Original Manuscript

MDM Medical Decision Making

Medical Decision Making
1–13

DOI: 10.1177/0272989X18793401
journals.sagepub.com/home/mdm
SAGE

Optimal Information Collection Policies in a Markov Decision Process Framework

Lauren E. Cipriano, Jeremy D. Goldhaber-Fiebert, Shan Liu, and Thomas A. Weber

Abstract

Background. The cost-effectiveness and value of additional information about a health technology or program may change over time because of trends affecting patient cohorts and/or the intervention. Delaying information collection even for parameters that do not change over time may be optimal. **Methods.** We present a stochastic dynamic programming approach to simultaneously identify the optimal intervention and information collection policies. We use our framework to evaluate birth cohort hepatitis C virus (HCV) screening. We focus on how the presence of a time-varying parameter (HCV prevalence) affects the optimal information collection policy for a parameter assumed constant across birth cohorts: liver fibrosis stage distribution for screen-detected diagnosis at age 50. **Results.** We prove that it may be optimal to delay information collection until a time when the information more immediately affects decision making. For the example of HCV screening, given initial beliefs, the optimal policy (at 2010) was to continue screening and collect information about the distribution of liver fibrosis at screen-detected diagnosis in 12 years, increasing the expected incremental net monetary benefit (INMB) by $169.5 million compared to current guidelines. **Conclusions.** The option to delay information collection until the information is sufficiently likely to influence decisions can increase efficiency. A dynamic programming framework enables an assessment of the marginal value of information and determines the optimal policy, including when and how much information to collect.

Keywords

Markov decision processes, value-of-information analysis, hepatitis C, optimal stopping, dynamic programming

Date received: March 26, 2018; accepted: July 11, 2018

Population health trends, such as increasing rates of obesity,[1] diabetes,[2,3] and cancer as well as changes in the distribution of cancer subtypes,[4] may influence the cost-effectiveness of health interventions. For example, the cost-effectiveness of lung cancer screening may be affected by decreasing smoking rates[5,6] or changes in the cost and efficacy of treatment.[7] Many drug prices decline over time, which can improve the cost-effectiveness of treatment.[8] Increasing adoption rates and coverage of the human papillomavirus (HPV) vaccine reduce HPV prevalence, thereby altering the cost-effectiveness of cytological screening.[9,10] In addition to assessing current cost-effectiveness, it is important to consider when it may become cost-effective to initiate a program that is currently not cost-effective or to terminate a program that will no longer be cost-effective. Such decisions often depend on an assessment of when, what kind, and how much information it is cost-effective to collect.[11] The optimal timing of information collection is clearly a factor for time-varying parameters, but the importance of timing for parameters that do not vary in time is less clear.

Expected value of sample information (EVSI) is increasingly advocated to determine optimal sample size and the societal return of proposed research.[11–20]

Corresponding Author
Lauren E. Cipriano, Ivey Business School, Western University, 1255 Western Road, London, Ontario, N6G 0N1, Canada.
(lcipriano@ivey.uwo.ca)

Current practice in health decision science assumes model parameters are constant across cohorts, and the value of additional information is calculated assuming the information collection effort occurs immediately.[16,21,22] Such static models may often provide good estimates of the cost-effectiveness of an intervention and the per-person value of information in current and near-term cohorts. However, the cost-effectiveness of a health technology or program—and, therefore, the value of collecting additional information about 1 or more model parameters—may be changing over time because of trends affecting the cohort, the intervention, or both.[8,23–25] Over long planning horizons, ignoring these trends in EVSI computations may result in recommending the collection of too much or too little information. Furthermore, collecting additional information immediately may not be the optimal strategy. Our study explores the importance of extending EVSI methods to consider the value of delaying information collection, even when the parameter of interest is not varying in time, showing that the presence of another time-varying parameter makes such an approach necessary over long planning horizons.

We consider a hypothetical medical intervention decision for a cohort of patients where one imperfectly observable parameter decreases across intervention cohorts. Our motivating example is hepatitis C virus (HCV) screening in 50-year-olds at a routine medical visit, where the prevalence of HCV is decreasing across successive cohorts of 50-year-olds. At each time, the policy maker can choose to screen the current cohort of 50-year-olds and to purchase sample information about parameters influencing the decision. Using a stochastic dynamic programming approach, we have previously shown that it may be optimal to delay information collection about the time-varying parameter (HCV prevalence) until the information is more likely to result in a change in policy.[26]

This article extends our previous framework to highlight the importance of considering delayed information collection in EVSI calculations even for time-invariant parameters to maximize the efficient use of research resources. First, we prove that in the presence of discounting (time value of money), it may be optimal to delay the collection of information for a time-invariant parameter (those that do not change across intervention cohorts). We illustrate the practicability and feasibility of applying our framework with the example of HCV screening using a previously published cost-effectiveness model.[27,28]

Application to HCV Screening

Affecting approximately 5 million Americans, chronic HCV infection is a slowly progressing disease causing liver fibrosis, cirrhosis, and liver cancer.[29,30] Two-thirds of chronically infected individuals were born between 1945 and 1965, and approximately half are unaware of their disease.[30,31] The advent of more effective therapy has changed the value of identifying infected individuals,[27,32–35] and as a result, the Centers for Disease Control and Prevention (CDC) and the US Preventive Services Task Force (USPSTF) recommended one-time screening of these cohorts in 2013.[36,37] HCV prevalence, and therefore the value of screening, is decreasing in post-1956 birth cohorts (Figure 1). The optimal duration for an HCV screening program remains uncertain; several studies indicate that screening individuals born after 1965 is likely cost-effective.[27,34,35]

Several highly uncertain model parameters influence the cost-effectiveness of HCV screening[27,32–35] and, therefore, the optimal duration of screening. Collecting additional information about these model parameters may therefore be valuable to inform this decision. We apply our framework to identify the optimal information collection policy for an example parameter that we assume is constant across birth cohorts—the liver fibrosis stage distribution at screen-detected diagnosis—alone and in combination with the opportunity to collect information about HCV prevalence, which is decreasing across birth cohorts. We compare the optimal information-collection policy to the standard approach for value-of-information assessment, focusing on how the optimal information collection policy is influenced by the presence of time-varying parameters. Using this example, we show that it may be optimal to delay information collection about model parameters that are not changing across cohorts.

Ivey Business School, Western University, London, ON, Canada (LEC); Center for Health Policy and Primary Care and Outcomes Research, Stanford University, Stanford, CA, USA (JDG-F); Industrial & Systems Engineering, College of Engineering, University of Washington, Seattle, WA, USA (SL); and Operations, Economics and Strategy, College of Management of Technology, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland (TAW).
Financial support for this study was provided in part by a doctoral fellowship from the Social Science and Humanities Research Council of Canada, a Discovery grant from the Natural Science and Engineering Research Counsel of Canada, and the Seth Bonder Scholarship for Applied Operations Research in Health Services from the Institute for Operations Research and the Management Sciences. Dr. Goldhaber-Fiebert was supported in part by a National Institutes of Health National Institute on Aging Career Development Award (K01 AG037593-01A1; principal investigator, Dr. Goldhaber-Fiebert). The authors had independence in designing the study, interpreting the data, writing, and publishing the report.

