





Dynamic Learning of Patient Response Types: An Application to Treating Chronic Diseases

Diana M. Negoescu,^a Kostas Bimpikis,^b Margaret L. Brandeau,^c Dan A. Iancu^b

^a Industrial and Systems Engineering Department, University of Minnesota, Minneapolis, Minnesota 55455; ^b Stanford Graduate School of Business, Stanford, California 94305; ^c Department of Management Science and Engineering, Stanford University, Stanford, California 94305

Contact: negoescu@umn.edu,  <http://orcid.org/0000-0002-8294-275X> (DMN); kostasb@stanford.edu,  <http://orcid.org/0000-0002-5289-0397> (KB); brandeau@stanford.edu,  <http://orcid.org/0000-0001-9331-8920> (MLB); daniancu@stanford.edu,  <http://orcid.org/0000-0001-9890-0928> (DAI)

Received: September 15, 2014

Revised: December 8, 2015; August 30, 2016; November 11, 2016

Accepted: November 19, 2016

Published Online in Articles in Advance: August 21, 2017

<https://doi.org/10.1287/mnsc.2017.2793>

Copyright: © 2017 INFORMS

Abstract. Currently available medication for treating many chronic diseases is often effective only for a subgroup of patients, and biomarkers accurately assessing whether an individual belongs to this subgroup typically do not exist. In such settings, physicians learn about the effectiveness of a drug primarily through experimentation—i.e., by initiating treatment and monitoring the patient’s response. Precise guidelines for discontinuing treatment are often lacking or left entirely to the physician’s discretion. We introduce a framework for developing adaptive, personalized treatments for such chronic diseases. Our model is based on a continuous-time, multi-armed bandit setting where drug effectiveness is assessed by aggregating information from several channels: by continuously monitoring the state of the patient, but also by (not) observing the occurrence of particular infrequent health events, such as relapses or disease flare-ups. Recognizing that the timing and severity of such events provide critical information for treatment decisions is a key point of departure in our framework compared with typical (bandit) models used in healthcare. We show that the model can be analyzed in closed form for several settings of interest, resulting in optimal policies that are intuitive and may have practical appeal. We illustrate the effectiveness of the methodology by developing a set of efficient treatment policies for multiple sclerosis, which we then use to benchmark several existing treatment guidelines.

History: Accepted by Noah Gans, stochastic models and simulation.

Funding: M. L. Brandeau was supported by the National Institute on Drug Abuse [Grant R01-DA15612].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2017.2793>.

Keywords: adaptive treatment • dynamic programming • optimal control • stochastic model applications • multiarmed bandits • continuous time

1. Introduction

Costs associated with the delivery of healthcare in the United States have risen sharply in recent years, both in terms of total expenditure (e.g., as a percentage of gross domestic product) and in spending recognized as wasteful, redundant, or inefficient (Young and Olsen 2010). In conjunction with advances in the field of medicine and the use of information technology, this has created increasing pressure for healthcare solutions that deliver better outcomes in a cost-effective manner.

Despite this impetus, however, the design of adaptive treatment policies for chronic conditions (e.g., multiple sclerosis, Crohn’s disease, and depression) has often been perceived as slow,¹ with some of the complicating factors intrinsically related to the specifics of disease progression and available medication.

Currently available disease-modifying therapies (DMTs) for several chronic diseases are only effective in a subset of the population (“responders”), and biomarkers that accurately assess a priori whether a

given patient belongs to this subgroup are not available.² In such cases, the main way to evaluate DMT efficacy is by initiating treatment and then continuously monitoring the patient through self-reported surveys, periodic check-ups, or more in-depth scans and evaluations.

If the role of treatment were the reversal of an obvious short-term abnormality, such monitoring would provide sufficient evidence for how well the patient is responding. However, the primary goal of DMTs for *chronic* diseases is to prevent disease progression in the long run, which often translates to limiting the occurrence of infrequent negative health events (e.g., disease flare-ups) that can severely diminish a patient’s quality of life. As such, the (non)occurrence or the exact *timing* and severity of such episodes often convey critical information concerning a DMT’s effectiveness for the patient. Quantifying the impact of such information and translating it into actionable guidelines for medical decision making is often not straightforward.

