, we formulate the problem of assessing the aggregate effect of model uncertainty in a Markov model as a pair of optimization problems. Using this optimization formulation, we show that if uncertainty sets do not have a row-wise structure, the problem of computing certainty intervals is generally computationally intractable (NP-hard). Conversely, if uncertainty sets possess a row-wise structure, we prove that the problem has an equivalent formulation as a MDP, and can be solved efficiently, either directly using robust optimization or iteratively through value or policy iteration. In particular, we showed how the iterative approaches are suitable for problems with a large state space, and that they can be made more efficient by exploiting special problem structures.

We applied our model to analyze the cost-effectiveness of FIT for CRC screening. Our results support the growing medical consensus (Whitlock et al. 2008, Parekh et al. 2008, Telford et al. 2010, Heitman et al. 2010) that annual FIT screening is a viable and cost-effective screening modality for CRC, nominally costing around two-thirds as much as 10-year colonoscopy screening per expected year of life. Moreover, our analysis is the first that guarantees that the cost-effectiveness of FIT relative to colonoscopy persists even after accounting for all the uncertainties within FIT's characteristics. Even assuming that repeated FIT tests are perfectly correlated, FIT still costs about 85% as much as colonsocopy per expected year of life after accounting for model uncertainty.

One caveat revealed by our analysis is that FIT has a higher incidence rate of localized CRC than colonoscopy, and its incidence rates for more advanced CRCs (regional and distant) are nominally lower, but within the range of aggregate model uncertainty. Therefore, our results must be interpreted with caution: It is *not* true that FIT is better for overall health, especially since the incidence of various CRCs are important outcomes in the own right. For example, getting diagnosed with localized CRC can itself be deeply psychologically stressful to patients. Our results suggest that if life expectancy is the primary health outcome of interest, then on average, FIT delivers better health per dollar spent, and should at least be considered as an alternative, if not preferred, screening modality by physicians, insurers, and policymakers. Of course, the most accurate assessment of the comparative value of FIT screening must be performed by randomized clinical trials, several of which are already in progress (Quintero et al. 2012, Dominitz and Robertson 2013).

Although our model was motivated by the problem of uncertainty in cost-effectiveness studies for medical innovations, it is not limited to this application. More generally, our methods can be applied to analyze cases where parameter uncertainties may be substantial in DTMCs. In particular, in applications where the transition parameters are estimated from data, there will be sampling error in the elements of the transition matrix, and our methods could be applied to assess the aggregate impact of these sampling errors on various quantities of interest.

Acknowledgements. We are grateful to Dan Iancu, Melvyn Sim, Martin Puterman, Evan Porteus, Jeffrey Pfeffer, and Yinyu Ye for helpful discussions on this subject and for pointing us to useful references and related papers.

References

Alizadeh, F., D. Goldfarb. 2003. Second-order cone programming. *Math Program* **95**(1) 3–51.

Allison, J. E., I. S. Tekawa, L. J. Ransom, A. L. Adrain. 1996. A comparison of fecal occult-blood tests for colorectal-cancer screening. *New Engl J Med* **334**(3) 155–160.

Arias, E. 2011. United States life tables, 2007. *National Vital Statistics Reports* **59**(9) 1–60.

Bandi, C., D. Bertsimas. 2012. Tractable stochastic analysis in high dimensions via robust optimization. *Math Program* 1–48.

Ben-Tal, A., D. den Hertog, A. De Waegenaere, B. Melenberg, G. Rennen. 2013. Robust solutions of optimization problems affected by uncertain probabilities. *Manage Sci* **59**(2) 341–357.

Ben-Tal, A., L. El Ghaoui, A. Nemirovski. 2009. *Robust Optimization*. Princeton University Press.

Ben-Tal, A., A. Nemirovski. 1998. Robust convex optimization. *Math Oper Res* **23**(4) 769–805.