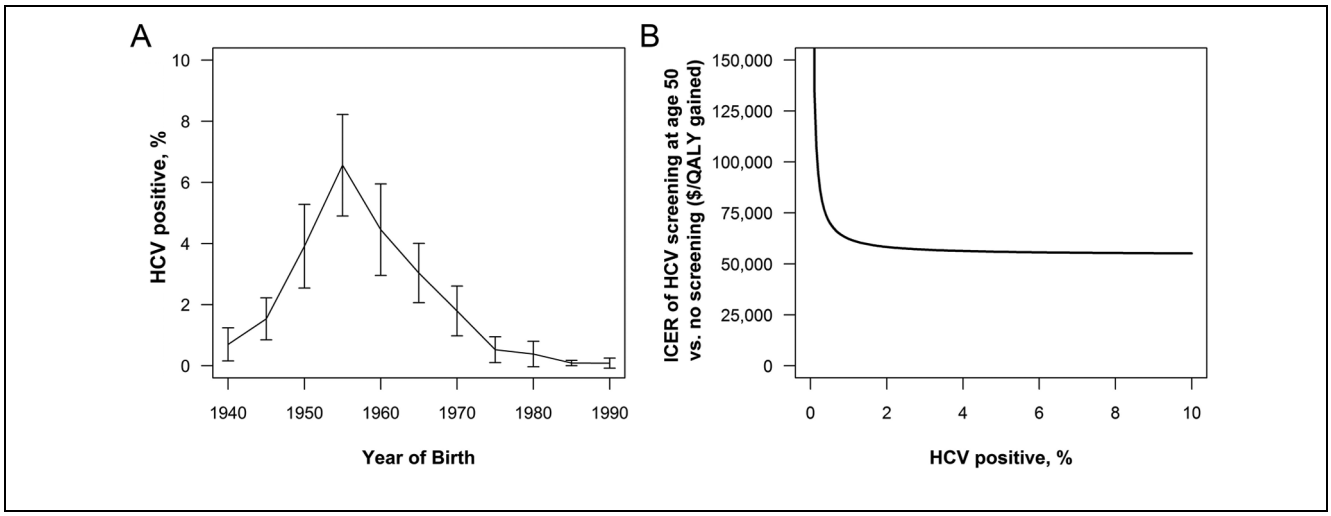

Figure 1 (A) Hepatitis C virus (HCV) prevalence in US men by birth year. Estimated using the National Health and Nutrition Examination Survey (NHANES) (1999–2010). (B) Incremental cost-effectiveness ratio (ICER) of screening for HCV at a routine preventive health exam at age 50 compared to a policy of no routine screening in incremental cost per quality-adjusted life year (QALY) gained. Estimated using the cost-effectiveness model developed by Liu et al.[27,28]

Methods

The General Model

The general model has been described previously.[26] A policy maker faces recurring decisions at times $t \in \{0, 1, 2, \ldots\}$ about whether to stop investment in a health intervention delivered at most once per cohort (of size $N_t > 0$). The policy maker's objective is to maximize the (expected) net monetary benefit. The incremental net monetary benefit (INMB) of the intervention compared to the next-best alternative is linear in an uncertain time-varying probability $\tilde{p}_t \in [0, 1]$,

$$\mathrm{INMB}_t = \tilde{\theta}_t \tilde{p}_t - \gamma,$$

where $\tilde{\theta}_t$ and γ represent the marginal and fixed INMB with respect to $\tilde{p}_t$, respectively.

The policy maker's prior belief about $\tilde{p}_t$ at time t follows a beta-distribution with parameters $x_t = (a_t, b_t)$, which describe the current belief state. The terms $\mu_p(x_t)$ and $\sigma_p(x_t)$ denote the mean and standard deviation of $\tilde{p}_t$, respectively. We assume that the distribution $\tilde{p}_t$ is changing over time with dynamics aimed at approximating the prevalence of a health condition that is becoming less common over time. Specifically, we assume

$$\mu_p(x_{t+1}) = z\mu_p(x_t) \quad \text{and}$$
$$\sigma_p^2(x_{t+1}) = \sigma_p^2(x_t)\left(z + z(1-z)\left(\frac{\mu_p(x_t)}{1-\mu_p(x_t)}\right)\right), \qquad (1)$$

where $z \in [0, 1]$ is a known decay rate. These dynamics imply a geometrically decreasing mean, decreasing variance (as long as $\mu_p(x_t) \leq \frac{1}{1+z}$), and an increasing coefficient of variation σ_p/μ_p.

We assume that the true values of $\theta > 0$ and $\gamma > 0$ are constant over time. The true value of θ is uncertain, and we denote our current belief, $\tilde{\theta}_t(y_t)$, with probability distribution parameters y_t, expectation $\mu_\theta(y_t)$, and standard deviation $\sigma_\theta(y_t)$. In the application to HCV screening, $\tilde{p}_t$ is the (decreasing) disease prevalence in the t-th cohort, $\tilde{\theta}_t$ is the benefit of early diagnosis and treatment initiation for an affected individual, and γ is the per-person cost of the screening program.

At each time t, the policy maker faces 3 simultaneous decisions. First, there is the decision $d_t \in \{0, 1\}$ to invest in the health intervention for cohort t, with $d_t = 0$ indicating "no intervention" and $d_t = 1$ indicating "intervention." The policy maker also has the option to collect sample information about $\tilde{p}_t$ and $\tilde{\theta}_t$ with sample of sizes $n_t \geq 1$ and $m_t \geq 1$ at (nonnegative and increasing) costs $\kappa_p(n_t)$ and $\kappa_\theta(m_t)$, respectively. Therefore, at each time t, the decision maker implements the control $u_t = (d_t, n_t, m_t)$ in the control-constraint set $\mathcal{U}$ containing the following elements: $u_t = (0, 0, 0)$, "no intervention (and do not sample)"; $u_t = (1, 0, 0)$, "do intervention and do not sample"; $u_t = (1, n_t, 0)$, "do intervention and sample n_t individuals to learn about $\tilde{p}_t$"; $u_t = (1, 0, m_t)$, "do intervention and sample m_t individuals to learn about $\tilde{\theta}_t$"; and $u_t = (1, n_t, m_t)$, "do intervention, sample n_t individuals to learn about $\tilde{p}_t$, and sample m_t individuals to learn about

$\tilde{\theta}_t$." Information arrives at the end of the current period and can only be used, in combination with the dynamics on $\tilde{p}_t$, to better inform the decision for the next cohort.

The current-period reward is then

$$g(\tilde{p}_t, \tilde{\theta}_t, u_t) = d_t(\tilde{\theta}_t \tilde{p}_t - \gamma) - \kappa_p(n_t) - \kappa_\theta(m_t). \quad (2)$$

We assume $\tilde{\theta}_t$ and $\tilde{p}_t$ are independent, so the expected reward for choosing "intervention" for the t-th cohort is

$$\mathbb{E}[g(\tilde{p}_t, \tilde{\theta}_t, u_t)|x_t, y_t, d_t = 1] = \mu_\theta(y_t)\mu_p(x_t) - \gamma - \kappa_p(n_t) - \kappa_\theta(m_t).$$

We assume a decision to discontinue the current health intervention applies to all future cohorts, so

$$\mathbb{E}[g(\tilde{p}_t, \tilde{\theta}_t, u_t)|x_t, y_t, d_t = 0] = 0 \text{ for all } t \geq T \text{ where } d_T = 0.$$

Information Collection Problem

We assume that the parametric distributions of $\tilde{p}_t$ and $\tilde{\theta}_t$ have conjugate-pair distributions (e.g., the beta and the beta-binomial distributions).[38] Therefore, the information states (captured by the parameters x_t and y_t) can be updated in a Bayesian manner while maintaining the parametric distributions of $\tilde{p}_t$ and $\tilde{\theta}_t$ over time. For a study about $\tilde{p}_t$, we denote the number of positive samples $\tilde{v}_t \in \{0, \ldots, n_t\}$ and the belief state of the posterior beta-distribution by $\hat{x}_t$ where $\hat{x}_t = \psi_x(x_t, v_t, n_t)$. For a study about $\tilde{\theta}_t$, we denote the sample information $\tilde{w}_t$ and the belief state of the posterior distribution $\hat{y}_t$, where $\hat{y}_t = \psi_y(y_t, w_t, m_t)$.

We assume that a study about $\tilde{p}_t$ does not provide any information about $\tilde{\theta}_t$ and vice versa. We assume, in general, the n_t individuals sampled to learn about $\tilde{p}_t$ and the m_t individuals sampled to learn about $\tilde{\theta}_t$ represent nonoverlapping groups because the criteria for inclusion are different. For example, to learn about disease prevalence, a study may randomly sample the population including individuals unaware of their infection status, whereas a study to learn about treatment effectiveness would only include individuals known to have the disease.

