

Population-level intervention and information collection in dynamic healthcare policy

Lauren E. Cipriano¹  · Thomas A. Weber²

Received: 23 February 2017 / Accepted: 10 August 2017 / Published online: 8 September 2017
© The Author(s) 2017. This article is an open access publication

Abstract We develop a general framework for optimal health policy design in a dynamic setting. We consider a hypothetical medical intervention for a cohort of patients where one parameter varies across cohorts with imperfectly observable linear dynamics. We seek to identify the optimal time to change the current health intervention policy and the optimal time to collect decision-relevant information. We formulate this problem as a discrete-time, infinite-horizon Markov decision process and we establish structural properties in terms of first and second-order monotonicity. We demonstrate that it is generally optimal to delay information acquisition until an effect on decisions is sufficiently likely. We apply this framework to the evaluation of hepatitis C virus (HCV) screening in the general population determining which birth cohorts to screen for HCV and when to collect information about HCV prevalence.

Keywords Medical decision making · Markov decision processes · Hepatitis C virus · Optimal stopping · Dynamic programming

1 Introduction

There is currently no guidance for determining the optimal schedule for collecting additional information regarding a decision to invest in a health program or technology [1, 2]. Current practice in the health decision science literature assumes that model parameters are fixed across cohorts and the value of additional information is calculated assuming the information-collection effort is initiated immediately [3–5]. However, in many cases the cost-effectiveness of a health program or technology – and, therefore, the value of additional information about one or more model parameters – may be changing over time because of trends affecting the cohort or the intervention [6]. In these cases, collecting additional information immediately may not be optimal and value-of-information calculations based on static parameter assumptions are likely to be biased. Planning over longer horizons is particularly important in health policy because, once established, clinical practice is difficult to change due to high switching costs (re-training and potentially new capital equipment expenditures), particularly if it appears that the level of service is being reduced [7].

In this paper we apply a stochastic dynamic programming approach to identify both the optimal time to change the current health intervention policy and the optimal time to collect decision-relevant information. We consider a hypothetical medical intervention for a cohort of patients. At each time, a new cohort of patients becomes eligible for the intervention and one parameter varies across the cohorts with imperfectly observable linear dynamics. We assume that the value of the intervention is linear in the dynamic parameter. In general, the (incremental) net monetary benefit of an intervention is linear in parameters with a one-time

✉ Lauren E. Cipriano
lcipriano@ivey.uwo.ca

Thomas A. Weber
thomas.weber@epfl.ch

¹ Ivey Business School, Western University, 1255 Western Road, London ON N6G 0N1, Canada

² Ecole Polytechnique Fédérale de Lausanne, CDM-ODY 3.01, Station 5, CH-1015 Lausanne, Switzerland

effect (e.g., the prevalence of a disease at one point in time or the outcome of a one-time screening test). When an effect accrues over time, such as for a reduction in the annual transition rate of a disease complication or death, linearity is often used as an approximation (see, e.g., [8]). At each time, the policy-maker can choose to invest in the medical intervention and/or to purchase sample information about the uncertain dynamic parameter. We demonstrate that information acquisition is best delayed until the signal is sufficiently likely to affect the optimal policy decision.

We apply this framework to the evaluation of hepatitis C virus (HCV) screening. Prior to the development of highly-effective treatments, HCV screening in the general population was not considered cost-effective [9] and universal screening was not recommended [10]. The advent of more effective therapy has changed the value of identifying infected individuals early to initiate treatment [11–15]. Recently released guidance by the Centers for Disease Control and Prevention (CDC) and the US Preventive Services Task Force (USPSTF) recommends one-time HCV screening for all individuals born between 1945 and 1965 [16, 17] although screening individuals born after 1965 may also be cost effective [13–15]. Based on our primary analysis of the National Health and Nutrition Examination Survey (NHANES), in the US general population, HCV prevalence is highest in people born around 1956 and declines thereafter at a rate of approximately 11% per birth year. Since HCV prevalence is decreasing across birth cohorts, HCV screening will only be cost-effective for a limited time or for a limited set of birth cohorts. We apply our model to simultaneously evaluate the optimal HCV-screening and information-acquisition policy.

Specifically, we apply our model to the policy decision of whether or not to perform one-time HCV screening in successive cohorts of healthy 50-year olds, who have not previously been tested for HCV, at a routine preventive health visit. Applying a traditional health economics framework, the policy-maker could decide today how many cohorts will be screened (e.g., each cohort of 50-year olds until those born in 1965 turn 50) or, to inform this decision, the policy-maker may seek additional information to be collected immediately. Our framework differs from the traditional paradigm in that each year the policy-maker makes a decision about whether to continue the one-time HCV screening program (whether or not to screen the new cohort of healthy 50-year olds) and whether to collect information about disease prevalence in this current cohort. If information is never collected, the optimal policy does not differ across frameworks. However, in our framework, the immediate decision is not limited to the decision of when to change policies, but it also includes when to collect information to inform a future change of policy. For example, the (immediately) optimal policy might be to screen each

cohort of 50-year olds for the next 6 years and then collect information about HCV prevalence to inform future decision making. Delaying information acquisition until a time that the information is sufficiently likely to affect the decision increases the value of the information. In addition, from a practical perspective, collecting information years before it is likely to influence a policy change wastes immediate resources and, should something occur in the lag-time between the information-acquisition effort and the policy change, implementing the pre-determined policy change may not be optimal.

1.1 Related literature and contribution

The relevant literature spans technology adoption, dynamic decisions in healthcare, and the value of information in healthcare.

Technology adoption In technology-adoption models, a decision-maker considers the adoption of a technology of unknown profitability. Jensen [18] introduced a model in which information about a new technology is costlessly observed and the decision-maker can decide to adopt the new technology at any point in time. McCardle [19] presented a model in which collecting information is associated with a fixed cost; in each period the decision-maker can defer and collect information, or make a final decision to accept or reject the new technology. The optimal policy in each period is characterized by two thresholds: if the expected benefit is above the upper threshold, it is optimal to adopt the technology; if the expected benefit is below the lower threshold, it is optimal to reject the technology; and, if the expected benefit is between these two thresholds the optimal strategy is to gather information. Uncertainty about the technology's value decreases over time and the two thresholds converge to the cost of adoption. Smith and McCardle [20] provided several meta-results, some of which we use, describing how properties of the value function of a stochastic dynamic program are preserved and propagated through finite-horizon Markov-reward and decision processes. Ulu and Smith [21] extended this work by relaxing the assumption that the decision-maker's value of the technology can be summarized by the expected benefit, and they use more general monotone-comparative-statics techniques in terms of likelihood orders to generalize the class of signals that are observed prior to making an adoption decision.

Another line of research considered technologies, like ours, with uncertain and changing value. Rosenberg [22] found that expectation of technological improvement may delay a firm's irreversible technology investments. Bessen [23] calculated the option value of delay for such a problem. Kornish [24] considered the choice between two uncertain

technologies where each is subject to a positive network effect and explored the impact of the network effect on the optimal adoption policy. Chambers and Kouvelis [25] formulated a technology-adoption problem incorporating expected learning-curve effects.

Stochastic dynamic programs in healthcare Sequential decisions under uncertainty are common in healthcare [26, 27]. Most healthcare applications of stochastic dynamic programs have focused on optimizing the timing of interventions for an individual patient: the decision to accept or reject an offered kidney for transplantation [28]; the optimal treatment plan for mild spherocytosis [29]; the optimal surveillance and management of ischemic heart disease [30]; the optimal time to perform a living-donor liver transplant [31, 32]; the optimal time to initiate treatment for HIV [33, 34]; the optimal timing and frequency of HCV testing from the patient perspective [35]; the optimal use of statins in patients with type 2 diabetes [36, 37]; the optimal prostate biopsy referral [38]; and, optimal cancer screening programs [39, 40]. Dynamic programming has also been applied to complex appointment scheduling problems in healthcare, including problems with patients of different clinical types/priority [41, 42]; incorporating patient no-shows [43]; and problems of sequential appointment scheduling with the objective of closely adhering to a prescribed schedule (e.g., sequential chemotherapy appointments [44]) or with the objective of satisfying patient preferences [45, 46]. Fewer examples of application to population-level policy exist. Kornish and Keeney [47] and Özaltın et al. [48] formulated the influenza-strain selection problem in a finite-horizon optimal-stopping framework. Similar to our problem, the influenza-vaccine composition decision is also an optimal-stopping problem with information acquisition; however, it has many unique characteristics that distinguish it from the problem discussed here such as an inventory deadline (finite horizon), a product useful for one season only, and a time-consuming production process. Similar to many of the technology-adoption models discussed above but unlike our framework, in the influenza-vaccine composition models, information is collected in every period in which a final decision has not yet been made.

Health economics and value of information in healthcare Cost-effectiveness analysis is an economic method for comparing the lifetime discounted costs and health benefits associated with two or more medical interventions or health programs [1, 2]. In theory, the optimal allocation of resources across a portfolio of health interventions is determined by solving a constrained optimization problem with the objective of maximizing health benefits subject to a budget constraint [49–51]. In reality, regional and national

health policy bodies routinely compare the incremental cost effectiveness ratio of candidate interventions to a pre-determined threshold intended to approximate the shadow price of the budget to determine if the intervention is ‘cost-effective’ as one component of their policy-making process [52]. Cost-effectiveness analysis is widely used to evaluate general population screening for relatively rare conditions because these programs impose a small cost on everyone who is screened and provide substantive healthcare gains for only a small number of individuals who are identified (or identified earlier than they would be otherwise); calculating the population-level costs and benefits can require detailed natural history models, extensive model calibration and validation, and thorough analysis.

Bayesian decision theory approaches to value-of-information assessment were first introduced by Raiffa and Schlaifer [53]. Weinstein [54] proposed the widespread adoption of value-of-information analysis to research priority setting in health policy and medicine. Hornberger et al. [55], Claxton and Posnett [56], and Claxton [57] introduced a Bayesian approach to identifying the optimal trial sample size and to assessing the value of additional information for technology-adoption assessments. Several approaches to increasing the accuracy of value-of-information calculations continued to relax assumptions implicit in the original formulation (see examples in [58–62]). One common assumption in these studies is that the currently estimated per-person value of information can be applied to individuals in all future cohorts. Recognizing some of the implications of this assumption, Philips et al. [6] discussed the impact that intervention-horizon uncertainty, price changes, and technological development can have on the per-person value of information for future cohorts. They find that delaying information collection may be desirable but do not provide a framework for determining the optimal time to collect information.

Contribution In this paper, we extend the technology-adoption literature by allowing for a technology that is changing in value over time, for the opportunity to ‘wait’ without collecting information, and for the possibility of optimally determining the collected amount of information in each period. We also incorporate the possibility of an imperfect information-collection technology. We broaden the scope of applications of stochastic dynamic programs in the area of healthcare in an important way – focusing on population policy rather than patient-level decisions. We extend the health decision science literature on value-of-information assessment by developing an approach to identify the optimal information-acquisition policy when model parameters are varying across cohorts. Finally, as an example, we apply our framework to the timely public policy problem of developing a population screening program

for HCV. We find that considering the opportunity to collect information in the future leads to a substantially different policy recommendation than current guidelines because it explicitly considers and addresses the parameter uncertainty which is changing over time.

2 The model

A policy-maker faces recurring decisions for cohorts arriving at times $t \in \{0, 1, 2, \dots\}$ about whether to invest in a health intervention delivered once per cohort (of size N). By cohort we mean a group of individuals with a certain medical presentation (i.e., individuals with a new diagnosis of cancer) or of a certain status (i.e., individuals who turned 50 this year). The policy-maker's objective is to maximize net monetary benefit from a societal perspective. The per-person incremental net monetary benefit (INMB) of performing the intervention compared to the *status quo* is assumed to be affine in an uncertain parameter \tilde{p}_t that varies across the cohorts, with realizations in $[0, 1]$ and known dynamics. So $\text{INMB}_t = \theta \tilde{p}_t - \gamma$, for all $t \geq 0$, where θ is the marginal INMB (with respect to the parameter \tilde{p}_t) and $-\gamma$ is the fixed INMB, both measured on a per-person basis.

At the beginning of period t , the policy-maker simultaneously decides whether to invest in a medical intervention for the individuals in cohort t and whether to conduct a study of sample size n_t over the period to obtain a better estimate of the uncertain parameter \tilde{p}_t . Information, if sought, arrives at the end of the current period and is used, together with the known dynamics of \tilde{p}_t , to inform the intervention decision for future cohorts. Let $d_t \in \mathcal{D} = \{0, 1\}$ denote the intervention decision at time t , where $d_t = 0$ indicates 'No intervention' and $d_t = 1$ indicates 'Intervention.' The amount of information collected is measured in terms of the sample size $n_t \in \mathcal{N} = \{0, \dots, N\}$; it is obtained at the cost $K(n_t)$, where $K(\cdot)$ is an increasing function including a fixed and a variable cost when $n_t > 0$ and $K(0) = 0$. Thus, at each time t the policy-maker implements the control $u_t = (d_t, n_t) \in \mathcal{D} \times \mathcal{N}$. The per-person current reward for the cohort in period t is

$$g(\tilde{p}_t, u_t) = d_t(\theta \tilde{p}_t - \gamma) - \frac{K(n_t)}{N}. \quad (1)$$

The application in Section 4 features the decision problem of when to stop a once-in-a-lifetime disease-screening program where \tilde{p}_t is the uncertain disease prevalence in the t -th cohort which, in expectation, is geometrically decreasing over time; $\theta > 0$ denotes the marginal benefit of early

diagnosis and treatment for an affected individual, $\gamma > 0$ is the per-person cost of the program, and the current-period INMB g is increasing in \tilde{p}_t . Beyond our leading example, the framework can accommodate a wide variety of problems. As formulated, the uncertain parameter needs to lie in a compact interval (which can be mapped via bijection to $[0, 1]$). Thus, the parameter can represent not only a probability but also other model parameters, such as a quality-of-life weight or cost. Additionally, our analysis assumes that the parameter value is decreasing over time. To model a situation where the expectation of the uncertain parameter is increasing (e.g., obesity prevalence), the problem can be formulated as one in which a parameter of opposite definition is decreasing (e.g., prevalence of individuals who are not obese). Our exposition involves an example of when to stop a health intervention. However, the framework can also be used in situations in which the decision-maker wishes to identify the optimal time to initiate a new intervention (e.g., when to adopt a new surgical technique). More broadly, our framework can be applied in settings in which the decision-maker wishes to identify the optimal time to stop the current intervention or initiate a new intervention; the uncertain parameter is geometrically increasing or decreasing across intervention cohorts; and the current-period reward function is linearly increasing or decreasing in the uncertain parameter. Examples are shown in Table 1.

2.1 The information-acquisition problem

The policy-maker's prior belief about \tilde{p}_t at $t = 0$ is beta-distributed with distribution parameters $x_0 = (a_0, b_0)$. The posterior distribution when a beta-density is updated in a Bayesian manner with information collected using an imperfect information-collection technology is a mixture of beta-densities. Thus, in general, the policy-maker's prior beliefs about \tilde{p}_t at time t are in \mathcal{P} where \mathcal{P} denotes the set of measures which are a mixture of beta-densities. Specifically if $\tilde{p}_t \in \mathcal{P}$, then there exists parameters $x_{t,i} = (a_{t,i}, b_{t,i}) \in \mathbb{R}_{++}^2$ for all i where $1 \leq i \leq m$, $m \in \mathbb{R}_{++}$, and a set of non-negative weights ω_i such that $\sum_{i=1}^m \omega_i = 1$, where the distribution of \tilde{p}_t is a mixture of beta-densities of the form $\sum_{i=1}^m \omega_i \text{beta}(a_{t,i}, b_{t,i})$.