A primary example of a chronic disease with these features is multiple sclerosis (MS), an autoimmune inflammatory disease of the central nervous system that is a leading cause of disability in young adults. MS is incurable; DMTs attempt to slow its progression by decreasing the frequency and severity of clinical attacks, known as “relapses” (see, e.g., Cohen et al. 2004, National Multiple Sclerosis Society 2017). While newly available drugs represent advances for MS management, none is fully effective (Rovaris et al. 2001), and the question of identifying patients who are not responsive to treatment is centrally important. In the words of the National Clinical Advisory Board of the National Multiple Sclerosis Society (2004, p. 1),

Whatever the relative merits of these drugs, all can only be considered partially effective agents. This reality raises the difficult problem of the identification of a suboptimal response or treatment failure in an individual case and, once identified, leads to consideration of the appropriate avenues for alternative treatments.

As the quote highlights, the problem of identifying patients who do not respond to DMTs is quite challenging. For a newly diagnosed patient, current guidelines recommend immediately starting treatment, and assessing effectiveness by continually monitoring the disease progression, through MRI scans and self-reported assessments of disability, such as the Expanded Disability Status Scale (EDSS) (National Multiple Sclerosis Society 2008). The guidelines emphasize the critical role of learning and explicitly recognize that the timing and frequency of relapses, as well as more continuous measurements such as EDSS and/or MRI, can all be informative.³ However, they stop short of providing a systematic way to use this information and suggest only simple rules for discontinuing treatment. To the best of our knowledge, these rules are not the outcome of a quantitative framework and have not been tested for efficiency (see, Cohen et al. 2004). Furthermore, while several studies have attempted to identify early predictors of nonresponse (Horakova et al. 2012, Romeo et al. 2013), the results have not been used to inform the design of optimal treatment plans in a quantitative fashion.

Further underscoring the need for fast and accurate identification of nonresponders is the fact that DMTs can cause significant side effects, such as persistent flu-like symptoms, injection-site necrosis, and liver damage, which result in poor compliance and large dropout rates (Prosser et al. 2004). Additionally and quite importantly, treatment is expensive, with mean annual costs of \$60,000 per diagnosed case in the United States (Hartung et al. 2015). This has resulted in a significant amount of debate around policies for MS treatment, in the United States and elsewhere.⁴

This example and the preceding discussion give rise to several natural research questions. Given the

available medications, what is the optimal treatment plan for chronic diseases such as MS? Does an optimal plan involve discontinuation rules—i.e., is it optimal to start a patient on treatment and then stop at a particular point in time? How can a medical decision maker optimally aggregate all of the information acquired during treatment to design optimal treatment plans? Would such optimized plans outperform current existing medical guidelines?

This paper can be viewed as one step toward answering such questions. We propose a framework that can be used to inform treatment decisions for chronic diseases that have the features described above: treatment is effective only for a subset of patients that is a priori unidentifiable; the frequency and/or severity of side effects and major health events depends on a patient’s response type; and information regarding the effectiveness of treatment is obtained gradually over time. Our main contributions can be summarized as follows:

- We formulate the problem of determining an optimal adaptive treatment policy as a continuous-time stochastic control problem. A key point of departure from other work in medical decision making is that we incorporate information from both the *day-to-day monitoring* of disease progression, as well as the *timing* of major health events. Our framework thus implicitly trades off the immediate and the long-term impact of treatment, providing a systematic way to incorporate new information in the design of optimal treatment plans.

- When choosing between two treatments with linear dose–response, our model can be analyzed in closed form, resulting in intuitive optimal policies that take the form of discontinuation rules. In an extension discussed in Online Appendix A, we show how these analytical results can be used to derive optimal policies for choosing among several treatments by solving very simple (one-dimensional, convex) optimization problems. We also discuss conditions under which the optimal policy is no longer a simple discontinuation rule, such as when dose–response curves are nonlinear or when the *severity* of major health events is indicative of treatment effectiveness.

- We apply our results to MS, for which we develop and test adaptive treatment policies for administering interferon- β . Our framework allows an explicit trade-off between the benefits of treatment and its associated costs, and can be used to generate an entire frontier of cost-effective treatment policies, depending on the amount a decision maker is willing to pay for one additional quality-adjusted life-year (QALY), expressed as a willingness-to-pay (WTP) value. We use these policies to benchmark and test the performance of three treatment guidelines: a “no-treatment” policy, which does not prescribe any interferon; a “standard” policy,

which administers interferon to all patients and discontinues treatment on progression to an EDSS score of 6–7.5 (Río et al. 2011); and a “consensus” policy, which discontinues treatment when patients experience two or more relapses in a year or progress to an EDSS score of 6–7.5 (Cohen et al. 2004).