Mathematical Representation of the Policy Maker's Problem

The optimal policy $\pi^*(\cdot)$ determines the optimal mapping π^* from states (x_t, y_t) to actions $u_t^* = (d_t^*, n_t^*, m_t^*) = \pi^*(x_t, y_t)$, solving the optimal control problem

$$\max_{\pi(\cdot)} \sum_{t=0}^{\infty} \delta^t \mathbb{E}\left[g(\tilde{p}_t, \tilde{\theta}_t, \pi(x_t, y_t))\middle|\mathcal{I}_t\right]$$
$$\text{subject to} \quad x_{t+1} = \phi(\psi_x(x_t, \tilde{v}_t, u_t)), \quad x_0 \text{ given},$$
$$y_{t+1} = \psi_y(y_t, \tilde{w}_t, u_t), \quad y_0 \text{ given},$$
$$u_t = \pi(x_t, y_t) \in \mathcal{U}$$

where $\delta \in (0, 1)$ denotes the per-period discount factor, $\phi(\cdot)$ denotes the function generating the system dynamics (see equation (1)), $\psi_x(\cdot)$ denotes the Bayesian updating of x_t, $\psi_y(\cdot)$ denotes the Bayesian updating of y_t, and $\mathcal{I}_t$ refers to all the information available up to time t. Given the optimal policy $\pi^*(\cdot)$, the value function $V(x_t, y_t)$ satisfies Bellman's equation,[39]

$$\begin{aligned} V(x_t, y_t) &= \max_{u_t \in \mathcal{U}} \{d_t(\mu_\theta(y_t)\mu_p(x_t) - \gamma) - \kappa_x(n_t) - \kappa_y(m_t) \\ &\quad + \delta \mathbb{E}[V(\phi(\psi_x(x_t, \tilde{v}_t, u_t)), \psi_y(y_t, \tilde{w}_t, u_t))]\} \\ &= d_t^*(\mu_\theta(y_t)\mu_p(x_t) - \gamma) - \kappa_x(n_t^*) - \kappa_y(m_t^*) \\ &\quad + \delta \mathbb{E}[V(\phi(\psi_x(x_t, \tilde{v}_t, u_t^*)), \psi_y(y_t, \tilde{w}_t, u_t^*))], \end{aligned} \quad (3)$$

where the first 3 terms represent the expected current period reward and the fourth terms represents the expected present value of the optimal policy in all future periods. For the special case in which information is not available about either $\tilde{p}_t$ or $\tilde{\theta}_t$, we present the optimal policy, specifically the optimal time to stop the intervention, in online Appendix B.1.

Application to HCV Screening

We consider the repeated decision problem of whether or not to perform one-time HCV screening of men at a routine medical appointment at age 50 for successive birth cohorts beginning with the 1960 birth cohort (Figure 2). We consider screening at age 50 because one-time screening at this age had the lowest incremental cost-effectiveness ratio in a cost-effectiveness analysis of single birth cohort screening.[27] At the beginning of each period, the policy maker must simultaneously decide whether to screen the current cohort of 50-year-olds for HCV and whether to conduct a study to better estimate current HCV prevalence, a study of an observable parameter that will help to better estimate the marginal benefit of early diagnosis and treatment for an individual with HCV, or to conduct both possible studies. Information arrives at the end of the current period and can be used (in combination with the prevalence dynamics) to better inform the screening decision for the next cohort.

The current-period INMB, $\tilde{\theta}_t \tilde{p}_t - \gamma$, represents the INMB of screening at age 50 compared to not screening in a cohort with prevalence $\tilde{p}_t$, where

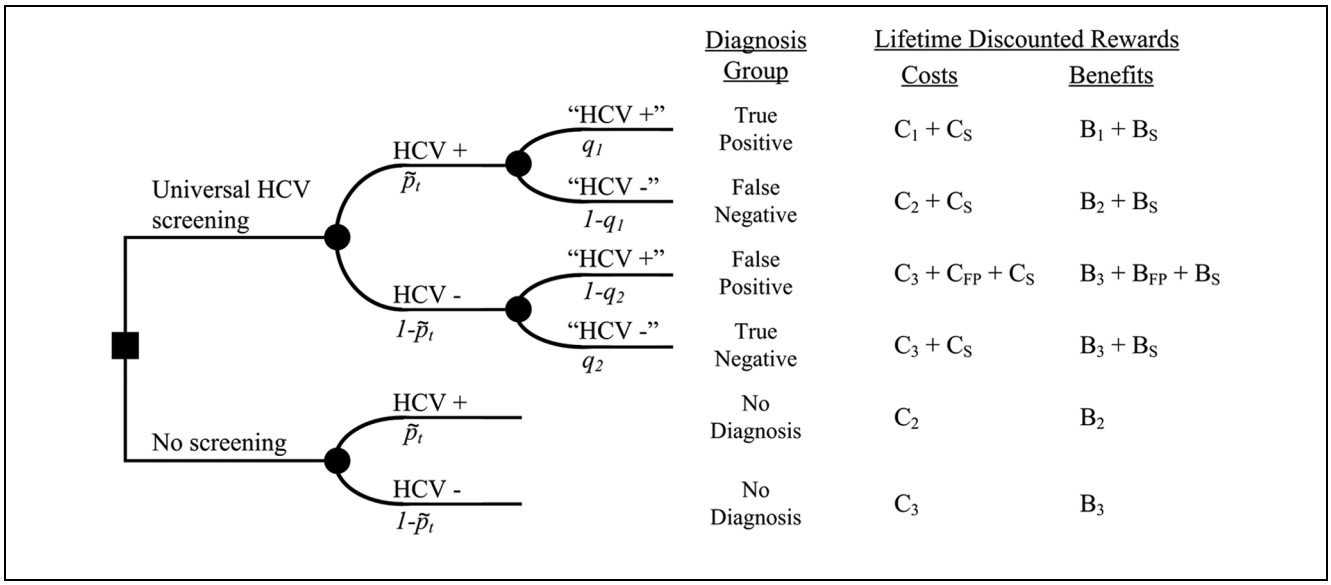

Figure 2 A schematic of the medical intervention decision. Each year, the policy maker must choose whether to screen for hepatitis C virus (HCV) infection in 50-year-olds who attend a routine preventive health exam. The HCV prevalence in the t-th cohort is $\tilde{p}_t$. The screening test has sensitivity q_1 and specificity q_2. The screening test costs $C_S > 0$ and screening test invasiveness may result in a quality-of-life loss, $B_S \leq 0$. An initially false-positive test result is corrected at a cost $C_{FP} \geq 0$ and quality-of-life loss, $B_{FP} \leq 0$. The lifetime discounted costs and benefits of the true-positive, false-negative, and true-negative screening outcomes are denoted by C_1, C_2, C_3 and B_1, B_2, B_3, respectively.

$$\tilde{\theta}_t = q_1[\lambda(B_1 - B_2) - (C_1 - C_2)] - (1 - q_2)[\lambda B_{FP} - C_{FP}], \tag{4}$$

and

$$\gamma = C_S - \lambda B_S - (1 - q_2)(\lambda B_{FP} - C_{FP}), \tag{5}$$

where λ is the willingness-to-pay threshold, q_1 and q_2 are the sensitivity and specificity of the screening test, $C_S > 0$ and $B_S \leq 0$ are the cost and quality-of-life loss associated with performing the screening test, $C_{FP} \geq 0$ and $B_{FP} \leq 0$ are the cost and quality-of-life loss associated with correcting a false-positive test result, and the lifetime discounted costs and benefits of the true-positive, false-negative, and true-negative screening outcomes are denoted by C_1, C_2, C_3 and B_1, B_2, B_3, respectively (derivations presented in online Appendix A.1). Because $\tilde{\theta}_t$, the marginal benefit of early diagnosis and treatment for an individual with HCV is not directly measurable, reducing uncertainty about $\tilde{\theta}_t$ requires information on factors that contribute to it, such as treatment efficacy, treatment adherence, access to treatment, and the distribution of disease severity at diagnosis (the example presented herein).