The policy-maker has the option to update his beliefs about the parameter \tilde{p}_t by testing n_t individuals at cost $K(n_t)$. The information-collection technology has binary test characteristics $q = (q_1, q_2)$, where q_1 is the sensitivity, q_2 is the specificity, and $q_1 + q_2 > 1$ (indicating the test is properly labeled). The terms 'sensitivity' and 'specificity' are often used to describe test accuracy in the medical literature. For clarity, we state their relation-

Table 1 Examples of alternative cases in which our framework applies

Case	Uncertain time-varying parameter $\tilde{p}_t = 1 - \tilde{q}_t$	INMB _t of “Intervention” vs. “No Intervention” [*]	Setting [†]	Example [‡]
A	\tilde{p}_t , decreasing in t	INMB _t is increasing in \tilde{p}_t , $\theta, \gamma > 0$	$\mu_p(x_0) > \frac{\gamma}{\theta}$ “Intervention” is currently implemented. Optimal stopping problem.	“Intervention”: General population HCV screening at age 50 (Section 4). Period reward function: INMB _t = $\theta\tilde{p}_t - \gamma$; \tilde{p}_t , prevalence of HCV in cohort t ; θ , marginal benefit of early diagnosis and treatment for an infected individual; γ , fixed cost of screening.
B	\tilde{p}_t , decreasing in t	INMB _t is decreasing in \tilde{p}_t , $\theta, \gamma < 0$	$\mu_p(x_0) > \frac{\gamma}{\theta}$ “Intervention” not currently implemented. Optimal starting problem.	“Intervention”: New surgical device vs. old device. Period reward function: INMB _t = $-\hat{\theta}\tilde{p}_t + \hat{\gamma}$; \tilde{p}_t , probability of an adverse event in device iteration t ; $\hat{\theta}$, incremental cost of an adverse event; $\hat{\gamma}$, benefit of surgical intervention without an adverse event. Problem transformation to framework: $\theta = -\hat{\theta}, \gamma = -\hat{\gamma}$.
C	\tilde{q}_t , increasing in t	INMB _t is decreasing in \tilde{q}_t (increasing in \tilde{p}_t), $\theta, \gamma > 0$	$\mu_q(x_0) < 1 - \frac{\gamma}{\theta} \Rightarrow \mu_p(x_0) > \frac{\gamma}{\theta}$ “Intervention” is currently implemented. Optimal stopping problem.	“Intervention”: Pap smear for early identification of pre-cancerous lesions on the cervix from HPV infection. Period reward function: INMB _t = $-\hat{\theta}\tilde{q}_t + \hat{\gamma}$; \tilde{q}_t , prevalence of HPV vaccine coverage in cohort t ; $\hat{\theta}$, difference in benefit of Pap smear in a vaccinated person (compared to an unvaccinated person); $\hat{\gamma}$, value of Pap smear in an unvaccinated person. Problem transformation to framework: $\tilde{p}_t = 1 - \tilde{q}_t, \theta = \hat{\theta}, \gamma = -\hat{\gamma} + \hat{\theta}$.
D	\tilde{q}_t , increasing in t	INMB _t is increasing in \tilde{q}_t (decreasing in \tilde{p}_t), $\theta, \gamma < 0$	$\mu_q(x_0) < 1 - \frac{\gamma}{\theta} \Rightarrow \mu_p(x_0) > \frac{\gamma}{\theta}$ “Intervention” not currently implemented. Optimal starting problem.	“Intervention”: Peanut-free spaces regulation (in schools, airplanes, etc.). Period reward function: INMB _t = $\hat{\theta}\tilde{q}_t - \hat{\gamma}$; \tilde{q}_t , prevalence of severe peanut allergy at time t ; $\hat{\theta}$, benefit of peanut-free spaces to individuals with peanut allergies; $\hat{\gamma}$, fixed cost of creating and enforcing peanut-free public spaces. Problem transformation to framework: $\tilde{p}_t = 1 - \tilde{q}_t, \theta = -\hat{\theta}, \gamma = \hat{\gamma} - \hat{\theta}$.

^{*}When $\theta \leq 0 \leq \gamma$, the “Intervention” is dominated by the alternative for all realizations of \tilde{p}_t . For $\gamma \leq 0 \leq \theta$, the “Intervention” dominates the alternative for all realizations of \tilde{p}_t

[†] $\mu_p(x_0)$ is the expectation of the initial belief \tilde{p}_0 ; $\mu_q(x_0)$ is the expectation of the initial belief \tilde{q}_0 ; $\mu_p(x_0) = 1 - \mu_q(x_0)$

[‡]In each of the examples, the period reward function is linear in the time-varying parameter. That the mean and variance of the time-varying parameter satisfy the dynamics presented in Section 2.3 should be verified empirically for each case

ship to Type I and Type II error: ‘Specificity’ = $1 -$ ‘Type I error’ = $1 -$ ‘False positive rate’ and ‘Sensitivity’ = $1 -$ ‘Type II error’ = $1 -$ ‘False negative rate’. The number of positive samples is an uncertainty \tilde{v}_t with realization $v_t \in \{0, \dots, n_t\}$. Based on the collected information the policy-maker updates his beliefs about \tilde{p}_t in a Bayesian manner.

Proposition 1 *If the policy-maker’s prior belief $f_p(\cdot)$ is a mixture of beta-densities, i.e., $f_p \in \mathcal{P}$, then for any number of positive observations $\tilde{v}_t = v_t$ from n_t samples, the Bayesian posterior belief $f_{p|v}(\cdot|v_t)$ is also a mixture of beta-densities, i.e., $f_{p|v}$ is in \mathcal{P} .*

Proof See Appendix A1. □

If \tilde{p}_t is a mixture of $m \geq 1$ beta-densities and if the information-collection technology is imperfect (i.e., $\min\{q_1, q_2\} < 1$), then the true posterior distribution is also a mixture of beta-densities, containing between $m + n_t$ and $m \times (n_t + 1)$ unique beta-distributions (see Appendix A.2.1). The resulting probability density function (pdf) is

$$f_{p|v}(p|x_t, n_t, v_t, q) = \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \sum_{i=1}^m \omega'_{j,k,i} \times \text{beta}(a_{t,i} + j + k, b_{t,i} + n_t - j - k), \quad (2)$$

with updated weights

$$\omega'_{j,k,i} = \frac{\omega_i q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \Gamma(a_{t,i}+b_{t,i}) \Gamma(a_{t,i}+j+k) \Gamma(b_{t,i}+n_t-j-k)}{\Gamma(j+1) \Gamma(v_t-j+1) \Gamma(k+1) \Gamma(n_t-v_t-k+1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i}+b_{t,i}+n_t)}$$

$$\sum_{r=0}^{v_t} \sum_{s=0}^{n_t-v_t} \sum_{i=1}^m \frac{\omega_i q_1^r (1-q_2)^{v_t-r} (1-q_1)^s q_2^{n_t-v_t-s} \Gamma(a_{t,i}+b_{t,i}) \Gamma(a_{t,i}+r+s) \Gamma(b_{t,i}+n_t-r-s)}{\Gamma(r+1) \Gamma(v_t-r+1) \Gamma(s+1) \Gamma(n_t-v_t-s+1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i}+b_{t,i}+n_t)}$$

Explicit expressions for the conditional mean and variance, $\mu_{p|v}$ and $\sigma_{p|v}^2$, are provided in Appendix A.2.2.

Remark 1 If \tilde{p}_t follows a mixture of $m \geq 1$ beta-densities and the information-collection technology is perfect (i.e., $q_1 = q_2 = 1$), then the distribution of sample information, \tilde{v}_t , is a mixture of m beta-binomial distributions with the same weights ω_i . Updating results in a posterior distribution that is a mixture of m beta-densities with pdf

$$f_{p|v}(p|x_t, n_t, v_t, q=(1, 1)) = \sum_{i=1}^m \omega_i \text{beta}(a_{t,i}+v_t, b_{t,i}+n_t-v_t), \quad (3)$$

with updated weights

$$\omega'_i = \frac{\omega_i \frac{\Gamma(a_{t,i}+b_{t,i}) \Gamma(a_{t,i}+v_t) \Gamma(b_{t,i}+n_t-v_t)}{\Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i}+b_{t,i}+n_t)}}{\sum_{j=1}^m \omega_j \frac{\Gamma(a_{t,j}+b_{t,j}) \Gamma(a_{t,j}+v_t) \Gamma(b_{t,j}+n_t-v_t)}{\Gamma(a_{t,j}) \Gamma(b_{t,j}) \Gamma(a_{t,j}+b_{t,j}+n_t)}}$$

for all $i \in \{1, \dots, m\}$.

2.2 Approximate Bayesian inference

For practically relevant sample sizes n_t and an imperfect information-collection technology, the number of beta-densities in the posterior distribution can become very large, thus requiring approximation. The need for distributional approximations in decision models has been recognized by Smith who proposed moment matching to replace continuous distributions by appropriate discrete ones [63]. More recently, moment-matching methods have also been used in a Markovian setting, to approximate vector-autoregressions [64]. In our Markov dynamic programming setting, we apply moment matching to approximate the exact posterior distribution which is a mixture of beta-densities with a single beta-distribution. This greatly simplifies the belief propagation compared to dealing with mixtures of beta-densities which feature an increasingly large number of coefficients with each information-collection effort and ultimately an infinite-dimensional state space.

Thus, instead of carrying forward full distribution information about the posterior mixture of beta-densities caused by an imperfect information-collection technology, the policy-maker's posterior belief about \tilde{p}_t is approximated by a single beta-distribution with the same mean and variance as the exact posterior distribution. The policy-maker's prior belief is represented by the distribution parameters $x_t = (a_t, b_t)$ and the posterior belief incorporating any information collected at time t is represented by the updated parameters $\hat{x}_t = (\hat{a}_t, \hat{b}_t)$. Using the mean and variance of the exact posterior distribution, $\mu_{p|v}$ and $\sigma_{p|v}^2$, the approximate posterior belief parameters are determined using the one-to-one relationship between the standard parameters of the beta-distribution and its mean and variance¹. We let $\psi(x_t, n_t, v_t, q)$ denote the function that generates the approximating parameters, with

$$\hat{x}_t = \begin{bmatrix} \hat{a}_t \\ \hat{b}_t \end{bmatrix} = \psi(x_t, n_t, v_t, q) = \left(\frac{\mu_{p|v}(1-\mu_{p|v})}{\sigma_{p|v}^2} - 1 \right) \begin{bmatrix} \mu_{p|v} \\ 1-\mu_{p|v} \end{bmatrix} \Big|_{(x_t, n_t, v_t, q)}$$

In the case of a perfect information-collection technology, the preceding relations describe the policy-maker's posterior beliefs exactly.

Mixtures of beta-distributions can be fitted to any continuous distribution on $[0, 1]$. Thus, a single beta-distribution with the same mean and variance as a distribution formed by the mixture of beta-densities, will not always provide a satisfactory approximation. However, we focus on the special case where the time- t belief \tilde{p}_t has been obtained via Bayesian updating from a single beta-prior. In this special case, approximating the mixture of beta-densities with a single beta-distribution with the same mean and variance

¹A beta-distribution with parameters (a, b) has mean $\mu = \frac{a}{a+b}$ and variance $\sigma^2 = \frac{ab}{(a+b)^2(a+b+1)}$. Through direct substitution and rearrangement, it can be shown that a beta-distribution with mean μ and variance σ^2 has parameters $a = \mu \left(\frac{\mu(1-\mu)}{\sigma^2} - 1 \right)$ and $b = (1 - \mu) \left(\frac{\mu(1-\mu)}{\sigma^2} - 1 \right)$.

maintains unimodality² and stationarity of the state space over time.

We assessed the approximation quality using simulation in the policy-relevant region for our application (Appendix A.3). We found that the maximum distance between the cumulative density function of the exact posterior distributions and that of the approximation with matching mean and variance were generally small (< 2%), but became large when the mean was approaching zero and the standard deviation was relatively large. The quality of the approximation was very good (< 0.5%) when the mean was greater than 2%. We deemed the approximation to be of sufficiently high quality for our numerical analysis because our initial conditions and predicted trajectory without information acquisition rely on the regions in which the approximation is good. Also, because of relatively high fixed costs associated with information acquisition, optimal sample sizes in our numerical analysis tended to be sufficiently large that information would likely only be collected once which reduces concerns about compounding the approximation error over successive information-collection efforts.

2.3 System dynamics

The belief state x_t , containing the parameters of the distribution of \tilde{p}_t , represents the policy-maker’s current beliefs about the uncertain parameter and follows a law of motion of the form

$$x_{t+1} = \phi(\hat{x}_t) = \begin{bmatrix} z & 0 \\ 1 - z & 1 \end{bmatrix} \hat{x}_t, \tag{4}$$

²The posterior distribution is a weighted sum of component beta-distributions; see Eq. 2. The weights of the exact posterior distribution are generated by the convolution of two binomial distributions; see Eq. 13. The convolution of two binomial distributions creates unimodal weights over the ordered set of component beta-distributions in the mixture (ordered in terms of increasing first parameter). For example, consider a prior $x_t = (3, 7)$. Given a sample size $n_t = 5$, there are 6 possible true outcomes, by which we mean the potentially unobservable number of actual positive samples in the study. These true outcomes correspond to 6 possible unique beta-distributions with parameters (3, 12), (4, 11), (5, 10), (6, 9), (7, 8), and (8, 7) forming the components of the exact posterior distribution. Because the information technology is imperfect, we have a belief over these possible outcomes equal to the distribution of the actual number of positives given a specific number of observed positives. We can compute this distribution using the weights in Eq. 2, where j is the number of actual positives among the v_t observed positives, and k is the number of actual positives among the $n_t - v_t$ observed negatives. The probability that the actual number of positives in the sample is W is determined by summing all the weights for which $j + k = W$. The unimodality of the weights over the ordered component distributions ensures the unimodality of the posterior distribution. To complete the numerical example, consider the case where there are 3 observed positives, $v_t = 3$, and $q = (0.9, 0.85)$, which results in weights of 0.028, 0.141, 0.346, 0.414, 0.067, 0.004 over the component beta-distributions in the mixture.

where $z \in (0, 1)$ is the decay rate. These dynamics imply a geometrically decreasing expected value, increasing coefficient of variation, and decreasing variance for $\mu(x_0) \leq \frac{1}{1+z}$ (Fig. 1). In the mean-variance space, the equivalent state dynamics become

$$\begin{aligned} \mu(x_{t+1}) &= z\mu(\hat{x}_t) \text{ and } \sigma^2(x_{t+1}) \\ &= \sigma^2(\hat{x}_t) \left(z + z(1 - z) \left(\frac{\mu(\hat{x}_t)}{1 - \mu(\hat{x}_t)} \right) \right). \end{aligned}$$

Derivations of these equations are presented in Appendix A.4. The features of these dynamics can represent a wide variety of settings in which the expectation of a parameter is geometrically decreasing over time (e.g., a health condition that is decreasing in prevalence over time; see Section 4). To model a situation where the expectation (and variance) of the uncertain parameter is increasing (e.g., obesity prevalence), the problem can be re-formulated as one in which a parameter of opposite definition is decreasing (e.g., prevalence of individuals who are not obese).

2.4 The policy-maker’s problem

Given a social discount factor $\delta \in (0, 1)$, the policy-maker’s objective is to maximize the net present value of the stream of expected INMBs, given the initial belief $x_0 = (a_0, b_0)$ and admissible policy decisions $U \in \mathcal{U} = \{(u_t)_{t \in \mathbb{N}} : u_t = (d_t, n_t) \in \mathcal{D} \times \mathcal{N}\}$. To achieve the objective, the policy-maker seeks to find the best of all possible policies $\pi_t(\cdot)$, $t \geq 0$, with $u_t = \pi_t(x_t)$ for all $x_t \in \mathbb{R}_{++}^2$, which at each time t maps the state space to admissible current-period actions u_t , so that the implemented path of actions $U = (u_0, u_1, \dots)$ lies in the control-constraint




Fig. 1 Sample state trajectory with decay $z = 0.8$, with (dashed) and without (solid) information acquisition at time t (for $n_t = 50$), respectively

set \mathcal{U} . The number of positive observations in the testing sample of n_t is a random variable $\tilde{v}_t(n_t)$ with realization $v_t \in \{0, \dots, n_t\}$. Based on the collected information the policy-maker updates his beliefs about \tilde{p}_t in an (approximate) Bayesian manner using the function $\psi(x_t, n_t, v_t, q)$. Because of the decreasing trend of the uncertain parameter ($z < 1$), it is never optimal to restart an optimally stopped program.³ We consider stationary policies $\pi : \mathbb{R}_{++}^2 \rightarrow \mathcal{D} \times \mathcal{N}$ to solve the optimal control problem

$$\begin{aligned} \max_{\pi(\cdot)} \quad & \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t g(\tilde{p}_t, \pi(x_t)) \mid x_0 \right], \\ \text{subject to} \quad & x_{t+1} = \phi(\psi(x_t, n_t, \tilde{v}_t(n_t), q)), \quad x_0 \text{ given,} \\ & u_t = \pi(x_t) \in \mathcal{D} \times \mathcal{N}. \end{aligned} \tag{5}$$

Provided the value function $V(x)$ satisfies the Bellman equation,

$$\begin{aligned} V(x) = \max_{(d,n) \in \mathcal{D} \times \mathcal{N}} & \left\{ d(\theta\mu(x) - \gamma) - \frac{K(n)}{N} \right. \\ & \left. + \delta \mathbb{E}[V(\phi(\psi(x, n, \tilde{v}(n), q))) \right. \\ & \left. \times |x, (d, n)] \right\}, \end{aligned} \tag{6}$$

for all admissible states $x \in \mathbb{R}_{++}^2$, the corresponding maximizer $\pi^*(x)$ on the right-hand side defines an optimal policy.