Our first finding suggests that a no-treatment policy is optimal if the WTP does not exceed \$150,000/QALY. Furthermore, the gains from interferon treatment are generally not large in absolute terms and come at steep costs: even the best adaptive policy, requiring a WTP exceeding \$800,000/QALY, can only increase QALYs by 3.45% relative to a no-treatment alternative, while increasing costs by 16.2%. Since interventions are generally considered cost-effective when the costs per QALY gained do not exceed three times the country’s per-capita GDP (Drummond 2005, Hunink et al. 2014), this suggests that interferon treatment is not necessarily cost-effective, and that no-treatment may be optimal under lower WTP. However, this finding should be interpreted with caution: even though the increases in QALYs may not be large in absolute terms, they may nonetheless be significant, particularly for a chronic disease as debilitating as MS.

Our results provide validation of the consensus criteria proposed by Cohen et al. (2004): the resulting policy is close to being efficient at intermediate values of WTP, and achieves net monetary benefits close to a fully adaptive policy. Thus, these simple discontinuation rules may represent a viable alternative to implementing a complex optimal adaptive policy, particularly at intermediate WTP values.

Finally, we find that none of the treatment guidelines are satisfactory at very large WTP: the consensus and no-treatment policies generate low QALYs, and the standard policy is inefficient. An adaptive policy derived from our framework under a WTP of \$800,000/QALY generates the most QALYs without incurring significant costs, and attains a good balance between administering sufficient treatment to responders and identifying nonresponders early.

While we apply our model to MS and interferon- β primarily because of data availability, we note that the treatment of many other chronic diseases could benefit from our framework. Examples include rheumatoid arthritis, where increased disability is associated with higher mortality (Pincus et al. 1984); Crohn’s disease, where treatment often involves the same classes of medications as MS; and depression and other mental illnesses, where psychiatrists must choose between various treatments without knowing a priori which one might be effective.

1.1. Relevant Literature

Our model builds on the theory of continuous-time multi-armed bandits (Berry and Fristedt 1985, Mandelbaum 1987, Bank and Küchler 2007, Cohen and Solan

2013, Harrison and Sunar 2015). Closest to our work are papers on strategic experimentation (Bolton and Harris 1999, Keller et al. 2005, Keller and Rady 2010), which study free riding among a team of agents in an experimentation context. We adapt their framework in a medical decision-making setting, and extend their model and analysis by allowing the decision maker to learn from observing the rewards generated by two stochastic processes whose parameters depend on the choice of treatment: a Wiener process (Brownian motion) that models the day-to-day side effects experienced by the patient, and a Poisson process that captures the arrival of major health events—i.e., disease flare-ups and progression.

Our paper is related to the clinical trials literature and in particular to the growing number of studies that consider adaptive rules for assigning patients to treatments (e.g., Berry 1978, Berry and Pearson 1985, Ahuja and Birge 2016, Bertsimas et al. 2016). These approaches typically assume that the outcome of a clinical trial is binary (success/failure), and that the decision maker can learn from multiple patients since outcomes are positively correlated. It is difficult to implement such an approach in the context of a chronic disease, however, as one patient’s response to treatment is independent from another’s, and information about the quality of treatment is obtained gradually over time, with no single event providing sufficient evidence for or against a given treatment plan.

Also closely related is a growing literature (e.g., Denton et al. 2009, Zhang et al. 2012, Mason et al. 2014, Helm et al. 2015) that uses Markov decision processes with fully or partially observed states and dynamic linear Gaussian systems to derive adaptive treatment policies. We formulate the problem using a continuous-time bandit model and derive closed-form expressions for optimal treatment decisions that, as we discuss in the paper, may have several advantages.

In the medical literature, Murphy (2003, 2005), Murphy and Collins (2007), Pineau et al. (2007), and Almirall et al. (2012), among others, propose adaptive treatment schemes in the context of psychiatric conditions such as depression, anxiety disorders, and drug/alcohol abuse. In particular, Murphy (2003) motivates and lays the foundation for developing *dynamic treatment regimes*—i.e., a set of rules for choosing effective treatments that are tailored to the individual characteristics of patients. The methodology in these studies varies from nonquantitative approaches (e.g., predefining a protocol to switch therapies after a certain time if a criterion is not met (Almirall et al. 2012)), to reinforcement learning (Pineau et al. 2007) and statistical frameworks for optimizing a general outcome while achieving a desired level of power or bias (Murphy 2005). Although we share the same motivation and use

similar methodological tools as some of these studies, i.e., dynamic programming (albeit, in continuous time), the emphasis in these papers is not on deriving explicit optimal treatment policies nor on preserving the computational tractability of the resulting framework. In contrast, our explicit characterization of the optimal adaptive policies allows us to compute useful comparative statics with respect to features of the underlying environment and to derive optimal policies for the case of multiple treatments by solving very simple (one-dimensional) convex optimization problems in an offline fashion (see Online Appendix A).