Bertsimas, D., M. Sim. 2004. The price of robustness. *Oper Res* **52**(1) 35–53.

Blanc, H., D. Den Hertog. 2008. On Markov chains with uncertain data. *Working Paper, Tilburg University, Center for Economic Research* 1–20.

Borkar, V. S., V. Ejov, J. A. Filar. 2009. On the Hamiltonicity gap and doubly stochastic matrices. *Random Struct Algor* **34**(4) 502–519.

Brandeau, M., D. Owens, C. Sox, R. Wachter. 1992. Screening women of childbearing age for human immunodeficiency virus: A cost-benefit analysis. *Arch Intern Med* **152**(11) 2229–2237.

Briggs, A., M. Sculpher, M. Buxton. 1994. Uncertainty in the economic evaluation of health care technologies: The role of sensitivity analysis. *Health Econ* **3**(2) 95–104.

De Cooman, G., F. Hermans, E. Quaeghebeur. 2009. Imprecise Markov chains and their limit behavior. *Probab Eng Inform Sc* **23**(4) 597–635.

Dominitz, J. A., D. J. Robertson. 2013. Colonoscopy versus fecal immunochemical test in reducing mortality from colorectal cancer (confirm). URL `http://clinicaltrials.gov/show/NCT01239082`. Clinical Trials Identifier: NCT01239082. Accessed Feb 26, 2013.

Duffy, S., H. Chen, L. Tabar, N. Day. 1995. Estimation of mean sojourn time in breast cancer screening using a Markov chain model of both entry to and exit from the preclinical detectable phase. *Stat Med* **14**(14) 1531–1543.

Eddy, D. 1990. Screening for colorectal cancer. *Ann Intern Med* **113**(5) 373–384.

Frazier, A., G. Colditz, C. Fuchs, K. Kuntz. 2000. Cost-effectiveness of screening for colorectal cancer in the general population. *J Amer Med Assoc* **284**(15) 1954–1961.

Goh, J., M. Sim. 2011. Robust Optimization Made Easy with ROME. *Oper Res* **59**(4) 973–985.

Greenberg, P. D., L. Bertario, R. Gnauck, O. Kronborg, J. Hardcastle, M. S. Epstein, D. Sadowski, R. Sudduth, G. R. Zuckerman, D. C. Rockey. 2000. A prospective multicenter evaluation of new fecal occult blood tests in patients undergoing colonoscopy. *Am J Gastroenterol* **95**(5) 1331–1338.

Heitman, S. J., R. J. Hilsden, F. Au, S. Dowden, B. J. Manns. 2010. Colorectal cancer screening for average-risk North Americans: An economic evaluation. *PLoS Med* **7**(11) e1000370.

Hermans, F., G. De Cooman. 2012. Characterisation of ergodic upper transition operators. *Int J Approx*

Reason **53**(4) 573–583.

Hillner, B., T. Smith. 1991. Efficacy and cost effectiveness of adjuvant chemotherapy in women with node-negative breast cancer. *New Engl J Med* **324**(3) 160–168.

Hutton, D., S. So, M. Brandeau. 2010. Cost-effectiveness of nationwide hepatitis B catch-up vaccination among children and adolescents in China. *Hepatology* **51**(2) 405–414.

Hutton, D., D. Tan, S. So, M. Brandeau. 2007. Cost-effectiveness of screening and vaccinating Asian and Pacific Islander adults for Hepatitis B. *Ann Intern Med* **147**(7) 460–469.

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

Jönsson, L., P. Lindgren, A. Wimo, B. Jönsson, B. Winblad, et al. 1999. The cost-effectiveness of donepezil therapy in Swedish patients with Alzheimer's disease: A Markov model. *Clin Ther* **21**(7) 1230–1240.

Joseph, D., J. King, J. Miller, L. Richardson. 2012. Prevalence of colorectal cancer screening among adults – Behavioral Risk Factor Surveillance System, United States, 2010. *Morbidity and Mortality Weekly Reports Supplement* **61**(2) 51–56.