Remark 2 To reflect the policy-maker’s ongoing concern for the health-intervention decision, the problem is formulated in an infinite-horizon setting. Given a time-invariant system, this implies that the optimal policy can be described as a mapping from states to actions, without explicit consideration of time. If more information about the system becomes available over time, for example, relating to the decay rate in the system dynamics (see Eq. 4), then it is possible for the policy-maker to re-solve the problem and update the policy accordingly.

3 Dynamic healthcare decisions

3.1 Policies without information acquisition

If information is prohibitively costly or practically infeasible to collect, Eq. 6 simplifies to

$$V_{\text{NoInfo}}(x) = \max_{d \in \{0,1\}} \{d(\theta\mu(x) - \gamma) + \delta V_{\text{NoInfo}}(\phi(x))\},$$

³Weber [65] uses global optimization to consider the general problem of switching between arbitrary streams of expected benefits allowing for multiple switches, which can be viewed as a deterministic equivalent of the multi-armed bandit problem.

for all $x \in \mathbb{R}_{++}^2$, as there is no Bayesian updating and therefore ψ reduces to an identity map. For all states x for which the optimal strategy is to not do the intervention, this action remains optimal in the future because of the decreasing trend of \tilde{p}_t . Indeed, since for $z \in (0, 1)$, $\mu(\phi(x)) = z\mu(x) < \mu(x)$, we have that for all states where $V(x) = 0$, it is also the case that $V(\phi(x)) = 0$. Hence, for $\mu(x) \leq \frac{\gamma}{\theta}$ it is optimal to stop the intervention. This defines a threshold policy of the form

$$d^* = \begin{cases} 0, & \text{if } \mu(x_t) \leq \frac{\gamma}{\theta}, \\ 1, & \text{otherwise,} \end{cases} \tag{7}$$

for all $t \geq 0$. Restricting attention to the interesting case where $\mu(x_0) \geq \frac{\gamma}{\theta}$ and using the fact that $\mu(x_t) = z^t \mu(x_0)$, we can identify the optimal time $T(x_0)$ to stop the intervention, which is the first period in which the intervention has a nonpositive expected INMB (see Appendix A.5):

$$T(x_0) = \left\lceil \frac{1}{\ln(z)} \ln \left(\frac{\gamma}{\theta \mu(x_0)} \right) \right\rceil. \tag{8}$$

Hence, given any initial state x , the value of implementing the optimal stopping policy for $t \in \{0, \dots, T(x) - 1\}$ is

$$\begin{aligned} V_{\text{NoInfo}}(x) &= \sum_{t=0}^{T(x)-1} \delta^t (\theta z^t \mu(x) - \gamma) \\ &= \theta \mu(x) \left(\frac{1 - (\delta z)^{T(x)}}{1 - \delta z} \right) - \gamma \left(\frac{1 - \delta^{T(x)}}{1 - \delta} \right). \end{aligned} \tag{9}$$

Proposition 2 *When information is prohibitively costly or practically infeasible to collect, the optimal value function $V_{\text{NoInfo}}(x_t)$ is non-decreasing and convex in $\mu(x_t)$.*

Proof See Appendix A.6. □

Remark 3 The above result depends only on the decay in the mean of the uncertain parameter distribution and is otherwise distribution-free. In other words, it does not depend on the policy-maker’s beliefs other than that \tilde{p}_t is expected to decrease over time.

3.2 Policies with information acquisition

When the policy-maker has the option to acquire information, the value function is determined by the Bellman equation (Eq. 6). Its properties in the no-information case (Proposition 2) carry over to the more general situation.

Proposition 3 *The optimal value function $V(x_t)$ is non-decreasing and convex in $\mu(x_t)$, and nondecreasing in $\sigma^2(x_t)$.*

Proof See Appendix A.7. □

3.2.1 Special case: one-time information collection

Assume for now that information can be collected at most once. Given a one-time size- η experiment (with $\eta \geq 1$) and, briefly, ignoring the cost of information collection $K(\eta)$, the value with information exceeds the no-information value,

$$\theta\mu(x) - \gamma + \delta\mathbb{E}[V_{\text{NoInfo}}(\phi(\psi(x, n_0 = \eta, \tilde{v}(\eta), q)))] - V_{\text{NoInfo}}(x) > 0,$$

as a consequence of Jensen’s inequality. This insight is also useful for the comparison of experiments. A higher confidence in the information, i.e., for a larger sample size and/or better test characteristics, produces a mean-preserving spread of the random-variable $\mu(\phi(\psi(x_t, n_0 = \eta, \tilde{v}_t(\eta), q)))$ in the original experiment, and thus, by the convexity of the (monotone) value function and second-order stochastic dominance, a larger value with information.

Because of the monotone system dynamics, the optimal time to collect information of sample size η , at cost $\kappa(\eta)$, is obtained by finding a period k where information acquisition is preferred to waiting until the next period, $k + 1$. In other words, find the smallest k for which

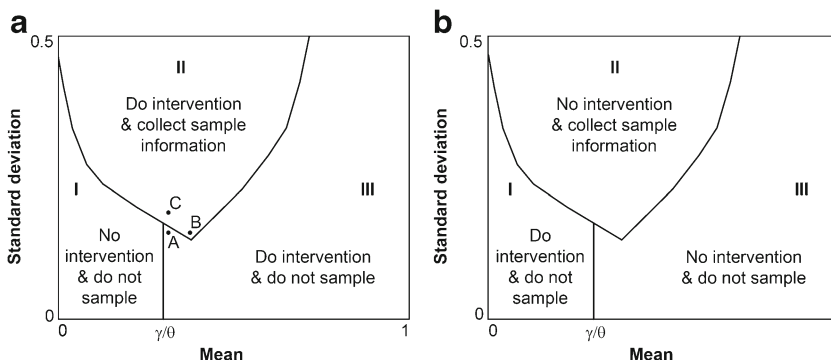
$$V(x_k | n_k = \eta) \geq V(x_{k+1} = \phi(x_k) | n_{k+1} = \eta),$$

or equivalently

$$\begin{aligned} &\theta z^k \mu(x_0) + \frac{\delta}{1 - \delta z} \\ &\times (\mathbb{E}[V_{\text{NoInfo}}(\phi(\psi(x_k, n_k = \eta, \tilde{v}(\eta), q)))] \\ &- \delta\mathbb{E}[V_{\text{NoInfo}}(\phi(\psi(\phi(x_k), n_{k+1} = \eta, \tilde{v}(\eta), q)))] \\ &\geq \frac{1 - \delta}{1 - \delta z} (\gamma + \kappa(\eta)). \end{aligned}$$

The positivity of the right-hand side of the last inequality indicates that information acquisition may, on certain trajectories, never be optimal. This is confirmed in our application in Section 4, where the stopping region and the region with information acquisition have a common boundary, transversal to expected state trajectories.

Fig. 2 Policy regions for **a** an optimal-stopping problem and **b** an optimal-starting problem. In either case, the initial belief is in region II or III; over time, the belief moves towards region I



3.2.2 General case: information collection in any period

Based on Proposition 3, the intervention is desirable for greater $\mu(x_t)$ and greater $\sigma(x_t)$; the latter increases the upside of the policy-maker’s asymmetric (convex) payoffs, as if holding a call option. The dynamics presented in Eq. 4, with decreasing expectation and decreasing variance, imply monotonicity of the intervention decision, $d_{t+1} \leq d_t$.

Corollary 1 Consider $x_t^{(1)}, x_t^{(2)}$ with $\mu(x_t^{(1)}) < \mu(x_t^{(2)})$ and $\sigma^2(x_t^{(1)}) = \sigma^2(x_t^{(2)})$, then if it is optimal to do the intervention with $\mu(x_t^{(1)})$, it is also optimal to do the intervention with $\mu(x_t^{(2)})$.

Proof See Appendix A.8. □

Corollary 2 Consider $x_t^{(1)}, x_t^{(2)}$ with $\mu(x_t^{(1)}) = \mu(x_t^{(2)})$ and $\sigma^2(x_t^{(1)}) < \sigma^2(x_t^{(2)})$, then if it is optimal to do the intervention with $\sigma(x_t^{(1)})$, it is also optimal to do the intervention with $\sigma(x_t^{(2)})$.

Proof See Appendix A.9. □

A direct consequence of Proposition 3 and Corollaries 1 and 2 is that an optimal policy, as a map from states to actions, features three regions (Fig. 2). We describe, in detail, the features of the optimal policy for the case of an optimal stopping problem (Fig. 2A). In region I, an optimal policy is ‘no intervention (and do not sample).’ In region II, an optimal policy is ‘do intervention and sample n_t individuals.’ In region III, an optimal policy is to ‘do intervention and do not sample.’

The boundary between regions I and III is $\frac{\gamma}{\theta}$ (Section 3.1). For $0 \leq \mu(x_t) \leq \frac{\gamma}{\theta}$, the policy-maker is indifferent between ‘no intervention (and do not sample)’ and ‘do intervention and sample n_t individuals’ when the rewards of the two regions are equal:

$$0 = \theta\mu(x_t) - \gamma - \kappa(n_t) + \delta\mathbb{E}[V(\phi(\psi(x_t, n_t, \tilde{v}_t, q)))] \tag{10}$$

Focusing on the region $\frac{\gamma}{\theta} \leq \mu(x_t) \leq 1$, the policy-maker is indifferent between ‘do intervention and sample

n_t individuals' and 'do intervention and do not sample' when the rewards of the two regions are equal. Removing common terms from each side, this occurs when

$$-\kappa(n_t) + \delta \mathbb{E}[V(\phi(\psi(x_t, n_t, \tilde{v}_t, q)))] = \delta V(\phi(x_t)). \quad (11)$$

For each $\sigma^2(x_t)$, there can exist more than one $\mu(x_t)$ where $\frac{\gamma}{\theta} < \mu(x_t) \leq 1$ satisfying Eq. 11 because $V(\phi(\psi(x_t, n_t, \tilde{v}_t, q)))$ is increasing, but neither concave or convex, in v_t . The existence of the section of region III between regions I and II (the location of point A) can be obtained using intuition. Consider two points, A and B, with the same standard deviation (Fig. 2). Compared to point B, if information were to be gathered at point A, the distribution of possible posterior states includes a higher proportion of states in region I (with a reward of 0) and a lower proportion of high-reward states (those with high mean and high standard deviation) and, therefore, information acquisition is less likely to yield a value exceeding its cost. Now consider two points, A and C, with the same mean. Compared to point C, if information were to be gathered at point A, the distribution of possible posterior states is narrower. In both of these cases, increased spread on the side of low mean has no impact on the expectation and increased spread into the high-reward states substantially increases expectation. Therefore, information acquisition is more likely to yield a value exceeding its cost for the state with higher standard deviation.

Proposition 4 For a fixed sample size η (so $n_t \in \{0, \eta\}$ for all t), misclassification in the information-collection technology decreases the value function and reduces the number of states for which information acquisition is optimal.

Proof See Appendix A.10. \square

This result is consistent with Blackwell's result that a less informative signal cannot increase the value of a single-person decision problem [66].

4 Application

4.1 Background and motivation

Chronic HCV infection is a slowly progressing blood-borne disease that causes liver fibrosis, cirrhosis, and liver cancer. It is the principal cause of death from liver disease and the leading indication for liver transplantation in the United States (US) [67, 68]. Between 2.7 and 5.2 million Americans (1.1% to 2.1% of the adult population) are chronically infected with HCV [69, 70]. In the non-injection drug using US population, prevalence peaks in the 1945 to 1965 birth cohorts and decreases thereafter (Fig. 3). Approximately half of all chronically infected individuals are unaware of their disease status [71].

Recent model-based analyses concluded that one-time screening of individuals born between 1945 and 1965 is cost-effective [11–15] and the CDC and USPSTF recently released new guidance in support of one-time screening of these birth cohorts [16, 17]. Several studies indicate that screening individuals born later than 1965 is also likely to be cost-effective [13–15]. Since HCV prevalence is decreasing in birth year after the 1956 birth cohort (Fig. 3), there may be a time at which screening is no longer cost-effective. To improve the decision about the best time to stop screening, additional information about prevalence of the current and future cohorts may be desirable. However, standard approaches to finding the value of information do not usually include the option to delay the information acquisition.

Note that the population we model were predominantly infected decades ago [72, 73] and do not have ongoing risk factors for HCV re-infection. Many historically significant modes of disease incidence have been virtually eliminated including transmission by surgical or other hospital equipment prior to modern sterilization procedures and blood transfusion [73, 74]. Injection equipment sharing among people who actively use injection drugs (PWID) is currently the principal cause of HCV transmission [76]. Although a history of injection drug use is relatively common among individuals with chronic HCV infection (approximately

Fig. 3 Prevalence of HCV by birth year in **a** men and **b** women. Estimated using the National Health and Nutrition Examination Survey (NHANES) (1999–2010). For details, see Appendix A.11.3




Table 2 Model parameter values and range used in sensitivity analysis

Variable, Description	Males		Females		Sources
	Value	(Range)	Value	(Range)	
<i>Annual Cohort: Individuals aged 50 years</i>					
Eligible for a preventive health exam (PHE)	2.1 million	(1.8 – 2.3 million)	2.1 million	(1.9 – 2.4 million)	[79]
Proportion who attend a PHE	24.4%	(19.3 – 29.5%)	43.3%	(37.3 – 49.3%)	[80]
<i>N</i> Annual number of PHE	508, 222	(386, 600 – 630, 000)	920, 706	(753, 000 – 1, 100, 000)	Calculated
<i>HCV Screening Test</i>					
<i>q</i> ₁ Sensitivity	0.97	(0.950 – 0.999)			[81]
<i>q</i> ₂ Specificity	0.9996	(0.990 – 1.0)			[82]
<i>C</i> _S Cost of HCV-antibody test	\$28	(\$20 – 40)			[83]
<i>C</i> _{FP} Cost of false positive	\$230	(\$200 – 250)			[83]
<i>B</i> _S Quality-of-life change, event of screening	0	(not varied)			Assumed
<i>B</i> _{FP} Quality-of-life change, false-positive result	0	(not varied)			Assumed
<i>Lifetime discounted costs per person</i>					
<i>C</i> ₁ HCV+, identified through screening	\$146, 928	(\$140, 000 – 154, 000)	\$161, 121	(\$153, 000 – 170, 000)	[14]
<i>C</i> ₂ HCV+, not identified through screening	\$126, 943	(\$120, 000 – 133, 000)	\$143, 900	(\$136, 000 – 152, 000)	[14]
<i>C</i> ₃ HCV- individual	\$181, 314	(\$172, 250 – 190, 000)	\$192, 135	(\$182, 500 – 202, 000)	[14]
<i>Lifetime discounted quality-adjusted life-years (QALYs) per person</i>					
<i>B</i> ₁ HCV+, identified through screening	10.42	(10.16 – 10.68)	11.71	(11.41 – 12.0)	[14]
<i>B</i> ₂ HCV+, not identified through screening	10.06	(9.80 – 10.31)	11.44	(11.15 – 11.72)	[14]
<i>B</i> ₃ HCV- individual	15.69	(15.30 – 16.08)	16.24	(15.83 – 16.64)	[14]
<i>Incremental net monetary benefit (INMB) per person</i>					
θ Variable component of INMB	\$7, 030	(\$4, 000 – 12, 000)	\$2, 962	(\$3, 000 – 10, 000)	Eq. 20
γ Fixed component of INMB	\$28.05	(\$22 – 40)	\$28.05	(\$22 – 40)	Eq. 21
<i>Cost of collecting information</i>					
<i>K</i> _F Fixed cost per sampling study	\$50, 000	(\$25, 000 – 250, 000)			Estimated
<i>K</i> _V Variable cost per sample	\$100	(\$50 – 500)			[83]
<i>K</i> (<i>n</i> _{<i>t</i>}) Cost of sampling	$(K_F + K_V n_t) \mathbf{1}_{\{n>0\}}$				Assumed
<i>Other</i>					
<i>x</i> ₀ Initial belief, HCV prevalence in undiagnosed individuals (1960 birth cohort in 2010)	$a_0 = 75.1, b_0 = 2350.5$		$a_0 = 51.5, b_0 = 3756.3$		Appendix A.11.3
<i>z</i> Rate of prevalence decay	$\mu(x_0) = 0.031, \sigma(x_0) = 0.0035$		$\mu(x_0) = 0.0135, \sigma(x_0) = 0.0019$		Appendix A.11.3
λ Willingness-to-pay threshold	0.893	(0.871 – 0.915)			[84]
<i>r</i> Annual discount rate	\$75, 000/QALY gained	(\$50, 000 – 100, 000/QALY gained)			[1]
	0.03	(0 – 0.05)			[1]

PHE – Preventive health exam; QALY – Quality-adjusted life-year; INMB – Incremental net monetary benefit

40% [71]), re-infection and disease transmission to others via injection drug use are not an ongoing risk for a large proportion of these individuals as three-quarters of HCV infected individuals with a self-reported history of injection drug use report last injecting greater than 5 years ago (median time since last injection = 20 years) [75]. Our model does not include PWID and so we do not consider the possibility of re-infection. PWID are a high-risk population and guidelines, separate from those otherwise discussed here, recommend routine annual HCV screening in this population [77].