Finally, our work is also related to empirical cost-effectiveness studies for MS, which have found DMTs to be very expensive for the benefits they provide, with costs of up to \$1.6 million per additional QALY gained (Phillips 2004, Tappenden et al. 2009, Noyes et al. 2011). Although we use a similar disease evolution model, the key point of departure is that we assess cost-effectiveness based on optimal adaptive treatments that utilize all available information, instead of heuristic treatment guidelines. We find that the optimal interferon treatment is not necessarily cost-effective, and we quantify the WTP values under which no-treatment is optimal.

2. Model Formulation

We first introduce our model in an abstract setting, and then discuss the connection and relevance to the medical applications motivating our work. In an effort to make the paper accessible to a broad audience, we deliberately keep the exposition style less formal, placing more emphasis on the intuition and connection with the applications. Readers interested in the mathematical details can refer to El Karoui and Karatzas (1994), Bolton and Harris (1999), Keller and Rady (2010), and references therein, which form the basis of our model.

2.1. Model Framework

We consider a continuous time frame, indexed by $t \in [0, \infty)$. A single decision maker (DM) is faced with the problem of choosing how to allocate the current period $[t, t + dt)$ between two possible alternatives (“arms”): a “safe” alternative, with known characteristics, and a “risky” alternative, which can be of either good (G) or bad (B) type, unbeknownst to the DM.

Each arm brings the DM immediate rewards that accrue continuously over time. More precisely, the “safe” arm generates instantaneous rewards governed by a Brownian motion with drift rate μ_0 and volatility σ , and a risky arm of type $\theta \in \{G, B\}$ generates instantaneous Brownian rewards with drift rate μ_θ and volatility σ . When the DM allocates a fraction $\alpha_t \in [0, 1]$ of the time interval $[t, t + dt)$ to the risky arm and the remaining fraction $1 - \alpha_t$ to the safe arm, the total

instantaneous rewards received are $d\pi^1(t) + d\pi^0(t)$, where

$$d\pi^1(t) \stackrel{\text{def}}{=} \alpha_t \mu_\theta dt + \sqrt{\alpha_t} \sigma dZ^1(t), \quad (1a)$$

$$d\pi^0(t) \stackrel{\text{def}}{=} (1 - \alpha_t) \mu_0 dt + \sqrt{1 - \alpha_t} \sigma dZ^0(t). \quad (1b)$$

Here, $dZ^0(t)$ and $dZ^1(t)$ are independent, normally distributed random variables, with mean zero and variance dt . To understand the scaling used, note that the DM’s instantaneous rewards from the risky and safe arm are normally distributed, with mean $\alpha_t \mu_\theta dt$ and variance $\alpha_t \sigma^2 dt$, and mean $(1 - \alpha_t) \mu_0 dt$ and variance $(1 - \alpha_t) \sigma^2 dt$, respectively. As such, the total instantaneous reward exactly equals a fraction α_t of the risky reward and $1 - \alpha_t$ of the safe reward.

In addition to the instantaneous rewards, each arm also induces relatively rare “life events,” as well as a special “stopping event” that terminates the decision process. The occurrence of any life event generates a deterministic “reward” of $-D$. The frequency of such events depends on the allocation used by the DM. More precisely, when $\alpha_t \in [0, 1]$ of the period is allocated to the risky arm, life events occur according to a Poisson process with rate $(1 - \alpha_t) \lambda_0 + \alpha_t \lambda_\theta$, where λ_0 (λ_θ) denotes the rate of life events under a safe arm (a risky arm of type θ , respectively). Similarly, the stopping event occurs at a time T that is exponentially distributed with rate $(1 - \alpha_t) \eta_0 + \alpha_t \eta_\theta$, and generates a reward of magnitude V . We assume that, conditional on the risky arm’s type θ and on the allocation α_t , the stopping event is independent from the Poisson process for life events.⁵

The DM knows all of the underlying parameters governing the arms and the reward structure (i.e., μ_0 , λ_0 , σ , D , V , μ_θ , and λ_θ , for $\theta \in \{G, B\}$), but does *not* know the type θ of the risky arm. At time $t = 0$, he starts with some initial belief p_0 that the risky arm is good, which he then updates during the rest of the planning horizon, depending on the observed instantaneous and lump-sum rewards. This generates an updated belief p_t at time t .