We estimated the base-case parameter values and ranges for sensitivity analysis from the National Health and Nutrition Examination Survey (NHANES), a previously published model developed to evaluate the cost-effectiveness of HCV screening in a single cohort,[27] and the medical literature. Parameter values are presented in Table 1. We assumed that the cohort size, the number of men who attend a preventive health exam at age 50, $N_t = Q$, is constant over time since there is less than 10% variation across cohorts currently 25 to 55 years old.[40] We assumed the cost of sample information about HCV prevalence is linear in sample size with a fixed cost of \$50,000 and a variable cost of \$100 for each participant. The HCV screening test has a sensitivity of 0.97 and a specificity of 0.9996[41,42]; therefore, the posterior distribution of $\tilde{p}_t$ after Bayesian updating is a mixture of different beta-distributions. To find stationary policies, we approximated the true posterior mixture distribution with a single beta-distribution with the same mean and variance.

We adopted a societal perspective, considered costs and benefits over a lifetime horizon for each cohort, and discounted future costs and benefits at 3% annually.[43] Costs are expressed in 2010 US dollars and inflation

Table 1 Base-Case Parameter Values and Sources

Variable, Description		Value[a]	Sources
Annual cohort: Men aged 50 years attending a PHE			
	Number eligible for a PHE	2.1 million	40
	Proportion who attend a PHE	24.4%	62
Q	Annual number of PHEs	508,222	Calculated
HCV screening test			
q_1	Sensitivity	0.97	41
q_2	Specificity	0.9996	42
C_S	Cost of screening test, HCV antibody test	$28	49
C_{FP}	Cost of false positive	$230	49
B_S	Quality-of-life change, event of screening	0	Assumed
B_{FP}	Quality-of-life change, false-positive result	0	Assumed
Lifetime discounted costs per person ($)			
C_1	HCV +, identified through screening	$\$141{,}675\tilde{F}_{0,t} + \$146{,}656\tilde{F}_{1,t} + \$149{,}961\tilde{F}_{2,t} + \$149{,}595\tilde{F}_{3,t} + \$148{,}104\tilde{F}_{4,t}$	27
C_2	HCV +, not identified through screening	$\$126{,}784\tilde{F}_{0,t} + \$128{,}625\tilde{F}_{1,t} + \$128{,}079\tilde{F}_{2,t} + \$125{,}498\tilde{F}_{3,t} + \$120{,}398\tilde{F}_{4,t}$	27
C_3	HCV– individual	$181,314	27
Lifetime discounted QALYs per person			
B_1	HCV +, identified through screening	$11.19\tilde{F}_{0,t} + 10.82\tilde{F}_{1,t} + 10.27\tilde{F}_{2,t} + 9.71\tilde{F}_{3,t} + 8.75\tilde{F}_{4,t}$	27
B_2	HCV +, not identified through screening	$11.09\tilde{F}_{0,t} + 10.53\tilde{F}_{1,t} + 9.81\tilde{F}_{2,t} + 9.19\tilde{F}_{3,t} + 8.06\tilde{F}_{4,t}$	27
B_3	HCV– individual	15.69	27
INMB per person ($)			
$\tilde{\theta}_t$	Variable component of INMB	$-\$6857\tilde{F}_{0,t} + \$3660\tilde{F}_{1,t} + \$12{,}708\tilde{F}_{2,t} + \$14{,}123\tilde{F}_{3,t} + \$23{,}059\tilde{F}_{4,t}$	Equation (4)
γ	Fixed component of INMB	$28.05	Equation (5)
Initial belief			
$\tilde{p}_t(x_0)$	HCV prevalence in undiagnosed individuals (1960 birth cohort in 2010)	$x_0 = (a_0,b_0) = (75.1,2350.5)$; $\mu_p(x_0) = 0.031$; $\sigma_p(x_0) = 0.0035$	NHANES; see Cipriano and Weber[26]
$\tilde{F}_t(y_0)$	Screen-detected fibrosis stage distribution	$y_0 = (5,20,5,4,5)$; $\tilde{F}_0(y_0) = (0.128,0.513,0.128,0.103,0.128)$	48
Cost of collecting information			
For a study about $\tilde{p}_t$			
$\kappa_{x,F}$	Fixed cost per sampling study	$50,000	Estimated
$\kappa_{x,V}$	Variable cost per sample	$100	49
$\kappa_x(n_t)$	Cost of sampling, per person in the cohort	$((\kappa_{x,F} + n_t\kappa_{x,V})/Q)\mathbf{1}_{\{n_t>0\}}$	Assumed
For a study about $\tilde{F}_t$			
$\kappa_{y,F}$	Fixed cost per sampling study	$400,000	Estimated
$\kappa_{y,V}$	Variable cost per sample	$2000	49
$\kappa_y(m_t)$	Cost of sampling, per person in the cohort	$\big((\kappa_{y,F} + m_t\kappa_{y,V})/Q\big)\mathbf{1}_{\{m_t>0\}}$	Assumed
Other			
z	Rate of prevalence decay	0.893	26
λ	Willingness-to-pay threshold	$75,000/QALY	45
r	Annual discount rate	0.03	43

HCV, hepatitis C virus; INMB, incremental net monetary benefit; NHANES, National Health and Nutrition Examination Survey; PHE, preventive health exam; QALY, quality-adjusted life year.

a. The lifetime discounted costs, QALYs, and INMB of screening at age 50 compared to not screening per person are linear functions of $\tilde{F}_t$. We estimated the parameters of this function using a previously published model of HCV screening.[27] Using $\tilde{\theta}_t = -\$6857\tilde{F}_{0,t} + \$3660\tilde{F}_{1,t} + \$12{,}708\tilde{F}_{2,t} + \$14{,}123\tilde{F}_{3,t} + \$23{,}059\tilde{F}_{4,t}$: specifically, -$6857 is the INMB of screening for HCV at age 50 compared to not screening for an individual in fibrosis stage F0 at diagnosis, $3660 is the INMB for an individual in fibrosis stage F1 at diagnosis, $12,708 is the INMB for an individual in fibrosis stage F2 at diagnosis, $14,123 is the INMB for an individual in fibrosis stage F3 at diagnosis, and $23,059 is the INMB for an individual in fibrosis stage F4 at diagnosis.

adjusted with the Consumer Price Index.[44] Benefits are measured in quality-adjusted life years (QALYs). For the purpose of illustration, we assumed a mid-range value for society's maximum willingness to pay of \$75,000 per QALY gained.[45]

Example: Liver fibrosis stage distribution at screen-detected diagnosis. An HCV-infected individual's degree of liver fibrosis affects the incremental lifetime costs and benefits of early diagnosis and treatment[46]; therefore, the liver fibrosis stage distribution in the screen-detected population affects the cost-effectiveness of population screening.[27,47] Variation in estimates of the liver fibrosis stage distribution has led to widely different estimates for the cost-effectiveness of HCV screening.[32–35]

We classified fibrosis stages using Metavir scores of F0 (no fibrosis) through F4 (cirrhosis). We denote the distribution of screen-detected individuals across the fibrosis stages $\tilde{F}_t = (\tilde{F}_{0,t}, \tilde{F}_{1,t}, \tilde{F}_{2,t}, \tilde{F}_{3,t}, \tilde{F}_{4,t})$, where $0 \leq \tilde{F}_{0,t}, \tilde{F}_{1,t}, \tilde{F}_{2,t}, \tilde{F}_{3,t}, \tilde{F}_{4,t} \leq 1$ and $\sum_{i=0}^{4} \tilde{F}_{i,t} = 1$. We decomposed the marginal benefit of early HCV diagnosis and treatment, $\tilde{\theta}_t$, into a linear function of $\tilde{F}_t$ (online Appendix A.2). We estimated the marginal benefit of early diagnosis and treatment conditional on being diagnosed at each Metavir score using a previously published cost-effectiveness model[27] (Table 1).