We now apply the stochastic dynamic programming framework developed in Section 2 to the case of one-time HCV screening at a routine medical appointment at age 50 for successive birth cohorts. We consider screening at age 50 because one-time screening at this age had the lowest incremental cost-effectiveness ratio in an analysis of single birth cohort screening [14]. Waiting to perform a one-time screening in older individuals is less cost-effective because their disease may have progressed further and treatment is less effective in more severe disease states. One-time screening of younger individuals is less cost-effective because younger individuals are further away from the long-term consequences of HCV which screening and treatment hope to avoid. We transform the unbounded state space in terms of $x_t = (a_t, b_t)$ to the compact policy-relevant space $\mu(x_t)$ and $\sigma(x_t)$. Using value iteration implemented in R version 2.15.0 [78], we numerically determine an optimal HCV-screening and information-collection policy for US adults.

At each time, we consider the actions of ‘do not screen for HCV and do not collect information about HCV prevalence in the current cohort;’ ‘screen for HCV and collect sample information about HCV prevalence in the current cohort;’ ‘screen for HCV and do not collect information about HCV prevalence in the current cohort.’ We compare this optimal strategy to the policies identified by various alternative approaches: a slightly modified version of the new CDC and USPSTF recommendation; an optimal policy without information acquisition; and an optimal policy with (possibly immediate) information acquisition. A policy of HCV screening does not inherently provide additional information about HCV prevalence to policy-makers, because only positive test outcomes are reported to the CDC and the reason for the medical test is private health information (the test may have been performed for a reason other than routine screening at age 50). Estimating prevalence among asymptomatic individuals seeking routine preventive medical care therefore requires a study with random sampling of those individuals. The (quasi-)linearity of $INMB_t$ for this example is established in Appendix A.11.1. Parameter values and ranges used in sensitivity analysis are presented in Table 2. Details of parameter estimation are presented in Appendix A.11.2 and A.11.3.

Table 3 Comparison of optimal policies indicated by various analytic approaches for men with initial belief $\mu(x_0) = 0.0310$ and $\sigma(x_0) = 0.0035$ and women with initial belief $\mu(x_0) = 0.0135$ and $\sigma(x_0) = 0.0019$

Case	Optimal Policy	Value (Expected INMB)	Increase in Expected INMB
Males			
CDC/USPSTF recommendation*	Screen until 1965 birth cohort turns 50	\$399, 140, 000	Reference
No information available	Screen until 1978 birth cohort turns 50	\$566, 470, 000	\$167, 330, 000
Information only available immediately	Sample 910 men now, then identify optimal action	\$566, 490, 000	\$167, 350, 000
Information available in all periods	Sample 4,000 men in 16 years (1976 birth cohort), then identify optimal action	\$567, 940, 000	\$168, 800, 000
Females			
CDC/USPSTF recommendation*	Screen until 1965 birth cohort turns 50	\$15, 390, 000	Reference
No information available	Screen until 1963 birth cohort turns 50	\$21, 720, 000	\$6, 330, 000
Information only available immediately	Sample 4,930 women now, then identify optimal action	\$22, 320, 000	\$6, 930, 000
Information available in all periods	Sample 4,500 in 1 year (1961 birth cohort), then identify optimal action	\$22, 500, 000	\$7, 110, 000

*The new guidelines recommend screening all individuals born between 1945 and 1965 for HCV at their next routine medical visit [16, 17]. We ignore the screening of individuals born prior to 1960; for all others, we assume HCV screening occurs at age 50

4.2 Results

For the purposes of our analysis, we assume the current time to be the year 2010 and the initial cohort to be born in 1960.

4.2.1 Policies identified by alternative approaches

The expected value of the CDC and USPSTF recommendation was obtained by substituting $T = 6$ into Eq. 9. The sum of the discounted expected INMBs for screening 6 cohorts at age 50, until the 1965 birth cohort turns 50 years of age, is \$399.1 million for men and \$15.4 million for women (Table 3). The large difference between men and women is attributable to higher HCV prevalence and higher marginal INMB of early diagnosis and treatment in men.

We identify the threshold prevalence value below which the HCV-screening program should be terminated and the best time to terminate the screening program, assuming no opportunity to collect information using Eqs. 7–8. In men, the program should be terminated when prevalence falls below 0.4%, which will occur in 18 years (95%CI: 16-19 years). In women, the program should be terminated when prevalence falls below 0.1%, which will occur in 3 years (95%CI: 0-5 years). The expected INMB of these policies is \$566.5 million for men and \$21.7 million for women (Table 3).

The traditional approach to value-of-information assessment in the health policy literature assumes immediate information collection [3, 4]. For men and women, we find the optimal sample sizes to be 910 and 4,930 individuals from the current cohort, respectively (Fig. 4). The expected INMB of immediate information followed by the optimal policy based on the information collected increases by \$20,000 for men and \$600,000 for women. Women have a greater value of immediate information because they are closer to the intervention stopping region threshold and,

therefore, immediate information is more likely to result in a policy change.

4.2.2 Model results

Implementing the full model, we considered the possibility of collecting sample information at each decision period. For computational and illustrative reasons, we restricted the policy-maker’s choice to two sample sizes $\mathcal{N} \in \{0, \eta\}$. We considered several possible values for η (2000, 2500, 3000, ..., 8000) and we present the results for the sample size that maximized the value at the initial condition for each gender. We also performed analyses using multiple study sample size levels available at each period. We do not present these analyses, as they led to the same optimal policies indicating that our restriction to two sample sizes was not material for this application.

The optimal policy is characterized by the three main regions described in Section 3.2.2 (Fig. 5a). At low prevalence and relatively low uncertainty, it is optimal to not screen and not collect information. At high prevalence, it is optimal to screen and not collect information. At prevalence close to the $\frac{\gamma}{\theta}$ threshold and relatively high uncertainty, it is optimal to both screen and collect information.

For each state in the region where it is optimal to screen without information acquisition, we can identify the optimal next action and the time when it should occur (Fig. 5b). We subdivide this region by a solid line. Above the solid line, which is the region with higher uncertainty, it is optimal to screen without information acquisition for a specified number of periods and then to collect information. In the region with lower uncertainty, it is optimal to screen without information acquisition for a specified number of periods and then to stop screening without ever collecting information. The current prevalence estimates for men

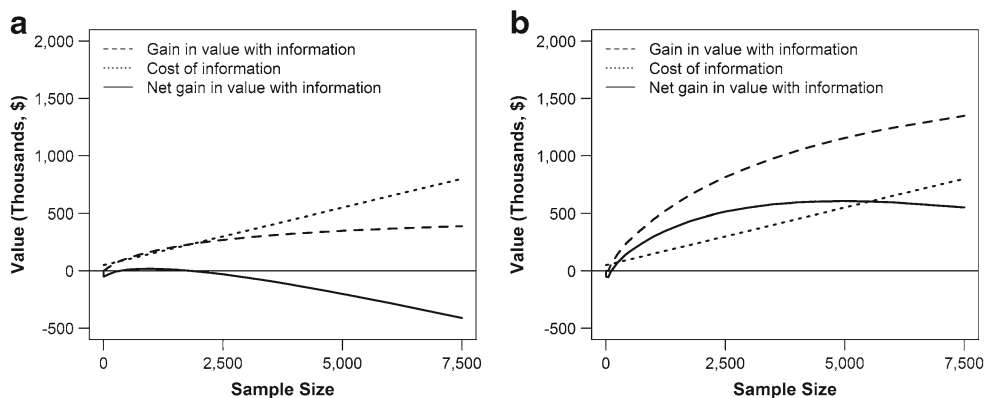


Fig. 4 The value of collecting sample information immediately for various sample sizes. The gain in the expected INMB of the policy $(\theta\mu(x_0) - \gamma + \delta E[V_{\text{NoInfo}}(\phi(\psi(x_0, n_0 = \eta, \tilde{v}(\eta), q)))])$, the cost of information $(\kappa(\eta))$, and the net gain in the expected INMB of an

HCV-screening policy if sample information is collected in the first period only as a function of sample size for **a** men with initial belief $\mu(x_0) = 0.031$ and $\sigma(x_0) = 0.0035$ and **b** women with initial belief $\mu(x_0) = 0.0135$ and $\sigma(x_0) = 0.0019$

and women indicate that it is optimal to screen without information collection for 16 years and 1 year, respectively, and then to collect sample information to inform the next action. The expected INMBs of these policies are \$567.9 million and \$22.5 million for men and women, respectively (Table 3).

For each state, we also computed the marginal value of collecting a specific amount of information (Fig. 5c). The marginal value of information in the current period is near-zero for states in which collecting information in the future is optimal. Consistent with our expectations, in the ‘Screen and Collect Information’ region, the marginal value




Fig. 5 a Optimal policy given any current belief about HCV prevalence and the opportunity to sample 4,000 men (left) and 4,500 women (right) at any time. b Time to the next policy action for men (left) and women (right). For states below the solid line, the next action is to

stop screening. For states above the solid line, the next action is to collect information. c The marginal value of collecting information, 4,000 samples for men (left) and 4,500 samples for women (right), in the current period.

of information is greatest close to the $\frac{\gamma}{\theta}$ -threshold and increases with uncertainty. In the ‘Screen and Do Not Collect Information’ region, the value of information is highest along the boundary that divides the region into points with trajectories leading to information collection and points with trajectories leading to ‘No Screening’ without information collection.

Sensitivity analysis identified that the general conclusions of our numerical analysis are robust to uncertainty in the inputs (details in Appendix A.11.4).

4.3 Discussion of application

Evaluating an HCV-screening policy over its entire lifecycle using a stochastic dynamic programming approach has led to several important policy-relevant insights. Our analysis indicates that recommendations by the CDC and USPSTF to screen individuals born between 1945 and 1965 at their next routine medical visit are conservative for men. Specifically, our analysis shows that, for men, screening should continue until at least the 1976 birth cohort turns 50 (in 2026), at which point 4,000 individuals should be sampled to inform about the continuation of the program. Screening men at least 10 years longer will enable early diagnosis in an estimated 50,500 additional individuals, thus preventing an expected 767 additional liver cancers and about 212 additional liver transplants. For women, we find that a large information-acquisition effort should take place when the 1961 birth cohort turns 50 (in 2011),⁴ as it is likely not cost-effective to screen women, per guidance, to the 1965 cohort because of relatively low prevalence (Fig. 3) and slower disease progression in women [85]. Compared to the CDC and USPSTF recommendation, our model increases the expected INMB by \$168.8 million in men and \$7.1 million in women.

Our analysis has several limitations. First, we assume only the current cohort can be sampled to learn about subsequent cohorts, relying on the correlation between cohorts (as implied by the system dynamics). In practice, for our example, it is possible to sample the next cohort (49-year olds) directly. We chose this assumption because the individuals who make up the ‘next cohort’ are typically unknown (e.g., the next cohort of patients with a heart attack, the next cohort of pregnant women, or the next cohort of cancer patients). Second, we consider one-time screening at age 50 based on a cost effectiveness analysis of once-in-a-lifetime HCV screening [14]. However, this analysis (and, consequently, ours) assumed that the cohort being screened has not been previously screened. Our model does not identify the optimal age at which to perform

one-time screening. Third, we assumed that the individuals who attend a preventive health exam and participate in recommended HCV screening are an unbiased sample from the cohort—that is, individuals are not more or less likely to attend their preventive health exam if they are HCV-positive. However, if individuals at higher-risk of HCV disproportionately self-select for general population screening, then we have underestimated the duration for which screening will be cost-effective. If individuals at lower-risk disproportionately self-select for screening (often called the “worried well”), then we have overestimated the duration for which screening will be cost-effective. Fourth, we focus on HCV screening policy in the non-injection drug using population only because they were the focus of the recent change in HCV screening policy. Finally, while uncertainty (and related information acquisition) with respect to model parameters other than prevalence can be treated in an analogous manner, the details are left for future work.

5 Conclusion

Our analysis shows that when parameters vary across intervention cohorts, it may be optimal to delay information acquisition. This is a significant improvement over the current paradigm which only considers one-time immediate information collection. More specifically, we provide a framework for optimal information acquisition, in terms of timing and precision of the acquired signal (sample size). Further, we incorporate misclassification from an imperfect information-collection technology into our framework, which is an important real-life complexity of information gathering that adds substantial analytical difficulty.

The common assumption that the per-person value-of-information remains constant for future cohorts may result in a significant error when estimating the population value of additional information. It may indicate immediate expensive information collection when, incorporating the system dynamics, the optimal action is to collect information in the future or never at all. When a parameter is evolving across intervention cohorts, ignoring the opportunity to wait and collect information in the future, when the information collected is more likely to result in action, is a missed opportunity for increased efficiency. As seen in our example, adding the option of delaying information acquisition until a time when the signal is more likely to justify a policy shift can increase the expected value compared to a policy of immediate information collection. The dynamic programming framework developed in this paper enables an accurate assessment of the marginal value of additional information and identifies an optimal information-acquisition policy.

In this work, we assumed that the dynamics are monotonically increasing or decreasing and that they are

⁴Our initial cohort is individuals born in 1960. This result can be interpreted as a recommendation for immediate information collection.

deterministic. In future work, we plan to consider the more realistic assumption of uncertainty in the dynamics. This would then enable learning about the evolution of the parameters, rather than just their current state. Furthermore, our model does not consider the possibility of intervening on a cohort at a different time in the course of their disease or lives (i.e., at an earlier or later age) or the possibility of the intervention modifying the population-level dynamics. Although true for our application, this latter assumption does not hold in general for an infectious disease. Including the additional benefits of reduced disease transmission from prevention and treatment interventions may generate more near-term benefits and may dramatically alter the value of the intervention over time.

With strained resources for health programs and population-health monitoring, this type of analysis may ensure an optimal implementation horizon for health programs together with guidance on when and how much information should be collected to inform health-program adjustments. Beyond health, many application areas face limited resources for investment and information acquisition, high-quality decision-relevant information is often difficult or expensive to collect, and population or environmental trends influence the preferences and behavior of customers across industries. Facing a dynamic consumer, competitive, or physical environment, the optimal timing of high-quality information acquisition may provide competitive advantage.