Karp, R. M. 1972. Reducibility among combinatorial problems. J. D. B. Raymond E. Miller, James W. Thatcher, ed., *Complexity of Computer Computations*. Plenum Press, New York, 85–103.

Kobelt, G., L. Jönsson, A. Young, K. Eberhardt. 2003. The cost-effectiveness of infliximab (remicade®) in the treatment of rheumatoid arthritis in Sweden and the United Kingdom based on the ATTRACT study. *Rheumatology* **42**(2) 326–335.

Ladabaum, U., C. Chopra, G. Huang, J. Scheiman, M. Chernew, A. Fendrick, et al. 2001. Aspirin as an adjunct to screening for prevention of sporadic colorectal cancer. A cost-effectiveness analysis. *Ann Intern Med* **135**(9) 769–781.

Ladabaum, U., K. Song, A. Fendrick. 2004. Colorectal neoplasia screening with virtual colonoscopy: when, at what cost, and with what national impact? *Clin Gastroenterol H* **2**(7) 554–563.

Löfberg, J. 2004. YALMIP: A toolbox for modeling and optimization in MATLAB. *IEEE International Symposium on Computer Aided Control Systems Design*. Taipei, Taiwan.

Manne, A. S. 1960. Linear programming and sequential decisions. *Manage Sci* **6**(3) 259–267.

Morikawa, T., J. Kato, Y. Yamaji, R. Wada, T. Mitsushima, Y. Shiratori. 2005. A comparison of the immunochemical fecal occult blood test and total colonoscopy in the asymptomatic population. *Gastroenterology* **129**(2) 422–428.

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

Parekh, M., A. Fendrick, U. Ladabaum. 2008. As tests evolve and costs of cancer care rise: Reappraising stool-based screening for colorectal neoplasia. *Aliment Pharm Therap* **27**(8) 697–712.

Quintero, E., A. Castells, L. Bujanda, J. Cubiella, D. Salas, Á. Lanas, M. Andreu, F. Carballo, J. D. Morillas, C. Hernández, et al. 2012. Colonoscopy versus fecal immunochemical testing in colorectal-cancer screening. *New Engl J Med* **366**(8) 697–706.

Ross, K., H. Carter, J. Pearson, H. Guess. 2000. Comparative efficiency of prostate-specific antigen screening strategies for prostate cancer detection. *J Amer Med Assoc* **284**(11) 1399–1405.

Rozen, P., J. Knaani, Z. Samuel. 1997. Performance characteristics and comparison of two immunochemical and two guaiac fecal occult blood screening tests for colorectal neoplasia. *Digest Dis Sci* **42**(10) 2064–2071.

Rozen, P., J. Knaani, Z. Samuel. 2000. Comparative screening with a sensitive guaiac and specific immunochemical occult blood test in an endoscopic study. *Cancer* **89**(1) 46–52.

Sanders, G., M. Hlatky, D. Owens. 2005. Cost-effectiveness of implantable cardioverter–defibrillators. *New Engl J Med* **353**(14) 1471–1480.

Škulj, D. 2007. Regular finite Markov chains with interval probabilities. *5th International Symposium on Imprecise Probability: Theories and Applications*. Prague, Czech Republic, 405–413.

Škulj, D., R. Hable. 2009. Coefficients of ergodicity for imprecise Markov chains. *6th International Symposium on Imprecise Probability: Theories and Applications*. Durham, United Kingdom, 377–386.

Sonnenberg, F., J. Beck. 1993. Markov models in medical decision making a practical guide. *Med Decis Making* **13**(4) 322–338.

Telford, J. J., A. R. Levy, J. C. Sambrook, D. Zou, R. A. Enns. 2010. The cost-effectiveness of screening for colorectal cancer. *Can Med Assoc J* **182**(12) 1307–1313.