We assumed that the policy maker's belief about $\tilde{F}_t$ at time t is Dirichlet distributed with parameters $y_t = (y_{0,t}, y_{1,t}, y_{2,t}, y_{3,t}, y_{4,t})$. As liver biopsies pose risks for serious adverse events, information on the liver fibrosis stage distribution among the screening-eligible population is scarce. The prior belief was therefore obtained using the count of individuals in each fibrosis stage in a study of risk factor–based screening in the Veterans Administration (VA) health care system, which identified 122 individuals with chronic HCV infection, 39 of whom underwent biopsy for fibrosis staging: $y_0 = (5, 20, 5, 4, 5)$.[48]

We assumed that the cost of a fibrosis stage distribution study in newly screen-detected individuals using ultrasound-guided liver biopsy is linear in sample size with fixed costs of \$400,000 (Barrett Levesque, personal communication 2013) and variable costs of \$2000 per study participant.[49] Therefore, a study with a sample size of 200 costs \$800,000; this cost is varied in a sensitivity analysis.

The number of sampled individuals in each category, $\tilde{w}_t = (w_{0,t}, w_{1,t}, w_{2,t}, w_{3,t}, w_{4,t})$, where $\sum_{i=0}^{4} w_{i,t} = m_t$, is Dirichlet-multinomial distributed. The policy maker updates her or his belief about $\tilde{F}_t$ in a Bayesian manner, resulting in a Dirichlet-distributed posterior belief where $y_{t+1} = (y_{0,t} + w_{0,t}, y_{1,t} + w_{1,t}, y_{2,t} + w_{2,t}, y_{3,t} + w_{3,t}, y_{4,t} + w_{4,t})$. With 5 fibrosis levels, there are $\binom{m_t + 5 - 1}{5 - 1} = 70{,}058{,}751$ possible outcomes from a study where $m_t = 200$, although many combinations yield very similar cost-effectiveness results. For computational feasibility, we used recursive partitioning regression to identify 24 classes of similar ("representative") posterior distributions, with the similarity metric defined in terms of the INMB of HCV screening[50] (details in online Appendix A.3). The expected value of the option to collect information about $\tilde{F}_t$ is obtained as the value of the optimal action given each possible representative posterior distribution weighted by the probability of that posterior.

Comparative Analysis

We began our analysis for year 2010 with the first cohort of screened individuals at age 50 born in 1960. For each example, we identified the optimal intervention and information collection policy given current beliefs comparing our approach to other EVSI approaches. First, we computed the expected INMB of the current recommendation to screen until the 1965 birth cohort. Second, we calculated the optimal intervention policy assuming that no information was available about either $\tilde{p}_t$ or $\tilde{F}_t$. Third, since the traditional EVSI approach is to only consider initiating information collection about model parameters immediately,[16,21,51] we considered the opportunity to collect information immediately about either or both parameters. Fourth, we identified the optimal intervention and information collection policy if information collection is available for one or both parameters and can be delayed using value iteration in 2 stages implemented in R version 2.15.0[39,52] (details in online Appendix A.4).

Results

In Cipriano and Weber,[26] we established properties of the optimal policy and of the value function when 1) information cannot be collected about $\tilde{p}_t$ or $\tilde{\theta}_t$ and 2) information can be collected about $\tilde{p}_t$ but not $\tilde{\theta}_t$, which continue to hold in the case of an uncertain $\tilde{\theta}_t$. We now consider 2 new cases: when information is available about $\tilde{\theta}_t$ but not $\tilde{p}_t$ and when information is available about both $\tilde{p}_t$ and $\tilde{\theta}_t$. First, we prove that, even though the true value of $\tilde{\theta}_t$ is constant, it may be optimal to delay information collection about $\tilde{\theta}_t$ until the information is more likely to change an immediate decision.

Table 2 Comparison of Optimal Policies When Information Is Available for $\tilde{p}_t$, $\tilde{F}_t$, or Both, Now Only or in All Future Periods Given Our Initial Belief about HCV Prevalence in Men Born in 1960, $\mu_p(x_0) = 0.031$ and $\sigma_p(x_0) = 0.0035$, and the Screen-Detected HCV Fibrosis Stage Distribution $F_0 = (0.128, 0.513, 0.128, 0.103, 0.128)$

Case	Optimal Policy	Expected INMB[a]	Increase in Expected INMB
CDC recommendation[b]	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			
$\tilde{p}_t$ only	$n_0 = 920$, then identify optimal stopping time	$566,490,000	$167,350,000
$\tilde{F}_t$ only	$m_0 = 200$, then identify optimal stopping time	$566,866,000	$167,726,000
$\tilde{p}_t$ and $\tilde{F}_t$	$(n_0, m_0) = (920, 200)$, then identify optimal stopping time	$566,867,000	$167,727,000
Information available in all periods[c]			
$\tilde{p}_t$ only	$n_{16} = 4000$ from the 1976 birth cohort in 16 years, then identify optimal next action	$567,940,000	$168,800,000
$\tilde{F}_t$ only	$m_{14} = 200$ from the 1974 birth cohort in 14 years, then identify optimal next action	$566,920,000	$167,780,000
$\tilde{p}_t$ and $\tilde{F}_t$	$m_{12} = 200$ from the 1972 birth cohort in 12 years, then identify optimal next action	$568,590,000	$169,450,000

CDC, Centers for Disease Control and Prevention; INMB, incremental net monetary benefit.

a. The expected value gained by a hepatitis C virus (HCV) screening program is the sum of the discounted INMB of screening compared to not screening over all future periods.

b. The CDC and US Preventive Services Task Force recommendation is to screen all individuals born between 1945 and 1965 for HCV at their next routine medical visit.[36,37] We ignore the screening of individuals born prior to 1960; for all others, we assume HCV screening occurs at age 50.

c. The optimal next action may be to terminate the screening program or to collect information about $\tilde{p}_t$, $\tilde{F}_t$, or both (depending on the scenario) in a specific future period.

Then, as an illustrative example, we demonstrate the application of our framework to HCV screening policy.

Analytical Results

When information is available in any period about $\tilde{\theta}_t$ only, it may be optimal to delay information collection about $\tilde{\theta}_t$ until a time when the information is more likely to change an immediate decision. To prove this, we consider the simplified case of only 2 possible realizations, $\tilde{\theta} \in \{\theta_{LOW}, \theta_{HIGH}\}$, neither of which would result in immediately switching from one intervention decision to the other. In this case, a policy maker would clearly prefer to delay spending on information collection until a time when the information is likely to become decision relevant. We present a formal proposition and proof in which we calculate the optimal time to collect information in online Appendix B.2. Following the same logic, the proof can be extended to consider more complicated $\tilde{\theta}$, including continuous distributions. Furthermore, this result holds even if the real (inflation-adjusted) cost of information is increasing, as long as the rate of information cost growth is less than the discount rate.

When information is available about $\tilde{p}_t$ and $\tilde{\theta}_t$ in any period, we use numerical methods to provide a solution. In online Appendix B.3, we show that the value function has the sufficient properties such that there is a unique solution.

Application to HCV Screening

The expected INMB HCV screening up to the 1965 birth cohort compared to a policy of no screening is $399 million. The optimal policy without information collection is to screen men for HCV until expected HCV prevalence equals $\frac{\gamma}{\mu_0(y_0)} = 0.4\%$ (see online Appendix B.1). Using equations (B.1) and (B.2), we can estimate that this occurs when the 1978 birth cohort turns 50 and that this policy increases the expected INMB by $167.3 million compared to screening until the 1965 cohort (Table 2).

If information can only be collected immediately, then the optimal strategy is to collect information about both HCV prevalence and the screen-detected fibrosis stage distribution at diagnosis ($(n_0, m_0) = (920, 200)$). This strategy increases the expected INMB by $397,000 compared to the policy of screening until the 1978 birth cohort (Table 2).