Acknowledgments The authors thank Laurie Barker, MSPH, Division of Viral Hepatitis, Centers for Disease Control and Prevention for assistance with NHANES analysis.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

A glossary of symbols is provided in Table 4.

A.1 Proof of Proposition 1

The special case where \tilde{p}_t is beta-distributed with parameters $x_t = (a_t, b_t)$ and the information-collection technology is perfect, i.e., where $q = (1, 1)$, is well known [53]. The sample information \tilde{v}_t , where \tilde{v}_t is the number of observed positives of n_t samples, is beta-binomially distributed, and updating the prior with the sample information results in a beta-distributed posterior belief with parameters $(a_t + v_t, b_t + n_t - v_t)$.

Table 4 Glossary of Symbols

Symbol	Definition
N	Annual number of preventive health exams
<i>HCV Screening Test</i>	
q_1	Sensitivity
q_2	Specificity
C_S	Cost of HCV-antibody test
C_{FP}	Cost of false positive
B_S	Quality-of-life change, event of screening
B_{FP}	Quality-of-life change, false-positive result
<i>Lifetime discounted costs per person</i>	
C_1	HCV+, identified through screening
C_2	HCV+, not identified through screening
C_3	HCV- individual
<i>Lifetime discounted quality-adjusted life-years (QALYs) per person</i>	
B_1	HCV+, identified through screening
B_2	HCV+, not identified through screening
B_3	HCV- individual
<i>Incremental net monetary benefit (INMB) per person</i>	
θ	Variable component of INMB
γ	Fixed component of INMB
<i>Cost of collecting information</i>	
K_F	Fixed cost per sampling study
K_V	Variable cost per sample
$K(n_t)$	Cost of sampling
<i>Other</i>	
x_0	Initial belief, HCV prevalence in undiagnosed individuals (1960 birth cohort in 2010)
z	Rate of prevalence decay
λ	Willingness-to-pay threshold
r	Annual discount rate

Consider the interesting case with a prior belief $f_p(x_t)$ which is a mixture of $m_t \geq 1$ beta-distributions where $x_t = (x_{t,1}, x_{t,2}, \dots, x_{t,m_t})$ and $x_{t,i} = (a_{t,i}, b_{t,i}) \in \mathbb{R}_{++}^2$ such that $1 \leq i \leq m$, and a set of positive weights ω_i such that $\sum_{i=1}^m \omega_i = 1$. The pdf is

$$f_p(p|x_t) = \sum_{i=1}^m \omega_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i)\Gamma(b_i)} p^{a_i-1} (1-p)^{b_i-1}.$$

At time t the policy-maker chooses $n_t > 0$, indicating that they will collect n_t Bernoulli trials. The information-collection technology, with test sensitivity q_1 and test specificity q_2 , is imperfect. The probability of observing a ‘positive’ signal from any single Bernoulli trial is $\tilde{p}_t q_1 + (1 - \tilde{p}_t)(1 - q_2)$. Therefore,

$$f_{v|p}(v_t|\tilde{p}_t = p, n_t, q) = \binom{n_t}{v_t} [p q_1 + (1-p)(1-q_2)]^{v_t} \times [1 - p q_1 - (1-p)(1-q_2)]^{n_t - v_t}.$$

The resulting distribution of sample information is effectively a weighted beta-binomial distribution, correcting

for the additional uncertainty introduced by the imperfect information-collection technology:

$$\begin{aligned}
 f_v(v_t|x_t, n_t, q) &= \int_0^1 f_{v|p}(v_t|\bar{p}_t = p_t, n_t, q) f_p(\bar{p}_t = p_t|x_t) dp \\
 &= \int_0^1 \binom{n_t}{v_t} [pq_1 + (1-p)(1-q_2)]^{v_t} [1-pq_1 - (1-p)(1-q_2)]^{n_t-v_t} f_p(\bar{p}_t = p_t|x_t) dp \\
 &= \int_0^1 \binom{n_t}{v_t} [pq_1 + (1-p)(1-q_2)]^{v_t} [p(1-q_1) + (1-p)q_2]^{n_t-v_t} \times \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} p^{a_{t,i}-1} (1-p)^{b_{t,i}-1} dp \\
 &= \int_0^1 \binom{n_t}{v_t} \left(\sum_{j=0}^{v_t} \binom{v_t}{j} p^j q_1^j (1-p)^{v_t-j} (1-q_2)^{v_t-j} \right) \left(\sum_{k=0}^{n_t-v_t} \binom{n_t-v_t}{k} p^k (1-q_1)^k (1-p)^{n_t-v_t-k} q_2^{n_t-v_t-k} \right) \\
 &\quad \times \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} p^{a_{t,i}-1} (1-p)^{b_{t,i}-1} dp \\
 &= \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \binom{n_t}{v_t} \binom{v_t}{j} \binom{n_t-v_t}{k} q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \times \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} \int_0^1 p^{a_{t,i}+j+k-1} (1-p)^{b_{t,i}+n_t-j-k-1} dp \\
 &= \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \binom{n_t}{v_t} \binom{v_t}{j} \binom{n_t-v_t}{k} q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \times \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} \frac{\Gamma(a_{t,i} + j + k) \Gamma(b_{t,i} + n_t - j - k)}{\Gamma(a_{t,i} + b_{t,i} + n_t)} \\
 &= \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \times \sum_{i=1}^m \omega_i \frac{\Gamma(n_t + 1) \Gamma(a_{t,i} + b_{t,i}) \Gamma(a_{t,i} + j + k) \Gamma(b_{t,i} + n_t - j - k)}{\Gamma(j + 1) \Gamma(v_t - j + 1) \Gamma(k + 1) \Gamma(n_t - v_t - k + 1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i} + b_{t,i} + n_t)}.
 \end{aligned}
 \tag{12}$$

Using Bayes' Theorem, the posterior distribution is a mixture of beta-distributions with weights summing to one:

$$\begin{aligned}
 f_{p|v}(p|x_t, n_t, v_t, q) &= \frac{f_{v|p}(\tilde{v}_t = v_t|\tilde{p}, n_t, q) f_p(\tilde{p}_t = p_t|x_t)}{f_v(\tilde{v}_t = v_t|x_t, n_t, q)} \\
 &= \frac{\binom{n_t}{v_t} [pq_1 + (1-p)(1-q_2)]^{v_t} [p(1-q_1) + (1-p)q_2]^{n_t-v_t} \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} p^{a_{t,i}-1} (1-p)^{b_{t,i}-1}}{f_v(\tilde{v}_t = v_t|x_t, n_t, q)} \\
 &= \frac{1}{f_v(\tilde{v}_t = v_t|x_t, n_t, q)} \binom{n_t}{v_t} \left(\sum_{j=0}^{v_t} \binom{v_t}{j} p^j q_1^j (1-p)^{v_t-j} (1-q_2)^{v_t-j} \right) \\
 &\quad \times \left(\sum_{k=0}^{n_t-v_t} \binom{n_t-v_t}{k} p^k (1-q_1)^k (1-p)^{n_t-v_t-k} q_2^{n_t-v_t-k} \right) \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} p^{a_{t,i}-1} (1-p)^{b_{t,i}-1} \\
 &= \frac{1}{f_v(\tilde{v}_t = v_t|x_t, n_t, q)} \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \binom{n_t}{v_t} \binom{v_t}{j} \binom{n_t-v_t}{k} q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \\
 &\quad \times \sum_{i=1}^m \omega_i \frac{\Gamma(a_{t,i} + b_{t,i})}{\Gamma(a_{t,i})\Gamma(b_{t,i})} p^{a_{t,i}+j+k-1} (1-p)^{b_{t,i}+n_t-j-k-1} \\
 &= \frac{1}{f_v(\tilde{v}_t = v_t|x_t, n_t, q)} \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \sum_{i=1}^m q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \omega_i \\
 &\quad \times \frac{\Gamma(n_t + 1) \Gamma(a_{t,i} + b_{t,i}) \Gamma(a_{t,i} + j + k) \Gamma(b_{t,i} + n_t - j - k)}{\Gamma(j + 1) \Gamma(v_t - j + 1) \Gamma(k + 1) \Gamma(n_t - v_t - k + 1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i} + b_{t,i} + n_t)} \times \text{beta}(a_{t,i} + j + k, b_{t,i} + n_t - j - k) \\
 &= \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \sum_{i=1}^m \omega'_{j,k,i} \text{beta}(a_{t,i} + j + k, b_{t,i} + n_t - j - k),
 \end{aligned}
 \tag{13}$$

with updated weights

$$\omega'_{j,k,i} = \frac{\omega_i q_1^j (1-q_2)^{v_t-j} (1-q_1)^k q_2^{n_t-v_t-k} \Gamma(a_{t,i}+b_{t,i}) \Gamma(a_{t,i}+j+k) \Gamma(b_{t,i}+n_t-j-k)}{\Gamma(j+1) \Gamma(v_t-j+1) \Gamma(k+1) \Gamma(n_t-v_t-k+1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i}+b_{t,i}+n_t)}$$

$$\sum_{r=0}^{v_t} \sum_{s=0}^{n_t-v_t} \sum_{i=1}^m \frac{\omega_i q_1^r (1-q_2)^{v_t-r} (1-q_1)^s q_2^{n_t-v_t-s} \Gamma(a_{t,i}+b_{t,i}) \Gamma(a_{t,i}+r+s) \Gamma(b_{t,i}+n_t-r-s)}{\Gamma(r+1) \Gamma(v_t-r+1) \Gamma(s+1) \Gamma(n_t-v_t-s+1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i}+b_{t,i}+n_t)}$$

The coefficients of each component of the mixture distribution sum to 1 and, therefore, the posterior distribution is a mixture of beta-distributions, which concludes our proof.

A.2 General properties of mixtures of beta-distributions

A.2.1 The number of component distributions in the posterior mixture distribution

Perfect information-collection technology (i.e., $q_1 = q_2 = 1$). If \tilde{p}_t is a mixture of $m \geq 1$ beta-distributions with weights ω_i and the information-collection technology is perfect, the distribution of sample information is a mixture of beta-binomial distributions with weights ω_i . Updating results in a posterior distribution that is a mixture of m beta-distributions with parameters $(a_{t,i} + v_t, b_{t,i} + n_t - v_t)$.

Consider the example where $m = 2, x_t = ((3, 7), (7, 3)), \omega = (0.8, 0.2), n_t = 10,$ and $v_t = 4$. Let \hat{x}_t denote the posterior belief state. Then, $\hat{x}_t = ((3 + 4, 7 + 6), (7 + 4, 3 + 6)) = ((7, 13), (11, 9))$ and $\omega' = (0.8, 0.2)$.

Imperfect information-collection technology (i.e., $\min\{q_1, q_2\} < 1$). If \tilde{p}_t is a mixture of $m \geq 1$ beta-distributions with weights ω_i and the information-collection technology is imperfect, the true posterior distribution presented in Eq. 13 is a mixture of beta-distributions. There are $n_t + 1$ possible outcomes of a study with sample size of n_t and so there are $n_t + 1$ possible values for $j + k$ in Eq. 13. However, not all the posterior distributions created by updating each of the m prior distributions are necessarily unique. Therefore, the true posterior distribution is a mixture of between $m + n_t$ and $m \times (n_t + 1)$ beta-distributions.

Consider the example where $m = 2, x_t = ((19, 20), (20, 19)), n_t = 5$. Again, let \hat{x}_t denote the posterior belief state. Given the sample size of $n_t = 5$ there are 6 possible true outcomes, by which we mean the unobservable number of actual positive samples in the study, which correspond to $n_t + m = 7$ possible unique beta-distributions $\hat{x}_t = ((19, 25), (20, 24), (21, 23), (22, 22), (23, 21), (24, 20), (25, 19))$ which each contribute to the posterior mixture distribution. When we observe a specific number of positives in the sample, the imperfect information-collection

technology results in a distribution over the true number of positives in the sample and, therefore, weights on each component in the posterior mixture distribution (given by Eq. 13).

A.2.2 Mean and variance of mixtures of beta-distributions

The posterior distribution of \tilde{p}_t given sample information v_t collected using an imperfect information-collection technology is a mixture of beta-distributions. In this section, we first derive equations for the mean and variance for a general mixture of beta-distributions (with simplified notation) to show their relationship to the mean, and more generally, to the parameters of the component distributions. Then, we do the appropriate substitutions to present the conditional mean and variance of the posterior distribution $f_{p|v}$.

Consider a distribution $f_Y(y)$ which is a mixture of M beta-distributions where the i -th component of the mixture has weight w_i , parameters a_i and b_i , and mean μ_i :

$$f_Y(y) = \sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i) \Gamma(b_i)} y^{a_i-1} (1-y)^{b_i-1} \text{ where } \sum_{i=1}^M w_i = 1.$$

We show that the mean of a mixture of beta-distributions is the weighted mean of each mixture component:

$$\begin{aligned} \mathbb{E}[f_Y(y)] &= \int_0^1 y f_Y(y) dy \\ &= \int_0^1 y \left(\sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i) \Gamma(b_i)} y^{a_i-1} (1-y)^{b_i-1} \right) dy \\ &= \sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i) \Gamma(b_i)} \int_0^1 y^{a_i} (1-y)^{b_i-1} dy \\ &= \sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i) \Gamma(b_i)} \frac{\Gamma(a_i + 1) \Gamma(b_i)}{\Gamma(a_i + b_i + 1)} \\ &= \sum_{i=1}^M w_i \frac{a_i}{a_i + b_i} = \sum_{i=1}^M w_i \mu_i. \end{aligned} \tag{14}$$

We also derive the variance of $f_Y(y)$:

$$\begin{aligned} \mathbb{V}[f_Y(y)] &= \int_0^1 y^2 f_Y(y) dy - (\mathbb{E}[f_Y(y)])^2 \\ &= \int_0^1 y^2 \left(\sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i)\Gamma(b_i)} y^{a_i-1} (1-y)^{b_i-1} \right) \\ &\quad \times dy - \left(\sum_{i=1}^M w_i \mu_i \right)^2 \\ &= \sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i)\Gamma(b_i)} \int_0^1 y^{a_i+1} (1-y)^{b_i-1} \\ &\quad \times dy - \left(\sum_{i=1}^M w_i \mu_i \right)^2 \\ &= \sum_{i=1}^M w_i \frac{\Gamma(a_i + b_i)}{\Gamma(a_i)\Gamma(b_i)} \frac{\Gamma(a_i + 2)\Gamma(b_i)}{\Gamma(a_i + b_i + 2)} - \left(\sum_{i=1}^M w_i \mu_i \right)^2 \\ &= \sum_{i=1}^M w_i \frac{(a_i + 1)a_i}{(a_i + b_i + 1)(a_i + b_i)} - \left(\sum_{i=1}^M w_i \mu_i \right)^2. \end{aligned} \tag{15}$$

Mean and variance of the posterior distribution $f_{p|v}$. The posterior distribution of \tilde{p}_t , $f_{p|v}$, is a mixture of beta-distribution (Proposition 1). Using Eqs. 14 and 15 with appropriate substitutions, we can identify:

$$\begin{aligned} \mathbb{E}[f_{p|v}(p|\tilde{v}_t = v_t)] &= \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \sum_{i=1}^m \omega'_{j,k,i} \\ &\quad \times \frac{a_{t,i} + j + k}{a_{t,i} + b_{t,i} + n_t} \end{aligned} \tag{16}$$

and

$$\begin{aligned} \mathbb{V}[f_{p|v}(p|\tilde{v}_t = v_t)] &= \sum_{j=0}^{v_t} \sum_{k=0}^{n_t-v_t} \sum_{i=1}^m \omega'_{j,k,i} \\ &\quad \times \frac{(a_{t,i} + j + k + 1)(a_{t,i} + j + k)}{(a_{t,i} + b_{t,i} + n_t + 1)(a_{t,i} + b_{t,i} + n_t)} \\ &\quad - (\mathbb{E}[f_{p|v}(p|\tilde{v}_t = v_t)])^2 \end{aligned} \tag{17}$$

where ω_i is the prior weight on the i -th component of the prior distribution, and

$$\omega'_{j,k,i} = \frac{\omega_i q_1^j (1 - q_2)^{v_t - j} (1 - q_1)^k q_2^{n_t - v_t - k} \frac{\Gamma(a_{t,i} + b_{t,i}) \Gamma(a_{t,i} + j + k) \Gamma(b_{t,i} + n_t - j - k)}{\Gamma(j + 1) \Gamma(v_t - j + 1) \Gamma(k + 1) \Gamma(n_t - v_t - k + 1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i} + b_{t,i} + n_t)}}{\sum_{r=0}^{v_t} \sum_{s=0}^{n_t - v_t} \sum_{i=1}^m \frac{\omega_i q_1^r (1 - q_2)^{v_t - r} (1 - q_1)^s q_2^{n_t - v_t - s} \Gamma(a_{t,i} + b_{t,i}) \Gamma(a_{t,i} + r + s) \Gamma(b_{t,i} + n_t - r - s)}{\Gamma(r + 1) \Gamma(v_t - r + 1) \Gamma(s + 1) \Gamma(n_t - v_t - s + 1) \Gamma(a_{t,i}) \Gamma(b_{t,i}) \Gamma(a_{t,i} + b_{t,i} + n_t)}}.$$

A.3 Quality of the Posterior Distribution Approximation

We performed extensive numerical simulations to test the accuracy of the approximation. We generated exact posterior distributions under the following conditions:

- Single beta-distribution priors with $\mu(x) \in \{0.001, 0.0025, 0.005, 0.0075, 0.01, 0.02, 0.03, 0.04\}$ and $\sigma(x) \in \{0.001, 0.002, \dots, 0.006\}$.
- Information-collection technologies with $q_1 \in \{0.7, 0.8, 0.9, 0.95, 0.97, 0.99, 1\}$ and $q_2 \in \{0.7, 0.8, 0.9, 0.95, 0.97, 0.99, 0.999, 0.9996, 1\}$.
- Sample sizes $n \in \{100, 500, 1000, 2500, 5000\}$.
- Observations v at percentiles of $\{0.00001, 0.0001, 0.001, 0.002, 0.0025, 0.005, \dots, 0.9925, 0.9975, 0.998, 0.999, 0.9999, 0.99999\}$.