U.S. Cancer Statistics Working Group. 2013. *United States Cancer Statistics: 1999 – 2009 Incidence and Mortality Web-based Report.*. U.S. Department of Health and Human Services, Centers for Disease Control and Prevention and National Cancer Institute. URL `http://www.cdc.gov/uscs`. Last Accessed Feb 11, 2013.

Whitlock, E. P., J. Lin, E. Liles, T. Beil, R. Fu, E. O'Connor, R. N. Thompson, T. Cardenas. 2008. Screening for colorectal cancer: An updated systematic review. Tech. Rep. 08-05-05124-EF-1, Agency for Healthcare Research and Quality, Rockville (MD).

Winawer, S., R. Fletcher, D. Rex, J. Bond, R. Burt, J. Ferrucci, T. Ganiats, T. Levin, S. Woolf, D. Johnson, et al. 2003. Colorectal cancer screening and surveillance: Clinical guidelines and rationale–update based on new evidence. *Gastroenterology* **124**(2) 544–560.

Wong, B.-Y., W. Wong, K. Cheung, T. Tong, P. Rozen, G. Young, K. Chu, J. Ho, W. Law, H. Tung, et al. 2003. A sensitive guaiac faecal occult blood test is less useful than an immunochemical test for colorectal cancer screening in a Chinese population. *Aliment Pharm Therap* **18**(9) 941–946.

Xu, H., S. Mannor. 2012. Distributionally robust Markov decision processes. *Math Oper Res* **37**(2) 288–300.

Zorzi, M., A. Barca, F. Falcini, G. Grazzini, R. Pizzuti, A. Ravaioli, d. B. P. Sassoli, C. Senore, A. Sigillito, M. Vettorazzi, et al. 2007. Screening for colorectal cancer in Italy: 2005 survey. *Epidemiol Prev* **31**(2-3 Suppl 2) 49–60.

Appendix: Proofs

Proof of Proposition 1.. Our proof proceeds in three steps. First, we will prove a lower bound on the optimal value of (3), namely,

$$\underline{f} \geq \frac{1}{1-\alpha^N}. \tag{10}$$

Suppose that (3) is feasible (otherwise, $\underline{f} = +\infty$ by convention and the inequality holds trivially) and let $\mathbf{P}$ be one such feasible solution. Let the probability law induced by $\mathbf{P}$ be $\mathbb{P}$ and denote the expectation under $\mathbb{P}$ by $\mathbb{E}$. From the second constraint of (3), $\mathbf{e}/N$ is a stationary distribution for the chain and therefore all states in V, including state 1, are positive recurrent states. Let $V_1 = V_1(\mathbf{P}) \subseteq V$ be the communicating class of the chain that contains state 1. Consider the Markov chain that is restricted to V_1, which has a stationary distribution $\boldsymbol{\pi}_{V_1} = \mathbf{e}/\left|V_1\right| \in \mathbb{R}^{|V_1|}$. For $k = 1, \ldots,$ let τ_{1k} be a random variable representing the kth return time to state 1. Because the restricted Markov Chain on V_1 is irreducible and positive recurrent, $\boldsymbol{\pi}_{V_1}$ is the unique stationary distribution, and the kth expected return time to state 1 is given by $\mathbb{E}\left(\tau_{1k}\right) = k\left|V_1\right|$. Since rewards are only obtained in state 1, we have

$$\mathbb{E}\left(\sum_{k=0}^{\infty} \alpha^{\tau_{1k}}\right) = \sum_{k=0}^{\infty} \mathbb{E}\left(\alpha^{\tau_{1k}}\right) \geq \sum_{k=0}^{\infty} \alpha^{k|V_1|} = \frac{1}{1-\alpha^{|V_1|}} \geq \frac{1}{1-\alpha^N}, \tag{11}$$

where the first equality is by the Monotone Convergence Theorem, the first inequality is by Jensen's inequality applied to the function $x \mapsto \alpha^x$, and the final inequality is because $|V_1| \leq |V| = N$. Since (11) holds for any feasible $\mathbf{P}$ (and its associated probability measure $\mathbb{P}$), we have proven (10).