The DM’s goal is to find a nonanticipative allocation policy $\{\alpha_t\}_{t \geq 0}$ that maximizes the total expected discounted rewards Π up to the stopping event—i.e.,

$$\Pi \stackrel{\text{def}}{=} \mathbb{E} \left[\int_0^T e^{-rt} [d\pi^1(t) + d\pi^0(t) - (N_{t+dt} - N_t) D] + e^{-rT} V \right], \quad (2)$$

where N_t denotes the total number of life events occurring in $[0, t]$.

Some observations regarding the problem formulation are in order. First, note that the integrand in the expression for Π contains three terms. The first two, $d\pi^1(t)$ and $d\pi^0(t)$, correspond to the instantaneous

rewards received from the risky and safe arms, given in (1a) and (1b), respectively. The third term corresponds to the expected lump-sum reward received on the occurrence of a life event during period $[t, t + dt]$. The integral is taken over the total (instantaneous and lump-sum) rewards discounted at a fixed rate $r > 0$, and the term outside of the integral corresponds to the reward received on the occurrence of the stopping event at time T . The expectation in (2) is with respect to the stochastic processes $dZ^0(t)$, $dZ^1(t)$, α_t , and also p_t . The latter reflects the DM's use of the belief p_t at time t with regard to the type θ of the risky arm.⁶

Note that in choosing a policy α_t to maximize the expected rewards, the DM is faced with the classical trade-off between “exploration” and “exploitation” (Powell and Ryzhov 2012)—i.e., between acquiring information about an unknown alternative, which may entail higher rewards, versus using a safe option. In this sense, α_t trades off the rate at which new information is gained with the risks entailed by the experimentation. With a choice $\alpha_t = 0$, the DM would only gain instantaneous and lump-sum rewards from the safe arm, hence completely eliminating the exposure to the risky arm as well as the ability to update the belief p_t . It is important to emphasize that new information in our model is acquired through two channels: (1) by observing the instantaneous rewards $d\pi^1(t)$ from the risky arm, and (2) by (not) observing life events and the stopping event. Whenever $\alpha_t > 0$, these channels all convey meaningful information to the DM, potentially tilting his belief p_t toward (or away from) deeming the risky arm as good.

2.2. Application in the Context of Chronic Diseases

We now discuss how our mathematical framework can be applied to the design of an adaptive treatment policy for chronic diseases such as MS.

The arms. In a medical context, the arms of our model correspond to available treatments, and the DM is a physician choosing the optimal treatment policy for the patient. Depending on the focus, the “rewards” could correspond to either a patient's health utility or a cost-adjusted health utility that also accounts for the cost of treatment (see our more detailed discussion in Section 4). As such, an arm's instantaneous reward denotes the impact of treatment on the patient's immediate (cost-adjusted) quality of life. “Life events” correspond to sudden health episodes associated with normal disease progression, which bring about immediate disutility (and costs) to the patient, without altering the fundamental underlying disease evolution or the efficacy of treatments. Examples of life events include *relapses* in MS or panic attacks in anxiety disorders. Depending on the circumstance and the exact disease modeled, the “stopping event” could be a special instance of a life event or an entirely separate

event, which changes the disease evolution or the treatment options (e.g., a heart attack, kidney failure, malignancy, or death). We elaborate on these distinctions further when discussing the *objective*. We note that our base-case model only allows choosing between two arms/treatments. We discuss the important extension to multiple arms in Online Appendix A.

Safe arm. A “safe” arm represents a treatment with homogenous response in the population. In MS, this typically consists of medication aimed at reducing or controlling MS-specific symptoms such as bowel and bladder function, spasticity, and pain, without modifying disease progression. In our model, such a treatment may still yield stochastic outcomes in terms of both instantaneous (cost-adjusted) health utility and life/stopping events, as one would expect in practice. The critical assumption is that the parameters governing these outcomes ($\mu_0, \sigma, \lambda_0, \eta_0$) are known to the physician. This is reasonable, since physicians often have more information about the natural disease progression when patients are not subjected to treatment (e.g., from studies of large historical cohorts of patients (Scalfari et al. 2010)).