In total, 757,781 exact distributions were generated, approximately 600,000 of which had means and variance in the policy-relevant region for our numerical analysis ($\mu(x) \in (0, 0.04) \times \sigma(x) \in (0, 0.008)$). We calculated the Kolmogorov-distance—the maximum distance between

the cumulative distribution functions—between each exact beta-mixture posterior distribution and a single beta-distribution with the same mean and standard deviation.

The Kolmogorov-distance between the cumulative density function of the exact posterior distributions and that of the approximation with matching mean and variance was generally small ($< 2\%$) (Fig. 6). The Kolmogorov distances only increased in magnitude for very small means or small means and large standard deviations. Kolmogorov distances above 2% typically appeared only in the stopping region of our numerical example or in the upper left-hand section of the information collection region with small means and high standard deviations ($\sigma(x) > 0.006$). Based on the numerical experiments, for the purposes of informing the practical policy decision, the approximate belief update using moment matching proves to be of reasonably high quality.

A.4 Derivation of Dynamics in Mean-Variance space

The state $x_t = (a_t, b_t)$, which contains the parameters of the distribution of \tilde{p}_t representing the policy-maker’s current beliefs, follows a law of motion of the form

$$x_{t+1} = \phi(\hat{x}_t) = \begin{bmatrix} z & 0 \\ 1 - z & 1 \end{bmatrix} \hat{x}_t,$$




Fig. 6 Average Kolmogorov-distance between an exact beta-mixture distribution and a single beta-distribution with the same mean and standard deviation

where $z \in (0, 1)$ is the decay rate and $\hat{x}_t = (\hat{a}_t, \hat{b}_t) = \psi(x_t, n_t, v_t, q)$ is the Bayesian update of x_t given v_t positive observations out of a test of n_t individuals in the current cohort.

First, we see that

$$x_{t+1} = \begin{bmatrix} a_{t+1} \\ b_{t+1} \end{bmatrix} = \begin{bmatrix} z\hat{a}_t \\ (1-z)\hat{a}_t + \hat{b}_t \end{bmatrix}.$$

Derivation of $\mu(x_{t+1})$ beginning with the expectation of a beta-distribution:

$$\begin{aligned} \mu(x_{t+1}) &= \frac{a_{t+1}}{a_{t+1} + b_{t+1}} = \frac{z\hat{a}_t}{z\hat{a}_t + (1-z)\hat{a}_t + \hat{b}_t} \\ &= \frac{z\hat{a}_t}{\hat{a}_t + \hat{b}_t} = z\mu(\hat{x}_t). \end{aligned}$$

Derivation of $\sigma^2(x_{t+1})$ beginning with the variance of a beta-distribution:

$$\begin{aligned} \sigma^2(x_{t+1}) &= \frac{a_{t+1}b_{t+1}}{(a_{t+1} + b_{t+1})^2(a_{t+1} + b_{t+1} + 1)} \\ &= \frac{(z\hat{a}_t)((1-z)\hat{a}_t + \hat{b}_t)}{(z\hat{a}_t + (1-z)\hat{a}_t + \hat{b}_t)^2(z\hat{a}_t + (1-z)\hat{a}_t + \hat{b}_t + 1)} \\ &= \frac{z(1-z)\hat{a}_t^2 + z\hat{a}_t\hat{b}_t}{(\hat{a}_t + \hat{b}_t)^2(\hat{a}_t + \hat{b}_t + 1)} \\ &= z(1-z) \left(\frac{\hat{a}_t^2}{(\hat{a}_t + \hat{b}_t)^2(\hat{a}_t + \hat{b}_t + 1)} \times \frac{\hat{b}_t}{\hat{b}_t} \times \frac{\hat{a}_t + \hat{b}_t}{\hat{a}_t + \hat{b}_t} \right) \\ &\quad + z \left(\frac{\hat{a}_t\hat{b}_t}{(\hat{a}_t + \hat{b}_t)^2(\hat{a}_t + \hat{b}_t + 1)} \right) \end{aligned}$$

$$\begin{aligned} &= z(1-z) \left(\frac{\hat{a}_t\hat{b}_t}{(\hat{a}_t + \hat{b}_t)^2(\hat{a}_t + \hat{b}_t + 1)} \right. \\ &\quad \left. \times \frac{\hat{a}_t}{\hat{a}_t + \hat{b}_t} \times \frac{\hat{a}_t + \hat{b}_t}{\hat{b}_t} \right) \\ &\quad + z\sigma^2(x_t) \\ &= z(1-z)\sigma^2(\hat{x}_t) \left(\frac{\mu(\hat{x}_t)}{1 - \mu(\hat{x}_t)} \right) + z\sigma^2(\hat{x}_t) \\ &= \sigma^2(\hat{x}_t) \left(z + z(1-z) \left(\frac{\mu(\hat{x}_t)}{1 - \mu(\hat{x}_t)} \right) \right). \end{aligned}$$

A.5 Derivation of Eq. 8

For $\mu(x_0) \geq \frac{\gamma}{\theta}$, given Eq. 1 and using the fact that $\mu(x_t) = z^t\mu(x_0)$, we can identify the optimal time to stop the intervention, $T(x_0)$, which is the first period in which the intervention has a nonpositive INMB. Specifically, we seek the value of t such that $E[g(x_t)] = 0$.

$$\begin{aligned} E[g(\tilde{p}_t)] &= \theta\mu(x_t) - \gamma = \theta z^t\mu(x_0) - \gamma = 0 \\ z^t &= \frac{\gamma}{\theta\mu(x_0)} \\ t &= \frac{1}{\ln(z)} \ln\left(\frac{\gamma}{\theta\mu(x_0)}\right) \end{aligned}$$

Since decisions can only be made at discrete time intervals, we identify the optimal time to stop the intervention, $T(x_0)$, as the first integer period in which the intervention has a nonpositive INMB

$$T(x_0) = \left\lceil \frac{1}{\ln(z)} \ln\left(\frac{\gamma}{\theta\mu(x_0)}\right) \right\rceil.$$

A.6 Proof of Proposition 2

Rewriting Eq. 9 as

$$V_{\text{NoInfo}}(x_0) = \sum_{t=0}^{\infty} \delta^t (\theta z^t \mu(x_0) - \gamma) \mathbf{1}_{\{t \leq T(x_0)\}},$$

where $T(x_0)$ is the optimal time to stop the intervention, the claim follows from the fact that each term is a nondecreasing, convex function. Specifically for each t , $\delta^t(\theta z^t \mu(x_0) - \gamma)$ is linear in $\mu(x_0)$ with $\theta \delta^t z^t \geq 0$ and $T(x_0)$ nondecreasing in $\mu(x_0)$. Noting that when $t = T(x_0) = t(x_0)$, $(\theta z^t \mu(x_0) - \gamma) = 0$, we can conclude that $\delta^t(\theta z^t \mu(x_0) - \gamma) \mathbf{1}_{\{t \leq T(x_0)\}}$ is a continuous, nondecreasing convex function.

A.7 Proof of Proposition 3

We apply Proposition 5 in Smith and McCardle [20], which states that if (a) the current-period reward function satisfies

the structural property (such as convexity and monotonicity in $\mu(x_t)$) for each action, and (b) the state transitions satisfy a stochastic version of the structural property for each action, then the value function satisfies the structural property in a finite-horizon setting. In our setting, which is in principle infinite-horizon up to the stopping time, if the structural property is satisfied in the final period, just before the optimal stopping time, then the previous-period value function is obtained via maximization over functions that each satisfy the structural property. Thus, if the structural property is preserved by maximization, then the previous-period value function also satisfies the structural property. For the proof we assume a perfect detection technology to simplify exposition. We extend to the general case at the end of the proof.

First, condition (a) is satisfied because the current-period expected reward, $\mathbb{E}[g(\tilde{p}_t, u_t)|x_t]$, is nondecreasing in $\mu(x_t)$ and $\sigma^2(x_t)$ and (at least weakly) convex in $\mu(x_t)$ for any action $u_t \in \mathcal{D} \times \mathcal{N}$.

Second, the Bayesian update ψ preserves stochastic dominance of the beta-distributed prior, in the sense that if one prior is stochastically dominated by another prior, the corresponding posteriors will exhibit the same dominance relationship, under first-order stochastic dominance (FOSD) and second-order stochastic dominance (SOSD).⁵ In particular, an increase of the mean $\mu(x_t)$ will result in an increase of the posterior mean, and an increase of the variance $\sigma^2(x_t)$ in an increase of the posterior variance. To see this for SOSD, consider a mean preserving spread, so $x_t^{(1)}, x_t^{(2)}$ with $\mu(x_t^{(1)}) = \mu(x_t^{(2)})$ and $\sigma^2(x_t^{(1)}) < \sigma^2(x_t^{(2)})$. Then $\sigma^2(x_t^{(1)}) < \sigma^2(x_t^{(2)})$ implies $a_t^{(1)} + b_t^{(1)} > a_t^{(2)} + b_t^{(2)}$, and $\mathbb{V}[\tilde{s}_t|x_t^{(1)}, n_t] \leq \mathbb{V}[\tilde{s}_t|x_t^{(2)}, n_t]$. The conditional variance of the next-period belief given $n_t = \eta \geq 0$ samples is $\mathbb{V}[x_{t+1}|x_t, \tilde{s}_t, n_t = \eta] = \frac{z(a_t + \tilde{s}_t)(a_t + b_t + \eta - z(a_t + \tilde{s}_t))}{(a_t + b_t + \eta)^2(a_t + b_t + \eta + 1)}$. The variance of the next-period belief is obtained using the law of total variance, so $\mathbb{V}[x_{t+1}|x_t^{(1)}, n_t = \eta] < \mathbb{V}[x_{t+1}|x_t^{(2)}, n_t = \eta]$; hence, $\phi \circ \psi$ is increasing in σ^2 . Thus, the Bayesian update is increasing in $\mu(x_t)$ and $\sigma^2(x_t)$, and the same holds true for its beta-approximation ψ introduced in Section 2.2. The state-transition function ϕ is linear (time-invariant) in $\mu(x_t)$ and linear (time-variant) in $\sigma^2(x_t)$, so that the Bayesian-updated state-transition function $\phi \circ \psi$ is increasing in $(\mu(x_t), \sigma^2(x_t))$.

Finally, the convexity in $\mu(x_t)$ survives the maximization in the Bellman equation given that the objective function is supermodular in $(u_t, \mu(x_t))$. Since the sum of nondecreasing convex functions is nondecreasing and convex, we only need a terminal condition to satisfy the backwards-induction

approach presented by Smith and McCardle [20]. Note that since $0 < z < 1$, $\lim_{t \rightarrow \infty} \mu(\phi(\psi(x_t, n_t, \tilde{v}_t, q))) = 0$ and $\lim_{t \rightarrow \infty} \sigma^2(\phi(\psi(x_t, n_t, \tilde{v}_t, q))) = 0$. Therefore, there exists a time $T < \infty$ for which, given any initial state, an optimal policy is to stop the intervention, i.e., $u_T = (0, 0)$, and $V(x_T) = 0$. Through the mean- and, ultimately, variance-reducing dynamics, with or without the variance-reducing acquisition of information, any initial state eventually approaches a ‘termination’ state over time. Since the reward of this state is zero, which is nondecreasing and convex, we conclude that $V(x_t)$ is nondecreasing and convex in $\mu(x_t)$; by a similar argument it is also increasing in $\sigma^2(x_t)$.

This proof relies mainly on the stochastic-dominance ordering, and it therefore directly extends to the case with misclassification. Condition (a) continues to be satisfied, since it relates only to the current period and is not influenced by information collection. Second, as before, Bayesian updating preserves stochastic dominance of the beta-mixture prior, in the sense that if one prior stochastically dominates another prior, the corresponding posteriors conserve the dominance ordering, for FOSD and SOSD. Finally, imperfect information collection does not affect the supermodularity of the objective function in $(u_t, \mu(x_t))$, so convexity in $\mu(x_t)$ survives the maximization in the Bellman equation. This allows for backward induction starting with the ‘stop intervention’-region at zero reward, as described above.

A.8 Proof of Corollary 1

For the case where no information is available, this corollary was already shown to be true with the derivation of a threshold policy in Section 3.1. For the general case, with or without information collection in the current or future periods, we rely on the properties of the value function demonstrated in Proposition 3: if $\mu(x_t^{(1)}) < \mu(x_t^{(2)})$ and $\sigma(x_t^{(1)}) = \sigma(x_t^{(2)})$, then $V(x_t^{(1)}) \leq V(x_t^{(2)})$. This directly implies that, if it is optimal to do the intervention with $\mu(x_t^{(1)})$, then it is also optimal to do the intervention at $\mu(x_t^{(2)})$. Furthermore, if it is *not* optimal to do the intervention with $\mu(x_t^{(2)})$ then, because $V(x_t^{(1)}) \leq V(x_t^{(2)})$, it is also not optimal to do the intervention at $\mu(x_t^{(1)})$.

A.9 Proof of Corollary 2

For the case where no information is available, the optimal policy does not depend on $\sigma(x_t)$ (see the threshold policy in Section 3.1). For the general case, with or without information collection in the current or future periods, we

⁵The monotone-likelihood-ratio property is satisfied [89].

rely on the properties of the value function demonstrated in Proposition 3: if $\mu(x_t^{(1)}) = \mu(x_t^{(2)})$ and $\sigma(x_t^{(1)}) < \sigma(x_t^{(2)})$, then $V(x_t^{(1)}) \leq V(x_t^{(2)})$. This directly implies, if it is optimal to do the intervention with $\sigma(x_t^{(1)})$, then it is also optimal to do the intervention at $\sigma(x_t^{(2)})$. Furthermore, if it is *not* optimal to do the intervention with $\sigma(x_t^{(2)})$ then, because $V(x_t^{(1)}) \leq V(x_t^{(2)})$, it is also not optimal to do the intervention at $\sigma(x_t^{(1)})$.

A.10 Proof of Proposition 4

Misclassification results in a posterior distribution which has greater variance than would occur with the same sample size and a perfect detection technology. When evaluating the expected value with information collection, greater variance in the posterior implies an expected value over a larger number of states where the optimal next action is to not do the intervention and, therefore, have a value of 0.