Second, we will prove that equality holds in (10) iff there exists an optimal $\mathbf{P}^*$ in (3) that induces a probability law $\mathbb{P}$ so that $\tau_{11} = N$ $\mathbb{P}$-a.s.. Because $x \mapsto \alpha^x$ is *strictly* convex, for the first inequality of (11) to hold with equality, we require $\mathbb{P}\left(\tau_{11} = |V_1|\right) = 1$. In order for the second inequality of (11) to hold with equality, we require $|V_1| = N$. Therefore, $\tau_{11} = N$, $\mathbb{P}$-a.s. is necessary for $\mathbf{P}^*$ to be optimal. Sufficiency follows immediately by direct substitution.

Third, we will prove that G contains a Hamiltonian cycle iff $\underline{f} = \frac{1}{1-\alpha^N}$. For the forward direction, suppose that G contains a Hamiltonian cycle $C \subseteq E$. Then, the binary matrix $\mathbf{P}^*$ such that $P^*_{ij} = \mathbb{1}_{(i,j)\in C}$ is both feasible in (3) and attains the lower bound, giving $\underline{f} = \frac{1}{1-\alpha^N}$. For the converse direction, suppose that $\underline{f} = \frac{1}{1-\alpha^N}$ holds. From the previous step, there exists some optimal $\mathbf{P}^*$ that induces a probability law $\mathbb{P}$ such that $\tau_{11} = N$, $\mathbb{P}$-a.s.. Pick any such optimal $\mathbf{P}^*$, and let $\mathbb{P}$ be the probability law it induces. Denote the state of the Markov chain in period $N-1$ by X_{N-1}, and let $V_{N-1} \subseteq V \setminus \{1\}$ represent the set of all possible states for X_{N-1}, $\mathbb{P}$-

a.s.. Since $X_N = 1$, $\mathbb{P}$-a.s. (to close the chain), it is necessary that $P^*_{i1} = 1$ for each $i \in V_{N-1}$. However, from the second constraint of (3), it is also necessary that $\sum_{i=1}^N P^*_{i1} = 1$. Since $\mathbf{P}^*$ is constrained to be nonnegative, this implies that V_{N-1} is a singleton, i.e., there exists some unique $v_{N-1} \in V \setminus \{1\}$ with $(v_{N-1}, 1) \in E$ such that $V_{N-1} = \{v_{N-1}\}$. Moreover, the chain cannot have a positive probability of being in v_{N-1} at any period before $N-1$. Otherwise, $\mathbb{P}(\tau_{11} < N) > 0$, which contradicts the assumption that $\tau_{11} = N$, $\mathbb{P}$-a.s.. Next, consider the state of the chain at period $N-2$, X_{N-2}, and let $V_{N-2} \setminus \{1, v_{N-1}\}$ be the set of all possible states for X_{N-2}, $\mathbb{P}$-a.s.. By an identical argument, we can show that V_{N-2} must be a singleton too, so that there is some unique $v_{N-2} \in V \setminus \{1, v_{N-1}\}$ with $(v_{N-2}, v_{N-1}) \in E$, such that $V_{N-2} = \{v_{N-2}\}$. By repeating this construction until termination, we end up with a set of *unique* states $v_1, \ldots, v_{N-1} \in V \setminus \{1\}$ such that $C = \{(1, v_1), (v_1, v_2), \ldots, (v_{N-1}, 1)\} \subseteq E$, so that C is a Hamiltonian cycle of G. $\square$