If we consider the illustrative situation of being able to collect information at any time but about only one parameter, it is optimal to delay information collection

in both cases. If information is only available about $\tilde{p}_t$, the optimal time to collect is in 16 years (when the 1976 birth cohort turns 50). If information is only available about $\tilde{F}_t$, the optimal time to collect that information is in 14 years (when the 1974 birth cohort turns 50). In each case, the information will guide the policy maker's next action.

Considering the more realistic scenario where information is available about $\tilde{p}_t$ in all periods and information is available about $\tilde{F}_t$ once, the optimal policy given current uncertainty about $\tilde{F}_t$ is composed of up to 5 regions (Figure 3A). When the policy maker is relatively confident that HCV prevalence is low, the optimal policy is to not screen. When the policy maker believes HCV prevalence is high, the optimal policy is to screen and not collect information. At intermediate prevalence, the optimal information collection strategy is primarily determined by the current belief about the mean of $\tilde{p}_t$. If the expected prevalence is very low but uncertainty about prevalence is preventing the policy maker from terminating the program, the optimal policy is to collect more information about prevalence. At slightly higher prevalence, the optimal policy is to collect information about both parameters or only about $\tilde{F}_t$.

This example also reveals 2 small pockets where the optimal policy is to screen and not collect information (Figure 3A). The first, close to the $\frac{\gamma}{\mu_\theta(y_0)}$ -threshold and with relatively low uncertainty about prevalence, exists because of the very high cost of fibrosis stage distribution information. States in this region are so close to the stopping region that information has little possible impact on the optimal screening policy. The second, beside the "screen and collect information about $\tilde{p}_t$" region at high uncertainty about prevalence, is increasing in width with increasing uncertainty about prevalence. This region exists because of the high cost of fibrosis stage distribution information relative to the contribution of $\tilde{F}_t$ to the overall decision uncertainty. As the contribution of $\tilde{p}_t$ to overall decision uncertainty grows, so does the preference for information about $\tilde{p}_t$ over $\tilde{F}_t$.

Given our current beliefs about $\tilde{p}_t$ and $\tilde{F}_t$, the optimal policy is to collect information about the screen-detected fibrosis stage distribution in 12 years. Compared to the optimal policy without information collection (screening until the 1978 birth cohort), this policy increased the expected INMB by \$2.12 million (Table 2). Once this information has been collected, the belief about $\tilde{F}_t$ can be updated and the optimal next action, the time to either stop the HCV screening program or to collect information about HCV prevalence, determined.

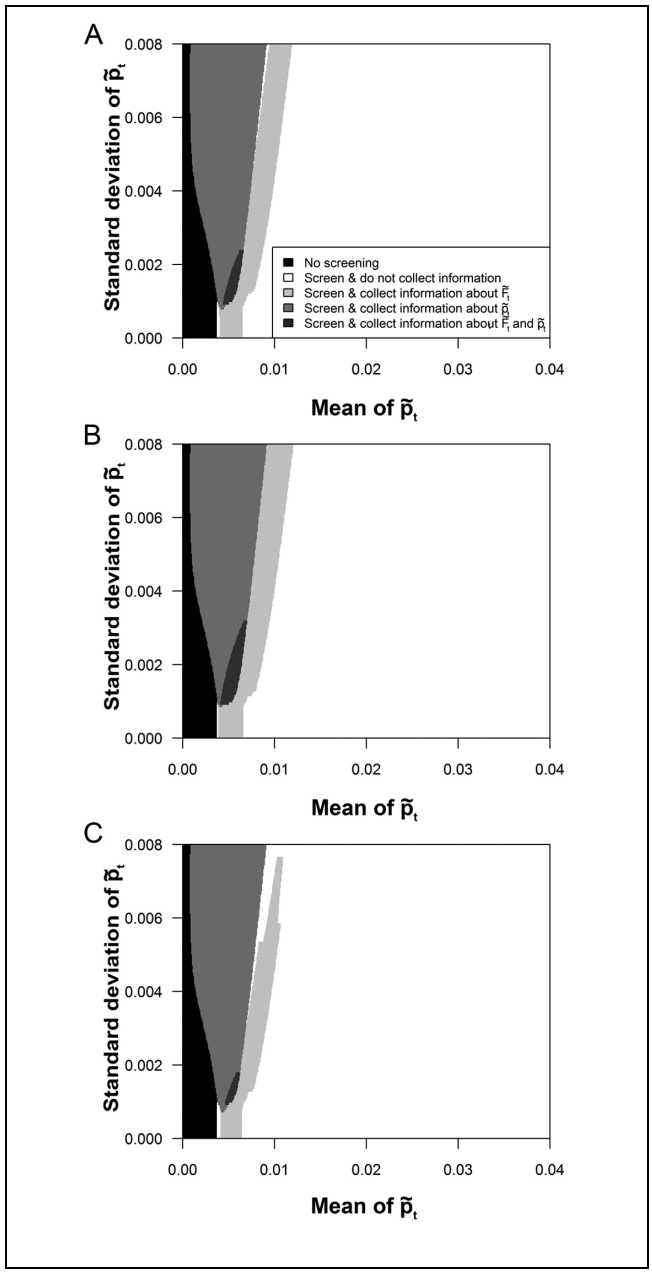

Figure 3 The optimal policy for any belief about hepatitis C virus (HCV) prevalence, $\tilde{p}_t$, given the current belief about the screen-detected fibrosis stage distribution, $\tilde{F}_t$. Decision alternatives include screening for HCV and not collecting any information, screening for HCV and performing a study to learn about HCV prevalence (sample size = 4500), screening for HCV and performing a study to learn about the screen-detected fibrosis stage distribution (sample size = 200), screening for HCV and performing both studies, or not screening for HCV. Panels vary in the cost of performing the fibrosis stage distribution study: (A) base case, \$800,000; (B) low cost, \$600,000; (C) high cost, \$1,000,000.

In our analysis, the per-participant costs are high ($2000 per person) because an ultrasound-guided liver biopsy is a relatively invasive and expensive procedure,[49] and the fixed costs of the study are high because it is difficult to identify newly diagnosed, screen-detected individuals willing to undergo liver biopsy.[48,53] However, some of these costs may be lower in a large centrally managed health system or if a noninvasive, but less accurate, liver fibrosis staging technology is used. When the cost of collecting information about the screen-detected fibrosis stage distribution is lower, the region in which it is optimal to collect information about fibrosis stage distribution is larger (Figure 3B). Conversely, this region becomes smaller as the costs of the study rise (Figure 3C).

Discussion

Evaluating a policy over its entire life cycle and considering the opportunity to collect information about both time-varying and time-invariant parameters using a stochastic dynamic programming approach yielded an interesting and policy-relevant insight: even when a parameter is not itself varying across intervention cohorts, it may be optimal to delay information collection about it. Because of the time value of money, it is potentially valuable to delay information collection until a time when the information gathered is more likely to result in a change in policy. Furthermore, we demonstrated how a stochastic programming approach can be used to simultaneously determine the optimal intervention and information collection policies.

We applied this framework to the timely problem of HCV screening. We focused on how the presence of a time-varying parameter, in this case HCV prevalence, $\tilde{p}_t$, affects the optimal information collection policy for an example time-invariant parameter (the liver fibrosis stage distribution at screen-detected diagnosis at age 50). Given our initial beliefs, if the opportunity to delay information collection is not considered, it appears optimal to collect information about both parameters now. However, if the opportunity to delay information collection is considered, then the optimal policy is to collect information only about the screen-detected fibrosis stage distribution in 12 years when the information is more likely to influence the HCV screening decision.

Our work extends the scope of applications of stochastic dynamic programs in the health care literature (reviewed in Schaefer et al.[54] and Alagoz et al.[55]), focusing on population policies rather than patient-level decisions and considering the timing of information collection as an integral part of optimal policy design. Previous applications of a Markov decision process framework to population policy, such as composition and timing decisions of flu shot design,[56,57] have assumed that information is collected in every period in which a final decision has not been made. However, information collection is costly, and so collecting information in every period may not be appropriate in other applications.