Given a perfect detection technology, smaller sample sizes will have greater variance in the posterior distribution than larger sample sizes. Therefore, misclassification can be thought of as an effective sample size reduction or, equivalently, an increase in cost for each full unit of information.

An increase in cost or a decrease in the expected value of the next state given information that was collected this period, decrease the value of the information-collection alternative. Therefore, there are fewer states for which information collection is the optimal action, i.e., the action providing the greatest expected value. Misclassification has no effect on the option not to do the intervention, and has a limited effect on the immediate option not to collect information this period (it would influence this option only

when the optimal action of a subsequent state is to collect information).

A.11 Supplemental methods and results for HCV screening example

A.11.1 Development of linear INMB

A schematic of the HCV screening decision problem for a single cohort is presented in Fig. 7.

We denote λ as the willingness-to-pay threshold, q_1 and q_2 as the test sensitivity and specificity, $C_S > 0$ as the cost of the screening test, $B_S \leq 0$ as the quality-of-life loss from the screening test, $C_{FP} \geq 0$ as the cost of correcting a false-positive test result, $B_{FP} \leq 0$ as the quality-of-life loss from a false-positive test result, and we denote the lifetime discounted costs and benefits of the true-positive, false-negative, and true-negative screening outcomes, C_1, C_2, C_3 , and B_1, B_2, B_3 , respectively.

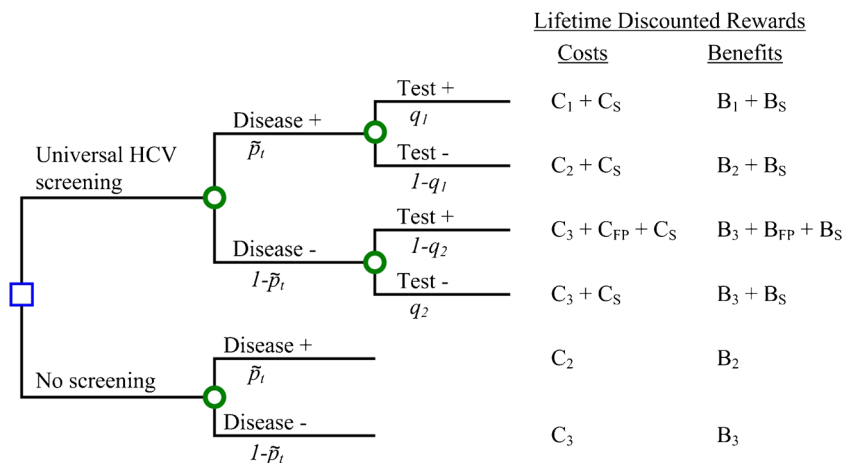
The net monetary benefit (NB) of the decision not to screen cohort t is

$$NB_{NoScreening} = \lambda (\tilde{p}_t B_2 + (1 - \tilde{p}_t) B_3) - (\tilde{p}_t C_2 + (1 - \tilde{p}_t) C_3). \tag{18}$$

The net monetary benefit (NB) of the decision not to screen cohort t is

$$NB_{Screening} = \lambda (\tilde{p}_t q_1 B_1 + \tilde{p}_t (1 - q_1) B_2 + (1 - \tilde{p}_t) \times (1 - q_2) (B_3 + B_{FP}) + (1 - \tilde{p}_t) q_2 B_3 + B_S) - (\tilde{p}_t q_1 C_1 + \tilde{p}_t (1 - q_1) C_2 + (1 - \tilde{p}_t) \times (1 - q_2) (C_3 + C_{FP}) + (1 - \tilde{p}_t) q_2 C_3 + C_S). \tag{19}$$

Fig. 7 A schematic of the medical intervention decision. Each year, the policy-maker must choose whether or not to screen for HCV infection in 50-year olds who attend a routine preventive health exam



The INMB of screening compared to the alternative of not screening is computed as the difference between Eqs. 19 and 18:

$$\begin{aligned}
 \text{INMB}_{\text{Screening}} &= \text{NB}_{\text{Screening}} - \text{NB}_{\text{NoScreening}} \\
 &= \lambda (\tilde{p}_t q_1 B_1 + \tilde{p}_t (1 - q_1) B_2 + (1 - \tilde{p}_t) \\
 &\quad \times (1 - q_2)(B_3 + B_{FP}) + (1 - \tilde{p}_t) \\
 &\quad \times q_2 B_3 + B_S) \\
 &\quad - (\tilde{p}_t q_1 C_1 + \tilde{p}_t (1 - q_1) C_2 + (1 - \tilde{p}_t) \\
 &\quad \times (1 - q_2)(C_3 + C_{FP}) + (1 - \tilde{p}_t) \\
 &\quad \times q_2 C_3 + C_S) \\
 &\quad - (\lambda (\tilde{p}_t B_2 + (1 - \tilde{p}_t) B_3) \\
 &\quad - (\tilde{p}_t C_2 + (1 - \tilde{p}_t) C_3)) \\
 &= \lambda \tilde{p}_t [q_1 B_1 + (1 - q_1) B_2 - (1 - q_2) \\
 &\quad \times (B_3 + B_{FP}) - q_2 B_3 - B_2 + B_3] \\
 &\quad + \lambda [(1 - q_2)(B_3 + B_{FP}) + q_2 B_3 \\
 &\quad + B_S - B_3] \\
 &\quad - \tilde{p}_t [q_1 C_1 + (1 - q_1) C_2 - (1 - q_2) \\
 &\quad \times (C_3 + C_{FP}) - q_2 C_3 - C_2 + C_3] \\
 &\quad - (1 - q_2)(C_3 + C_{FP}) - q_2 C_3 - C_S + C_3 \\
 &= \tilde{p}_t (q_1 [\lambda (B_1 - B_2) - (C_1 - C_2)] \\
 &\quad \times - (1 - q_2) [\lambda B_{FP} - C_{FP}]) \\
 &\quad + \lambda B_S - C_S + (1 - q_2)(\lambda B_{FP} - C_{FP})
 \end{aligned}$$

With terms collected, the INMB of screening at age 50 compared to not screening at time t in a cohort with HCV prevalence \tilde{p}_t can be written $\text{INMB}_t = \theta \tilde{p}_t - \gamma$. The marginal INMB of early diagnosis and treatment for an individual with HCV, θ , and the fixed INMB of screening, γ , are

$$\theta = q_1 [\lambda (B_1 - B_2) - (C_1 - C_2)] - (1 - q_2) [\lambda B_{FP} - C_{FP}], \tag{20}$$

and

$$\gamma = C_S - \lambda B_S - (1 - q_2)(\lambda B_{FP} - C_{FP}). \tag{21}$$

A.11.2 Parameter estimation

Consistent with the recommendations of the US Panel on Cost-Effectiveness in Health and Medicine, we adopted a societal perspective, considered costs and benefits over a lifetime horizon for each cohort, and discounted future costs and health benefits at 3% annually [1]. We measured costs in 2010 US dollars and adjusted for inflation using the Consumer Price Index when appropriate [90]. Benefits are measured in quality-adjusted life-years (QALYs).

We assumed a mid-range value for society’s maximum willingness to pay of \$75,000 per QALY gained [84].

Estimating the lifetime costs and benefits of each HCV screening outcome for cohorts of asymptomatic 50-year old men and women requires a detailed natural history model of HCV. We used the model by Liu et al. [14] to estimate the lifetime costs and benefits of each HCV screening outcome.

We assumed that the cohort size at each period, the number of people who attend a preventive health exam at age 50, is constant over time, N , since there is less than 10% variation from the average population size across cohorts currently aged between 25 and 55 years of age [79]. At the beginning of each period t , the policy-maker simultaneously decides whether to screen the current cohort for HCV and whether to conduct a study of sample size n_t to better estimate the current prevalence of HCV. Information arrives at the end of the current period and is used, together with the prevalence dynamics, to inform the screening decision at $t + 1$ for the next cohort. We assumed that the cost of sample information, $K(n_t)$, is affine in the sample size with a fixed cost of \$50,000 and variable cost of \$100 [83].

We used the National Health and Nutrition Examination Survey (NHANES) to estimate birth-cohort-specific HCV prevalence, HCV-prevalence dynamics, and the proportion of individuals currently unaware of their infection status. Ultimately, we estimated the HCV prevalence for our initial cohorts, men and women born in 1960 who are currently unaware of their infection status, to be 3.1% (95% CI: 2.4-3.8%) and 1.4% (95% CI: 1.0-1.7%), respectively. Restricting the analysis to individuals born between 1956 and 1980 ($n = 12,607$), we identified the rate of prevalence decay to be 0.893 (95% CI: 0.871-0.915) using logistic regression, controlling for race and gender. We present the detailed methods and results of this primary analysis below.

A.11.3 National health and nutrition examination survey (NHANES) analysis

Overview The National Center for Health Statistics periodically conducts NHANES to compile representative statistics on the health of the US population [88]. Our analysis includes data collected from 1999 through 2010. Participants were chosen according to a stratified multistage algorithm to produce a representative sample of the civilian, non-institutionalized population of all 50 states and the District of Columbia. Only participants at least age 6 years old were eligible for HCV testing because of low blood sample volume in younger children. Birth years for individuals younger than 85 years, for survey years 1999-2006, and for individuals younger than 80 years old, for survey years

Table 5 Logistic regression predicting HCV prevalence for 1956-1980 birth cohorts using NHANES 1999-2010 ($n = 12,607$)

Variable	Estimate	SE	p-value	Odds Ratio (95% CI)
Intercept	217.7	23.85	< .0001	
Birth year	-0.113	0.012	< .0001	0.893 (0.872 – 0.915)
Race: Black (vs. non-Black)	0.226	0.073	0.0021	1.57 (1.18 – 2.10)
Gender: Female (vs. Male)	-0.256	0.071	0.0003	0.6 (0.46 – 0.79)

SE – Standard error; CI – Confidence interval

2008-2010, were estimated using their age in months and the 6-month window of the survey. Age in months is not provided for older individuals. Beginning with the 2001/2002 survey, participants testing HCV-positive were informed in writing of their test results, and four months later, they received a follow-up telephone questionnaire.

Statistical Analysis All statistical analyses were performed in SAS version 9.3 (SAS Institute, Cary, North Carolina) according to National Center for Health Statistics guidelines [86, 87]. We accounted for the complex survey design using the appropriate study design variables, sampling weights, and by using SAS Survey procedures. Logistic regression analysis was used to identify the rate of HCV-prevalence decay over successive birth cohorts for individuals born between 1956 and 1980. Finally, using the follow up survey in HCV-positive individuals, we estimated the proportion of HCV-positive individuals who were unaware of their infection status prior to participation in NHANES.

Results Of 51,587 participants of at least age 6 years surveyed between 1999 and 2010, 45,153 gave a blood sample suitable for HCV-antibody testing (final response rate for testing, 87.5%). Restricting analysis to individuals born between years 1956 and 1980 ($n=12,607$), we identified the rate of prevalence decay to be 0.893 (95% CI: 0.871-0.915) using logistic regression controlling for race and gender

(Table 5). Using the regression, the predicted HCV prevalence for men and women born in 1960 are 4.7% (95% CI: 3.8-5.7%) and 2.9% (95% CI: 2.3-3.6%), respectively.

Since the 2001/02 survey, 500 subjects were identified as HCV-positive and contacted for follow-up which included asking if they were previously aware of their HCV-infection status. The response rate to the follow-up questionnaire was 206 (41%). Using logistic regression, we estimated the proportion of men and women who were unaware of their HCV-infection status prior to participating in the NHANES study to be 55% (95% CI: 46-65%) and 39% (95% CI: 30-39%), respectively (Table 6). Because of the small sample size, we did not stratify analysis by birth year; race was excluded from the final regression model because it was not a significant predictor of prior infection-status awareness.

To compute the HCV prevalence among those who are currently unaware of their infection status, we also needed an estimate of the proportion of individuals who are aware of their HCV-negative status, which is unknown. We assumed it to be 15%, consistent with Liu et al. [14]. Using the logistic regression model to predict birth-cohort-specific HCV prevalence and adjusting for the number of individuals who are unaware of their infection status, we estimated the HCV prevalence for our initial cohorts, men and women born in 1960 who are currently unaware of their infection status, is 3.1% (95% CI: 2.4-3.8%) and 1.4% (95% CI: 1.0-1.7%), respectively.

Table 6 Logistic regression predicting the proportion of HCV-positive individuals unaware of their infection status using NHANES 1999-2010 ($n = 206$)

Variable	Estimate	SE	p-value	Odds Ratio (95% CI)
<u>Model 1:</u>				
Intercept	0.00022	0.167	0.998	
Race: Black (vs. non-Black)	0.161	0.157	0.304	1.40 (0.75 – 2.55)
Gender: Female (vs. Male)	-0.332	0.140	0.017	0.51 (0.30 – 0.89)
<u>Model 2 (Final Model):</u>				
Intercept	-0.111	0.144	0.444	
Gender: Female (vs. Male)	-0.329	0.140	0.019	0.52 (0.30 – 0.90)

SE – Standard error; CI – Confidence interval

A.11.4 Sensitivity analysis

We performed sensitivity analysis to evaluate the robustness of the optimal policy to uncertainty in model inputs (Table 7). We identified that the general conclusions of our analysis are robust to the uncertainty inputs. Specifically, for women, we find that the optimal time to collect information

ranges from immediately to 7 years. For men, we find that the optimal time to collect information ranges from 11-21 years with the exception of scenarios in which we considered a low value of z . A very low value of z implies the prevalence of HCV is rapidly decreasing across birth cohorts. If this is the case, it is optimal to collect additional information immediately.

Table 7 Comparison of optimal policies indicated by various analytic approaches

Case	Adjusted Parameter Value	Optimal Policy*	Increase in expected INMB compared to screening for 5 years**
<u>Males</u>			
Base Case		Sample 4,000 men in 16 years (1976 birth cohort), then identify optimal action	\$168, 800, 000
Low θ	$\theta = 3000$	Sample in 11 years (1971 birth cohort)	\$51, 410, 000
High θ	$\theta = 10, 000$	Sample in 20 years (1980 birth cohort)	\$390, 240, 000
Low γ	$\gamma = 22$	Sample in 18 years (1978 birth cohort)	\$197, 990, 000
High γ	$\gamma = 40$	Sample in 13 years (1973 birth cohort)	\$120, 800, 000
Low z	$z = 0.871$	Sample immediately	\$185, 690, 000
High z	$z = 0.915$	Sample in 21 years (1981 birth cohort)	\$261, 390, 000
High κ_F	$\kappa_F = 250, 000$	Sample in 16 years (1976 birth cohort)	\$168, 600, 000
Low $\mu(x_0)$	$\mu(x_0) = 0.024$	Sample in 13 years (1973 birth cohort)	\$106, 040, 000
High $\mu(x_0)$	$\mu(x_0) = 0.038$	Sample in 18 years (1978 birth cohort)	\$235, 620, 000
Scenario 1	$\theta = 8500$; $\gamma = 25$; $z = 0.88$	Sample immediately	\$257, 240, 000
Scenario 2	$\theta = 5500$; $\gamma = 35$; $z = 0.905$	Sample in 14 years (1974 birth cohort)	\$108, 430, 000
<u>Females</u>			
Base Case		Sample 4,500 women in 3 years (1963 birth cohort), then identify optimal action	\$7, 110, 000
Low θ	$\theta = 1500$	Never initiate screening / Stop screening	\$63, 330, 000
High θ	$\theta = 6000$	Sample in 7 years (1967 birth cohort)	\$22, 550, 000
Low γ	$\gamma = 22$	Sample in 4 years (1964 birth cohort)	\$610, 000
High γ	$\gamma = 40$	Sample in 1 years (1961 birth cohort)	\$46, 040, 000
Low z	$z = 0.871$	Sample immediately	\$14, 570, 000
High z	$z = 0.915$	Sample in 3 years (1963 birth cohort)	\$2, 460, 000
High κ_F	$\kappa_F = 250, 000$	Sample in 1 years (1961 birth cohort)	\$6, 900, 000
Low $\mu(x_0)$	$\mu(x_0) = 0.024$	Sample immediately	\$29, 320, 000
High $\mu(x_0)$	$\mu(x_0) = 0.038$	Sample in 4 years (1964 birth cohort)	\$340, 000
Scenario 1	$\theta = 8500$; $\gamma = 25$; $z = 0.88$	Sample immediately	\$7, 090, 000
Scenario 2	$\theta = 5500$; $\gamma = 35$; $z = 0.905$	Never initiate screening / Stop screening	\$69, 020, 000