Risky arm. The “risky” arm is only effective for a subset of the population (i.e., when the type is good ($\theta = G$)). We assume that the physician is unable to determine a priori whether a new patient belongs to this subset. This is in keeping with the fact that precise biomarkers do not exist for many chronic diseases. For instance, treatments for MS such as interferon- β are effective only in a subgroup of patients (Cohen et al. 2004, Prosser et al. 2003, Horakova et al. 2012). In such cases, the only way to assess the impact of a drug or therapy is by subjecting the patient to treatment and relying on periodic examinations or self-reported assessments, such as the EDSS for MS. When patients respond to treatment, their condition may improve (i.e., $\mu_G > \mu_0$), the frequency of life events may be diminished (i.e., $\lambda_G < \lambda_0$), and the likelihood of a major health event or disease progression may also decrease (i.e., $\eta_G < \eta_0$). When patients do *not* respond, their condition may remain the same or even deteriorate slightly (e.g., due to side effects from treatment). A central assumption underlying our model is that physicians are able to separately assess the parameters governing how responders and nonresponders are impacted by treatment (i.e., $\mu_\theta, \lambda_\theta$, and η_θ). This is reasonable since medical studies often track patients for a relatively long period of time and retrospectively assign them to responder and nonresponder groups (e.g., Horakova et al. 2012).

Lump-sum rewards. Our assumption that the lump-sum “reward” $-D$ received on life events is independent of the type θ is particularly pertinent for diseases such as MS and anxiety disorder. For instance,

relapses in MS correspond to periods of acute disease activity when patients experience neurological symptoms such as sudden paralysis or loss of vision. Such episodes generate immediate disutility and have similar severity/consequences in all patients, but occur less often among patients responding to treatment (Kremenutzky et al. 2006, Horakova et al. 2012). Our framework can be extended to *stochastic* rewards that are independent of θ , as only their expected value would matter. We discuss the extension to rewards dependent on θ in Online Appendix B.

Fractional allocations. Our model assumes that fractional allocations of treatment are possible, and that the response (i.e., the reward) is directly proportional to the allocation. Fractional allocations allow modeling cocktails of drugs (Rudick et al. 2006) or administering a lower dosage of a drug (e.g., by adjusting the frequency and/or the magnitude of doses). The assumption that the response is linear renders our model analytically tractable and is a reasonable first-order approximation, as dose–response functions are often S-shaped and thus linear in a central band of values (see, e.g., the MS study of OWIMS 1999). We discuss this limitation further in Section 5, and we examine its impact numerically in Online Appendix D.

Objective. Considering a planning horizon T that corresponds to an exponentially distributed “stopping” event allows for modeling flexibility without sacrificing analytical tractability. The horizon T could capture the first occurrence of a life event, which is appropriate when the risky treatment improves the immediate quality of life of a patient but incurs a higher risk of severe side effects. For instance, studies have shown that certain rheumatoid arthritis treatments improve pain and disability, but may cause malignancies or severe infections (Mariette et al. 2011). More broadly, T could correspond to any major event that permanently alters the state of the patient, the disease evolution, or the response/rewards from treatment. Examples could include progression to severe disease (e.g., in MS, transitioning from the relapsing–remitting phase to the secondary–progressive phase (Lee et al. 2012)) or the release of a new drug that alters the set of feasible treatment options or drastically reduces the cost of treatment, impacting the reward rates in a cost-adjusted objective. The rewards V received on the stopping event can be interpreted as continuation values, which allows using our model as a building block for studying diseases with more complex dynamics, involving potentially nonstationary reward rates or phase transitions. For more details, we refer the reader to our case study in Section 4, which implements this idea.

Our model includes a fixed discount rate $r > 0$, in keeping with the recommendations of the U.S. Panel on Cost-Effectiveness in Health and Medicine that costs

and quality-adjusted life years should be discounted when estimating the cost-effectiveness of healthcare interventions (Gold 1996).

Simplifying assumptions. To preserve analytical tractability, our model makes a number of simplifying assumptions: arms/treatments are characterized by Brownian rewards and Poisson arrivals with known and stationary parameters; dose–response curves are linear; information collection and treatment updating can be conducted very frequently; and patients fully adhere to treatment recommendations. In Section 5, we discuss these limitations more extensively, providing several extensions and robustness checks, and outlining interesting directions for future research.

Although our model simplifies the reality of chronic diseases, it has the advantage of allowing exact analytical results, with simple and intuitive interpretations, as we discuss next.

3. Analysis

Letting \mathcal{F}_t denote the information set available to the DM at time t ,⁷ it can be seen that the belief that the risky arm is good, $p_t \stackrel{\text{def}}{=} \mathbb{P}\{\theta = G \mid \mathcal{F}_t\}$, is a sufficient statistic of the history up to time t . Thus, we take p_t as the state of the system, and we take the fraction of treatment allocated to the risky arm α_t as the DM’s action (control). Finally, with $\mathcal{A} \stackrel{\text{def}}{=} \{(\alpha_t)_{t \geq 0} \mid \alpha_t: \mathcal{F}_t \rightarrow [0, 1]\}$ denoting the set of all sequential, nonanticipative policies that are adapted to the available information, the DM’s problem can be compactly formulated as

$$\max_{\alpha \in \mathcal{A}} \mathbb{E}^\alpha \left[\int_0^T e^{-rt} [d\pi^1(t) + d\pi^0(t) - (N_{t+dt} - N_t)D] + e^{-rT}V \right],$$

where the expectation is with respect to the stochastic processes $dZ^0(t)$, $dZ^1(t)$, N_t , α_t and p_t . As a first step in our analysis, we characterize the evolution of the DM’s belief during time interval $[t, t + dt]$, as a function of the current belief p_t and action α_t . We start with the case when no life event or stopping event occurs during the interval.