Proof of Theorem 1.. First, we prove the lower bound for general $\mathcal{P}$. Consider any $\mathbf{v}$ that is feasible in (4) and any $\mathbf{P} \in \mathcal{P}$. Then, since $(\mathbf{I} - \alpha\mathbf{P})^{-1}$ is component-wise nonnegative, the inequality $\mathbf{v} \leq (\mathbf{I} - \alpha\mathbf{P})^{-1}\mathbf{r}$ holds component-wise for any $\mathbf{P} \in \mathcal{P}$. Since $\boldsymbol{\pi_0} \geq \mathbf{0}$, this implies that $\boldsymbol{\pi_0}'\mathbf{v} \leq \boldsymbol{\pi_0}'(\mathbf{I} - \alpha\mathbf{P})^{-1}\mathbf{r}$ for any $\mathbf{P} \in \mathcal{P}$. Hence, $\boldsymbol{\pi_0}'\mathbf{v} \leq \inf_{\mathbf{P}\in\mathcal{P}} \boldsymbol{\pi_0}'(\mathbf{I} - \alpha\mathbf{P})^{-1}\mathbf{r} = \underline{f}$. Maximizing over $\mathbf{v}$ on the LHS implies the result.

To prove the reverse inequality, we now assume that $\mathcal{P}$ satisfies Assumption 1. Let $\mathbf{v^*}$ satisfy $v^*_i = r_i + \alpha \sum_{k=1}^K u_{ki} \min_{\mathbf{z_k}\in\mathcal{Z}_k} \{\mathbf{z_k}'\mathbf{v^*}\}$. A unique such $\mathbf{v}^*$ is guaranteed to exist by the contraction mapping principle. Next, let $\mathbf{z^*_k} \in \arg\min_{\mathbf{z_k}\in\mathcal{Z}_k} \mathbf{z_k}'\mathbf{v^*}$, for each $k = 1, \ldots, K$, and define $\mathbf{P^*} := \sum_{k=1}^K \mathbf{u_k}\mathbf{z_k^{*\prime}}$. Note that $\mathbf{v^*}$ is feasible in (4), which implies $\boldsymbol{\pi_0}'\mathbf{v^*} = \boldsymbol{\pi_0}'(\mathbf{I} - \alpha\mathbf{P^*})^{-1}\mathbf{r} \leq f_*$. Finally, because $\mathbf{P^*} \in \mathcal{P}$, we have $\underline{f} = \min_{\mathbf{P}\in\mathcal{P}} \boldsymbol{\pi_0}'(\mathbf{I} - \alpha\mathbf{P})^{-1}\mathbf{r} \leq \boldsymbol{\pi_0}'(\mathbf{I} - \alpha\mathbf{P^*})^{-1}\mathbf{r} \leq f_*$. $\square$

Proof of Proposition 2.. Throughout this proof, we will fix some $k = 1, \ldots, K$ and t, and therefore omit explicitly writing the dependence on k and t for brevity. Without loss of generality, assume that $\mathbf{v}$ is already sorted in ascending order. We also assume that $\underline{\mathbf{z}} < \overline{\mathbf{z}}$ component-wise (otherwise, we can just remove the degenerate indices, i.e., those i where $\underline{z}_i = \overline{z}_i$ and rescale). Each update for (6) solves the optimization problem

$$\begin{aligned} \min_{\mathbf{z}} \quad & \mathbf{v}'\mathbf{z} \\ \text{s.t.} \quad & \underline{\mathbf{z}} \leq \mathbf{z} \leq \overline{\mathbf{z}}, \quad \mathbf{e}'\mathbf{z} = 1, \end{aligned} \tag{12}$$

where we recall $\mathbf{e}$ is the vector of all ones. By a change-of-variable $y_i = \frac{z_i - \underline{z}_i}{\overline{z}_i - \underline{z}_i}$, we observe that (12) becomes the linear relaxation of the one-dimensional knapsack problem, and is well-known to have an optimal solution with a "threshold" structure, which, after transforming back to the original decision variables, yields (7). $\square$