*Sample size for men = 4,000; Sample size for women = 4,500

**Expected value of the optimal policy with the parameter change compared to the expected value of the CDC/USPSTF recommendation of screening for 5 years (calculated with the adjusted input parameters)

References

- Gold MR, Siegel JE, Russell LB, Weinstein MC (1996) *Cost-Effectiveness in Health and Medicine*. Oxford University Press, Oxford
- Drummond MF, Sculpher MJ, Torrance GW (2005) *Methods for the Economic Evaluation of Health Care Programs*, 3rd edn. Oxford University Press, Oxford
- Ades AE, Lu G, Claxton KP (2004) Expected value of sample information calculations in medical decision modeling. *Med Decis Making* 24(2):207–227
- Claxton KP, Sculpher MJ (2006) Using value of information analysis to prioritise health research: Some lessons from recent UK experience. *Pharmacoeconomics* 24(11):1055–1068
- Eckermann S, Karnon J, Willan AR (2010) The value of value of information. *Pharmacoeconomics* 28(9):699–709
- Phillips Z, Claxton K, Palmer S (2008) The half-life of truth: what are appropriate time horizons for research decisions? *Med Decis Making* 28(3):287–299
- Eckermann S, Willan AR (2008) The option value of delay in health technology assessment. *Med Decis Making* 28(3):300–305
- Juusola JL, Brandeau ML (2016) HIV treatment and prevention: a simple model to determine optimal investment. *Med Decis Making* 36(3):391–409
- Singer ME, Younossi ZM (2001) Cost effectiveness of screening for hepatitis C virus in asymptomatic, average-risk adults. *Am J Med* 111(8):614–621
- Chou R, Clark EC, Helfand M (2004) Screening for hepatitis C virus infection: a review of the evidence for the U.S. Preventive Services Task Force. *Ann Intern Med* 140(6):465–479
- Rein DB, Smith BD, Wittenborn JS, Lesesne SB, Wagner LD, Roblin DW, Patel N, Ward JW, Weinbaum CM (2012) The cost-effectiveness of birth-cohort screening for hepatitis C antibody in US primary care settings. *Ann Intern Med* 156(4):263–270
- Coffin PO, Scott JD, Golden MR, Sullivan SD (2012) Cost-effectiveness and population outcomes of general population screening for hepatitis C. *Clin Infect Dis* 54(9):1259–1271
- McGarry LJ, Pawar VS, Panchmatia HR, Rubin JL, Davis GL, Younossi ZM, Capretta JC, O'Grady MJ, Weinstein MC (2012) Economic model of a birth cohort screening program for hepatitis C virus. *Hepatology* 55(5):1344–1355
- Liu S, Cipriano LE, Holodniy M, Goldhaber-Fiebert JD (2013) Cost-effectiveness analysis of risk-factor guided and birth-cohort screening for chronic hepatitis C infection in the United States. *PLoS One* 8(3):e58975
- Eckman MH, Talal AH, Gordon SC, Schiff E, Sherman KE (2013) Cost-effectiveness of screening for chronic hepatitis C infection in the United States. *Clin Infect Dis* 56(10):1382–1393
- Smith BD, Morgan RL, Beckett GA, Falck-Ytter Y, Holtzman D, Ward JW (2012) Hepatitis C virus testing of persons born during 1945–1965: Recommendations from the Centers for Disease Control and Prevention. *Ann Intern Med* 157(11):817–822
- Moyer VA, on behalf of the U.S. Preventive Services Task Force (2013) Screening for Hepatitis C Virus Infection in Adults: U.S. Preventive Services Task Force Recommendation Statement. *Ann Intern Med* 159(5):349–357
- Jensen R (1983) Innovation adoption and diffusion when there are competing innovations. *J Econ Theory* 29(1):161–171
- McCardle KF (1985) Information acquisition and the adoption of new technology. *Manag Sci* 31(11):1372–1389
- Smith JE, McCardle KF (2002) Structural properties of stochastic dynamic programs. *Oper Res* 50(5):796–809
- Ulu C, Smith JE (2009) Uncertainty, information acquisition, and technology adoption. *Oper Res* 57(3):740–752
- Rosenberg N (1982) *Inside the black box: technology and economics*. Cambridge University Press, Cambridge
- Bessen J (1999) Real options and the adoption of new technologies. *Research on Innovation*. <http://www.researchoninnovation.org>
- Kornish LJ (2006) Technology choice and timing with positive network effects. *Eur J Oper Res* 173(1):268–282
- Chambers C, Kouvelis P (2003) Competition, learning and investment in new technology. *IIE Trans* 35(9):863–878
- Schaefer AJ, Bailey MD, Shechter SM, Roberts MS (2005) Chapter 23: Modeling medical treatment using Markov decision processes. In: *Operations Research and Health Care: A handbook of methods and applications*, Springer US, volume 70 of International Series in Operations Research and Management Science. pp. 593–612
- Alagoz O, Hsu H, Schaefer AJ, Roberts MS (2010) Markov decision processes: a tool for sequential decision making under uncertainty. *Med Decis Making* 30(4):474–483
- Ahn JH, Hornberger JC (1996) Involving patients in the cadaveric kidney transplant allocation process: A decision-theoretic perspective. *Manag Sci* 42(5):629–641
- Magni P, Quaglini S, Marchetti M, Barosi G (2000) Deciding when to intervene: a Markov decision process approach. *Int J Med Inform* 60(3):237–253
- Hauskrecht M, Fraser H (2000) Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artif Intell Med* 18(3):221–244
- Alagoz O, Maillart LM, Schaefer AJ, Roberts MS (2004) The optimal timing of living–donor liver transplantation. *Manag Sci* 50(10):1420–1430
- Alagoz O, Maillart LM, Schaefer AJ, Roberts MS (2007) Choosing among cadaveric and living–donor livers. *Manag Sci* 53(11):1702–1715
- Shechter SM, Bailey MD, Schaefer AJ, Roberts MS (2008) The optimal time to initiate HIV therapy under ordered health states. *Oper Res* 56(1):20–33
- Shechter SM, Bailey MD, Schaefer AJ, Roberts MS (2008) A modeling framework for replacing medical therapies. *IIE Trans* 40(9):861–869
- Kirkizlar E, Faissol DM, Griffin PM, Swann JL (2010) Timing of testing and treatment for asymptomatic diseases. *Math Biosci* 226(1):28–37
- Kurt M, Denton B, Schaefer AJ, Shah N, Smith S (2011) The structure of optimal statin initiation policies for patients with type 2 diabetes. *IIE Trans* 1(1):49–65
- Mason JE, Denton BT, Shah ND, Smith SA (2014) Optimizing the simultaneous management of blood pressure and cholesterol for type 2 diabetes patients. *Eur J Oper Res* 233(3):727–738
- Zhang J, Denton BT, Balasubramanian H, Shah ND, Inman BA (2012) Optimization of prostate biopsy referral decisions. *Manuf Serv Op* 14(4):529–547
- Ayer T, Alagoz O, Stout NK, Burnside ES (2014) Designing a new breast cancer screening program considering adherence. *Manag. Sci.* Forthcoming
- Erenay FS, Alagoz O, Said A (2014) Optimizing colonoscopy screening for colorectal cancer prevention and surveillance. *Manuf Serv Op* 16(3):381–400
- Patrick J, Puterman ML, Queyranne M (2008) Dynamic multi-priority patient scheduling for a diagnostic resource. *Oper Res* 56(6):1507–1525
- Gocgun Y, Bresnahan BW, Ghate A, Gunn ML (2011) A Markov decision process approach to multi-category patient scheduling in a diagnostic facility. *Artif Intell Med* 53:73–81
- Patrick J (2012) A Markov decision model for determining optimal outpatient scheduling. *Health Care Manag Sci* 15:91–102

44. Gocgun Y, Puterman ML (2014) Dynamic scheduling with due dates and time windows: an application to chemotherapy patient appointment booking. *Health Care Manag Sci* 17:60–76
45. Gupta D, Wang L (2008) Revenue management for a primary-care clinic in the presence of patient choice. *Oper Res* 56(3):576–592
46. Wang J, Fung RYK (2015) Adaptive dynamic programming algorithms for sequential appointment scheduling with patient preferences. *Artif Intell Med* 63:33–40
47. Kornish LJ, Keeney RL (2008) Repeated commit-or-defer decisions with a deadline: the influenza vaccine composition. *Oper Res* 56(3):527–541
48. Özaltın OY, Prokopyev OA, Schaefer AJ, Roberts MS (2011) Optimizing the societal benefits of the annual influenza vaccine: a stochastic programming approach. *Oper Res* 59(5):1131–1143
49. Weinstein M, Zeckhauser R (1973) Critical ratios and efficient allocation. *J Public Econ* 2:147–157
50. Culyer AJ (1989) The normative economics of health care finance and provision. *Oxf Rev Econ Policy* 5:34–58
51. Stinnett AA, Paltiel AD (1996) Mathematical programming for the efficient allocation of health care resources. *J Health Econ* 15:641–653
52. Williams I, McIver S, Moore D, Bryan S (2008) The use of economic evaluations in NHS decision-making: a review and empirical investigation. *Health Technol Assess* 12(7):1–175
53. Raiffa H, Schlaifer R (1961) *Applied Statistical Decision Theory*. Harvard University Press, Cambridge
54. Weinstein MC (1983) Cost-effective priorities for cancer prevention. *Science* 221(4605):17–23
55. Hornberger JC, Brown BW, Halpern J (1995) Designing a cost-effective clinical trial. *Stat Med* 14(20):2249–2259
56. Claxton K, Posnett J (1996) An economic approach to clinical trial design and research priority-setting. *Health Econ* 5(6):513–524
57. Claxton K (1999) The irrelevance of inference: a decision-making approach to the stochastic evaluation of health care technologies. *J Health Econ* 18(3):341–364
58. Eckermann S, Willan AR (2008) Time and expected value of sample information wait for no patient. *Value Health* 11(3):522–526
59. McKenna C, Claxton K (2011) Addressing adoption and research design decisions simultaneously: the role of value of sample information analysis. *Med Decis Making* 31(6):853–865
60. Hall PS, Edlin R, Kharroubi S, Gregory W, McCabe C (2012) Expected net present value of sample information from burden to investment. *Med Decis Making* 32(3):E11–E21
61. Fenwick E, Claxton K, Sculpher M (2008) The value of implementation and the value of information: combined and uneven development. *Med Decis Making* 28(1):21–32
62. Willan AR, Eckermann S (2010) Optimal clinical trial design using value of information methods with imperfect implementation. *Health Econ* 19(5):549–561
63. Smith JE (1993) Moment methods for decision analysis. *Manag Sci* 39(3):340–358
64. Gospodinov N, Lkhagvasuren D (2014) A moment-matching method for approximating vector autoregressive processes by finite-state Markov chains. *J Appl Econom* 29(5):843–859
65. Weber TA (2017) Optimal switching between cash-flow streams. *Math. Method. Oper. Res.* Forthcoming. <https://doi.org/10.1007/s00186-017-0586-0>
66. Blackwell D (1951) Comparison of experiments. In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press. pp 93–102
67. Kim W (2002) The burden of hepatitis C in the United States. *Hepatology* 36(Suppl 1):S30–S34
68. Ghany MG, Strader DB, Thomas DL, Seeff LB (2009) Diagnosis, management, and treatment of hepatitis C: an update. *Hepatology* 49(4):1335–1374
69. Armstrong GL, Wasley A, Simard EP, McQuillan GM, Kuhnert WL, Alter MJ (2006) The prevalence of hepatitis C virus infection in the United States, 1999 through 2002. *Ann Intern Med* 144(10):705–714
70. Chak E, Talal AH, Sherman KE, Schiff ER, Saab S (2011) Hepatitis C virus infection in USA: an estimate of true prevalence. *Liver Int* 31(8):1090–1101
71. Denniston MM, Klevens RM, McQuillan GM, Jiles RB (2012) Awareness of infection, knowledge of hepatitis C, and medical follow-up among individuals testing positive for hepatitis C: National Health and Nutrition Examination Survey 2001–2008. *Hepatology* 55(6):1652–1661
72. Armstrong GL (2007) Injection drug users in the United States, 1979–2002: an aging population. *Arch Intern Med* 167(2):166–173
73. Armstrong GL, Alter MJ, McQuillan GM, Margolis HS (2000) The past incidence of hepatitis C virus infection: implications for the future burden of chronic liver disease in the United States. *Hepatology* 31(3):777–782
74. Joy JB, McCloskey RM, Nguyen T, Liang RH, Khudyakov Y, Olmstead A, Krajden M, Ward JW, Harrigan PR, Montaner JS, Poon AF (2016) The spread of hepatitis C virus genotype 1a in North America: a retrospective phylogenetic study. *Lancet Infect Dis* 16(6):698–702
75. Barker L Personal communication, August 2, 2016. Based on analyses using the Centers for Disease Control and Prevention, National Health and Nutrition Examination Survey (NHANES) (2005–2012)
76. Klevens RM, Hu DJ, Jiles R, Holmberg SD (2012) Evolving epidemiology of hepatitis C virus in the United States. *Clin Infect Dis Suppl* 55(1):S3–9
77. American Association for the Study of Liver Diseases and the Infectious Diseases Society of America (AASLD-IDSA). HCV testing and linkage to care. Recommendations for testing, managing, and treating hepatitis C. Available at: <http://www.hcvguidelines.org/full-report/hcv-testing-and-linkage-care>. Accessed: February 23, 2017
78. R Development Core Team (2012) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna. <http://www.R-project.org/>
79. US Census Bureau (2010) QT – P1: Age groups and sex: 2010. Available at: <http://www.census.gov/2010census/>
80. Mehrotra A, Zaslavsky AM, Ayanian JZ (2007) Preventive health examinations and preventive gynecological examinations in the United States. *Arch Intern Med* 167(17):1876
81. Gretch DR (1997) Diagnostic tests for hepatitis C. *Hepatology* 26(Suppl 3):43S–47S
82. Hyland C, Kearns S, Young I, Battistutta D, Morgan C (1992) Predictive markers for hepatitis C antibody ELISA specificity in Australian blood donors. *Transfusion Med* 2(3):207–213
83. Center for Medicare and Medicaid Services (CMS). 2010. Medicare fee schedule. U.S. Department of Health and Human Services. <http://www.cms.gov/home/medicare.asp>
84. Weinstein MC, Skinner JA (2010) Comparative effectiveness and health care spending – implications for reform. *New Engl J Med* 362(5):460–465
85. Thein HH, Yi Q, Dore GJ, Krahn MD (2008) Estimation of stage-specific fibrosis progression rates in chronic hepatitis C virus infection: A meta-analysis and meta-regression. *Hepatology* 48(2):418–431
86. Centers for Disease Control and Prevention (CDC) (2006) National Health and Nutrition Examination Survey (NHANES):

- analytic and reporting guidelines. http://www.cdc.gov/nchs/nhanes/nhanes2003-2004/analytical_guidelines.htm. Accessed: August 27, 2012. Last updated: September 2006
87. Centers for Disease Control and Prevention (CDC) (2011) Analytic note regarding 2007-2010 survey design changes and combining data across other survey cycles. http://www.cdc.gov/nchs/nhanes/nhanes2003-2004/analytical_guidelines.htm. Accessed: August 27, 2012. Last updated: September 2011
88. Centers for Disease Control and Prevention (CDC) (2012) National Health and Nutrition Examination Survey (NHANES) (1999–2010). <http://www.cdc.gov/nchs/nhanes.htm>. Accessed: August 27, 2012
89. Milgrom PR (1981) Good news and bad news: Representation theorems and applications. *Bell J Econ* 12(12):380–391
90. US Bureau of Labor Statistics (2011) Consumer Price Index (CPI):1913–present. Division of consumer prices and price indexes. Available at: <http://www.bls.gov/cpi/>