Lemma 1. *When no life event or stopping event occurs during time interval $[t, t + dt]$:*

(i) *the posterior belief p_{t+dt} conditional on an observed instantaneous reward from the risky arm $d\pi^1(t) = y$ is given by Bayes’ rule and takes a value of*

$$p_{t+dt} = (p_t F(\mu_G/\sigma) e^{-(\bar{\lambda}_G + \bar{\eta}_G)dt}) \cdot [p_t F(\mu_G/\sigma) e^{-(\bar{\lambda}_G + \bar{\eta}_G)dt} + (1 - p_t) F(\mu_B/\sigma) e^{-(\bar{\lambda}_B + \bar{\eta}_B)dt}]^{-1}, \quad (3)$$

where

$$\bar{\xi}_\theta \stackrel{\text{def}}{=} (1 - \alpha_t)\xi_0 + \alpha_t\xi_\theta, \quad \forall \xi \in \{\lambda, \eta\}, \forall \theta \in \{G, B\}, \quad (4a)$$

$$F(\mu) \stackrel{\text{def}}{=} \frac{1}{\sqrt{2\pi}dt} \exp\left(-\frac{1}{2dt}(y/\sigma - \sqrt{\alpha}\mu dt)^2\right); \quad (4b)$$

(ii) the change in the DM's belief $p_{t+dt} - p_t$ is normally distributed, with mean $\alpha_t p_t (1 - p_t) ((\lambda_B + \eta_B) - (\lambda_G + \eta_G)) dt$ and variance $\alpha_t ((p_t (1 - p_t) (\mu_G - \mu_B)) / \sigma)^2 dt$.

We make several observations about this result. First, note that when no life event or stopping event occurs during $[t, t + dt)$, the belief evolution only depends on the characteristics of the merged process of life and stopping events, which is Poisson under our assumptions. Thus, all the results depend on the sum of the rates of life events and stopping events (i.e., $\lambda_i + \eta_i$, $\forall i \in \{0, G, B\}$).

Part (i) of the result provides the update rule for the DM's belief. Note that changes only occur when the risky arm is used (i.e., when $\alpha_t > 0$), and the posterior only depends on the observed instantaneous reward from the risky arm ($d\pi^1$), but not from the safe arm ($d\pi^0$). This is intuitive, since the safe arm conveys no information about the risky arm's type. Note that the result implicitly requires the ability to separately observe the risky rewards, which may be problematic when the DM only observes the total instantaneous rewards $d\pi^0 + d\pi^1$ and $\alpha_t \in (0, 1)$. As our later results will show, this issue becomes moot in our setting, since the optimal policy will always entail $\alpha_t \in \{0, 1\}$, so the DM will never observe a mix of safe and risky instantaneous rewards.

Part (ii) establishes that the belief change is normally distributed, with parameters that depend on the arms' characteristics. The belief drifts upward—i.e., the risky arm is deemed more likely to be good—if and only if events under a good arm are less likely than under a bad arm (i.e., $\lambda_G + \eta_G < \lambda_B + \eta_B$). This is intuitive, since the absence of an event under such conditions can be viewed as “good news” for the DM. Consistent with this observation, note that “more learning” occurs (i.e., the mean belief update grows) as the difference between the rates of events under a good and a bad arm increases. More learning also occurs as $p_t(1 - p_t)$ grows (i.e., as the DM has more uncertainty a priori about the arm's type, as measured through the variance of the prior): as p_t gets closer to the extremes (zero or one), it takes a much stronger signal to alter the belief as compared to when p_t is close to 0.5. Lastly, as expected, more learning occurs as the DM experiments more aggressively with the risky arm (i.e., as α_t grows). However, such aggressive experimentation also leads to a larger variance in the updates (i.e., “more noise”). Updates also get noisier as the DM has more uncertainty a priori concerning the arm's type (i.e., as $p_t(1 - p_t)$ grows), as the difference in mean rewards under a good and bad arm is larger (i.e., $(\mu_G - \mu_B)^2$ grows), or as the rewards get less noisy (i.e., σ decreases).