The focus of this article was to demonstrate that it may be optimal to delay information collection about time-invariant parameters and to present a framework for identifying the optimal time to collect information. We make several simplifying assumptions related to HCV, and we did not perform an exhaustive analysis of parameters involved in determining an optimal HCV screening policy. In addition, our work has several other limitations. First, for this example, we assumed that the screen-detected fibrosis stage distribution is constant over time. This may not be true because of changes in the age of infection and risk factors for infection.[58] Our objective was to demonstrate how the presence of a time-varying parameter, in this case HCV prevalence, influenced the optimal information collection policy for a time-invariant parameter, and the screen-detected fibrosis stage distribution is likely changing more slowly than HCV prevalence. Furthermore, for computational and illustrative reasons, we assumed information could only be collected once about the time-invariant parameter. We believe this simplification is acceptable, since performing such a study incurs significant fixed costs and, therefore, a single study is likely preferred to performing several studies.

Our framework assumes that information can be collected within a single decision period (in our application, a year). If the study requires more time, it may be possible to adjust the time of study initiation so that the information arrives at the period indicated with an appropriate adjustment to the study cost to account for the timing of study expenditures. However, studies with longer patient accrual horizons or longer observation horizons resulting in the risk of greater participant loss to follow-up may experience nonlinear costs, uncertainty in costs, and uncertainty in the time it takes for research to report. Hall et al.[59] present a framework for incorporating uncertainty in the time it takes for research to report into value-of-information analysis. For difficult-to-study parameters, an integration of their framework with our own may be necessary to identify the optimal time to initiate a study given the presence of time-varying parameters affecting the study population, uncertain

study costs over the study horizon, and an uncertain time to study completion. We also assume that a change in policy can be implemented immediately; implementation can be slow, and this may influence the optimal time to collect information.[60,61] In our example, we only consider one time-varying parameter and one time-invariant parameter. The greater the number of simultaneous and sequential information collection efforts, the more computationally intensive the problem is to solve; however, our framework can readily accommodate extension to additional parameters.

Finally, our simplified model structure assumes that to gain information about the next cohort, we can collect information only about the current cohort and rely on the correlation between cohorts (as described by the system dynamics) to provide information about the next cohort. In the case of HCV screening, however, it is possible to directly sample the next cohort (49-year-olds). We chose this model structure because of its generality, as the individuals who make up the "next cohort" are typically unknown (e.g., the next cohort of patients with a heart attack, of pregnant women, or of cancer patients).

We extend the literature on value-of-information assessment by developing a model to identify the optimal information collection policy for time-varying and time-invariant model parameters when at least one parameter is varying across intervention cohorts. For a typical health policy decision, many population-level parameters are changing over time (e.g., decreasing rates of smoking, increasing rates of obesity and type II diabetes). Such changes will affect the cost-effectiveness of medical interventions and health programs that are sensitive to these parameters. Furthermore, a technology itself may increase in effectiveness over time through explicit revision (technological improvement to a medical test or device) or through learning curve effects. Trends affecting the cost-effectiveness of a technology affect the value of information about the parameter that is changing and parameters that are not changing. Over long assessment horizons, ignoring the influence of time-varying parameters on the value of information may lead to over- or underestimates of the value of additional information and the optimal sample size. Furthermore, ignoring the opportunity to wait and collect information in the future, at a time when the information collected is more likely to support an immediate action, is a missed opportunity for increased efficiency. A dynamic programming framework, such as the one we present, increases the efficient use of scarce resources for health care and health research by more accurately assessing the marginal value of information and identifying when and how much information to collect.

Supplementary Material

Supplementary material for this article is available on the *Medical Decision Making* Web site at http://journals.sagepub.com/home/mdm.

ORCID iD

Lauren E. Cipriano https://orcid.org/0000-0001-5568-4516
Thomas A. Weber https://orcid.org/0000-0003-2857-4754

References

1. Wolk R, Shamsuzzaman AS, Somers VK. Obesity, sleep apnea, and hypertension. *Hypertension*. 2003;42(6):1067–74.
2. Lipscombe LL, Hux JE. Trends in diabetes prevalence, incidence, and mortality in Ontario, Canada 1995–2005: a population-based study. *Lancet*. 2007;369(9563):750–6.
3. Cowie CC, Rust KF, Ford ES, et al. Full accounting of diabetes and pre-diabetes in the US population in 1988–1994 and 2005–2006. *Diabetes Care*. 2009;32(2):287–94.
4. Cheng L, Eng C, Nieman LZ, Kapadia AS, Du XL. Trends in colorectal cancer incidence by anatomic site and disease stage in the United States from 1976 to 2005. *Am J Clin Oncol*. 2011;34(6):573–80.
5. Nelson DE, Giovino GA, Shopland DR, Mowery PD, Mills SL, Eriksen MP. Trends in cigarette smoking among US adolescents, 1974 through 1991. *Am J Public Health*. 1995;85(1):34–40.
6. Jia H, Lubetkin EI. Recent trends and geographic patterns of the burden of disease attributable to smoking. *Value Health*. 2010;13(8):958–64.
7. Cipriano LE, Romanus D, Earle CC, et al. Lung cancer treatment costs, including patient responsibility, by disease stage and treatment modality, 1992 to 2003. *Value Health*. 2011;14(1):41–52.
8. Hoyle M. Future drug prices and cost-effectiveness analyses. *Pharmacoeconomics*. 2008;26(7):589–602.
9. Chesson HW, Flagg EW, Koutsky L, et al. Modeling the impact of quadrivalent HPV vaccination on the incidence of Pap test abnormalities in the United States. *Vaccine*. 2013;31(29):3019–24.
10. Kulasingam SL, Myers ER. Potential health and economic impact of adding a human papillomavirus vaccine to screening programs. *JAMA*. 2003;290(6):781–9.
11. Thorn J, Coast J, Andronis L. Interpretation of the expected value of perfect information and research recommendations: a systematic review and empirical investigation. *Med Decis Making*. 2016;36(3):285–95.
12. Uthman RT, Sutton AJ, Jackson LJ, Uthman OA. Does directly administered antiretroviral therapy represent good value for money in sub-Saharan Africa? A cost-utility and

value of information analysis. *PLoS ONE*. 2018;13(1): e0191465.
13. Bennette CS, Veenstra DL, Basu A, Baker LH, Ramsey SD, Carlson JJ. Development and evaluation of an approach to using value of information analyses for real-time prioritization decisions within SWOG, a large cancer clinical trials cooperative group. *Med Decis Making*. 2016;36(5):641–51.
14. Hall PS, Smith A, Hulme C, et al. Value of information analysis of multiparameter tests for chemotherapy in early breast cancer: the OPTIMA Prelim Trial. *Value Health*. 2017;20(10):1311–8.
15. Gc VS, Suhrcke M, Hardeman W, Sutton S, Wilson EC. Cost-effectiveness and value of information analysis of brief interventions to promote physical activity in primary care. *Value Health*. 2018;21(1):18–26.
16. Eckermann S, Karnon J, Willan AR. The value of value of information. *Pharmacoeconomics*. 2010;28(9):699–709.
17. Carlson JJ, Thariani R, Roth J, et al. Value-of-information analysis within a stakeholder-driven research prioritization process in a US setting: an application in cancer genomics. *Med Decis Making*. 2013;33(4):463–71.
18. Siebert U, Rochau U, Claxton K. When is enough evidence enough? Using systematic decision analysis and value-of-information analysis to determine the need for further evidence. *Z Evid Fortbild Qual Gesundheitswes*. 2013;107(910): 575–84.
19. Tuffaha HW, Gordon LG, Scuffham PA. Value of information analysis in healthcare: a review of principles and applications. *J Med Econ*. 2014;17(6):377–83.
20. Wallner K, Shapiro AJ, Senior PA, McCabe C. Cost effectiveness and value of information analyses of islet cell transplantation in the management of unstable type 1 diabetes mellitus. *BMC Endocrine Disord*. 2016;16(1):17.
21. Ades A, Lu G, Claxton K. Expected value of sample information calculations in medical decision modeling. *Med Decis Making*. 2004;24(2):207–27.
22. Claxton K, Ginnelly L, Sculpher M, Philips Z, Palmer S. *A Pilot Study on the Use of Decision Theory and Value of Information Analysis as Part of the NHS Health Technology Assessment Programme*. York, UK: York Publishing; 2004.
23. Philips Z, Claxton K, Palmer S. The half-life of truth: what are appropriate time horizons for research decisions? *Med Decis Making*. 2008;28(3):287–99.
24. Hoyle M, Anderson R. Whose costs and benefits? Why economic evaluations should simulate both prevalent and all future incident patient cohorts. *Med Decis Making*. 2010;30(4):426–37.
25. Grimm SE, Dixon S, Stevens JW. When future change matters: modeling future price and diffusion in health technology assessments of medical devices. *Value Health*. 2016;19(6):720–6.
26. Cipriano LE, Weber TA. Population-level intervention and information collection in dynamic healthcare policy. *Health Care Manag Sci*. doi: 10.1007/s10729-017-9415-5.
27. Liu S, Cipriano LE, Holodniy M, Goldhaber-Fiebert JD. Cost-effectiveness analysis of risk-factor guided and birth-cohort screening for chronic hepatitis C infection in the United States. *PLoS ONE*. 2013;8(3):e58975.
28. Liu S, Barnett PG, Holodniy M, et al. Cost-effectiveness of treatments for genotype 1 hepatitis C virus infection in non-VA and VA populations. *MDM Policy Pract*. 2016;1(1):2381468316671946.
29. Chak E, Talal AH, Sherman KE, Schiff ER, Saab S. Hepatitis C virus infection in USA: an estimate of true prevalence. *Liver Int*. 2011;31(8):1090–101.
30. Armstrong GL, Wasley A, Simard EP, McQuillan GM, Kuhnert WL, Alter MJ. The prevalence of hepatitis C virus infection in the United States, 1999 through 2002. *Ann Intern Med*. 2006;144(10):705–714.
31. Denniston MM, Klevens RM, McQuillan GM, Jiles RB. 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*. 2012;55(6):1652–61.
32. Rein DB, Smith BD, Wittenborn JS, et al. The cost-effectiveness of birth-cohort screening for hepatitis C antibody in US primary care settings. *Ann Intern Med*. 2012;156(4):263–70.
33. Coffin PO, Scott JD, Golden MR, Sullivan SD. Cost-effectiveness and population outcomes of general population screening for hepatitis C. *Clin Infect Dis*. 2012;54(9): 1259–71.
34. McGarry LJ, Pawar VS, Panchmatia HR, et al. Economic model of a birth cohort screening program for hepatitis C virus. *Hepatology*. 2012;55(5):1344–55.
35. Eckman MH, Talal AH, Gordon SC, Schiff E, Sherman KE. Cost-effectiveness of screening for chronic hepatitis C infection in the United States. *Clin Infect Dis*. 2013;56(10): 1382–93.
36. Smith BD, Morgan RL, Beckett GA, Falck-Ytter Y, Holtzman D, Ward JW. Hepatitis C virus testing of persons born during 1945–1965: recommendations from the Centers for Disease Control and Prevention. *Ann Intern Med*. 2012;157(11):817–22.
37. Moyer VA. Screening for hepatitis C virus infection in adults: US Preventive Services Task Force recommendation statement. *Ann Intern Med*. 2013;159(5):349–357.
38. Raiffa H, Schlaifer R. *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University Press; 1961.
39. Bellman R. *Dynamic Programming*. Princeton, NJ: Princeton University Press; 1957.
40. US Census Bureau. QT–P1: Age Groups and Sex: 2010; 2010. Available from: http://www.census.gov/2010census/
41. Gretch DR. Diagnostic tests for hepatitis C. *Hepatology*. 1997;26(suppl 3):43S–7.
42. Hyland C, Kearns S, Young I, Battistutta D, Morgan C. Predictive markers for hepatitis C antibody ELISA specificity in Australian blood donors. *Transfusion Med*. 1992;2(3):207–13.

43. Gold MR, Siegel JE, Russell LB, Weinstein MC. *Cost-Effectiveness in Health and Medicine*. New York: Oxford University Press; 1996.
44. US Bureau of Labor Statistics. Consumer Price Index. Division of Consumer Prices and Price Indexes. Available from: http://www.bls.gov/cpi/
45. Weinstein MC, Skinner JA. Comparative effectiveness and health care spending: implications for reform. *N Engl J Med*. 2010;362(5):460–5.
46. Innes H, Goldberg D, Dusheiko G, et al. Patient-important benefits of clearing the hepatitis C virus through treatment: a simulation model. *J Hepatol*. 2014;60(6):1118–26.
47. Liu S, Cipriano LE, Holodniy M, Owens DK, Goldhaber-Fiebert JD. New protease inhibitors for the treatment of chronic hepatitis C: a cost-effectiveness analysis. *Ann Intern Med*. 2012;156(4):279–90.
48. Mallette C, Flynn MA, Promrat K. Outcome of screening for hepatitis C virus infection based on risk factors. *Am J Gastroenterol*. 2008;103(1):131–7.
49. Centers for Medicare & Medicaid Services. Medicare fee schedule. US Department of Health and Human Services; 2010. Available from: http://www.cms.gov/home/medicare.asp
50. Therneau T, Atkinson B, Ripley B. *rpart: Recursive partitioning*. R package version 4.1-1. Available from: http://CRAN.R-project.org/package=rpart
51. Claxton KP, Sculpher MJ. Using value of information analysis to prioritise health research: some lessons from recent UK experience. *Pharmacoeconomics*. 2006;24(11):1055–68.
52. R Development Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria; 2012. Available from: http://www.R-project.org/
53. Allison RD, Conry-Cantilena C, Koziol D, et al. A 25-year study of the clinical and histologic outcomes of hepatitis C virus infection and its modes of transmission in a cohort of initially asymptomatic blood donors. *J Infect Dis*. 2012;206(5):654–61.
54. Schaefer AJ, Bailey MD, Shechter SM, Roberts MS. Modeling medical treatment using Markov decision processes. In: Brandeau ML, Sainfort F, Pierskalla WP, eds. *Operations Research and Health Care: A Handbook of Methods and Applications*. Springer US; 2005. p 593–612.
55. Alagoz O, Hsu H, Schaefer AJ, Roberts MS. Markov decision processes: a tool for sequential decision making under uncertainty. *Med Decis Making*. 2010;30(4):474–83.
56. Kornish LJ, Keeney RL. Repeated commit-or-defer decisions with a deadline: the influenza vaccine composition. *Oper Res*. 2008;56(3):527–41.
57. Özaltın OY, Prokopyev OA, Schaefer AJ, Roberts MS. Optimizing the societal benefits of the annual influenza vaccine: a stochastic programming approach. *Oper Res*. 2011;59(5):1131–43.
58. Alter MJ. Prevention of spread of hepatitis C. *Hepatology*. 2002;36(5B):s93–8.
59. Hall PS, Edlin R, Kharroubi S, Gregory W, McCabe C. Expected net present value of sample information from burden to investment. *Med Decis Making*. 2012;32(3):E11–E21.
60. Andronis L, Barton PM. Adjusting estimates of the expected value of information for implementation: theoretical framework and practical application. *Med Decis Making*. 2016;36(3):296–307.
61. Grimm SE, Dixon S, Stevens JW. Assessing the expected value of research studies in reducing uncertainty and improving implementation dynamics. *Med Decis Making*. 2017;37(5):523–33.
62. Mehrotra A, Zaslavsky AM, Ayanian JZ. Preventive health examinations and preventive gynecological examinations in the United States. *Arch Intern Med*. 2007;167(17):1876.