Our next result completes the characterization of the DM's belief update, by focusing on the case when a life event occurs during the interval $[t, t + dt)$.

Lemma 2. When a life event occurs during $[t, t + dt)$:

(i) the posterior belief p_{t+dt} conditional on an observed instantaneous reward from the risky arm $d\pi^1 = y$ is given by Bayes' rule and takes a value of

$$p_{t+dt} = (p_t F(\mu_G / \sigma) (1 - e^{-\bar{\lambda}_G dt}) e^{-\bar{\eta}_G dt}) \cdot [p_t F(\mu_G / \sigma) (1 - e^{-\bar{\lambda}_G dt}) e^{-\bar{\eta}_G dt} + (1 - p_t) F(\mu_B / \sigma) (1 - e^{-\bar{\lambda}_B dt}) e^{-\bar{\eta}_B dt}]^{-1}; \quad (5)$$

(ii) the change in the DM's belief $p_{t+dt} - p_t$ is normally distributed, with

$$\begin{aligned} \text{mean} &= \frac{\alpha_t p_t (1 - p_t) (\lambda_G - \lambda_B)}{\bar{\lambda}(p_t)} + \alpha_t p_t (1 - p_t) \bar{\lambda}_G \bar{\lambda}_B \\ &\cdot \frac{\eta_B - \eta_G + (\mu_G - \mu_B) \mu(p_t) / \sigma^2}{(\bar{\lambda}(p_t))^2} dt, \\ \text{variance} &= \alpha_t \left(\frac{p_t (1 - p_t) \bar{\lambda}_G \bar{\lambda}_B (\mu_G - \mu_B)}{\sigma} \right)^2 dt, \end{aligned}$$

where $\bar{\lambda}_G, \bar{\lambda}_B, \bar{\eta}_G, \bar{\eta}_B, F$ are defined in (4a)–(4b), and $\xi(p) \stackrel{\text{def}}{=} p \xi_G + (1 - p) \xi_B, \forall \xi \in \{\bar{\lambda}, \mu\}$.

Part (i) of the result provides an expression for the posterior of the DM's belief, which now depends separately on the rates of life events and stopping events under a good/bad arm. Part (ii) shows that the belief update remains normally distributed, but with modified mean and variance.

The mean update in (ii) now involves two terms. The first term is independent of dt and constitutes a jump in the belief caused by the occurrence of a life event. It can be readily verified that the posterior belief accounting for this jump (i.e., $j(\alpha_t, p_t) \stackrel{\text{def}}{=} p_t + \alpha_t p_t (1 - p_t) (\lambda_G - \lambda_B) / \bar{\lambda}(p_t)$), is increasing in p_t and λ_G , and decreasing in λ_B . This confirms the intuition that, ceteris paribus, the occurrence of a life event makes it relatively more likely that a risky arm is good when the prior belief that it was good was larger, or when life events become more (less) likely under a good (bad) arm. Note that $j(\alpha_t, p_t) < p_t$ if and only if $\lambda_G < \lambda_B$, so that a life event makes it more likely that the arm is good if only if life events are more likely under a good arm than under a bad one. When $\lambda_G < \lambda_B$, it can also be verified that $j(\alpha_t, p_t)$ is decreasing in α_t , so that a DM who experiments more aggressively becomes more skeptical about the risky arm on the occurrence of a life event.

The second term, which is directly proportional to dt , is a further drift in the belief caused by the instantaneous rewards. These rewards also give rise to variability (i.e., variance) in the belief update, and it can be checked that this grows as the DM has more uncertainty a priori concerning the arm's type (i.e., as $p_t(1 - p_t)$ grows), as the good and bad arm differ more in their instantaneous rewards (i.e., as $(\mu_G - \mu_B)^2$

grows), as the processes describing life events get more noisy (i.e., $\lambda_0, \lambda_G, \lambda_B$ grow), or as the rewards get less noisy (i.e., σ decreases).

With these results, we can now provide a characterization of the DM's optimal policy. We restrict our subsequent analysis to the “interesting” case: we allow belief updates to be noisy (i.e., $\mu_G \neq \mu_B$), and we assume that no arm can be eliminated a priori (i.e., a good arm dominates the safe arm, which in turn dominates a bad arm). This is summarized in Assumption 1.

Assumption 1. *The primitives for the framework satisfy*

$$\mu_G \neq \mu_B \quad \text{and} \quad A_B < A_0 < A_G,$$

where $A_\theta \stackrel{\text{def}}{=} (\mu_\theta - \lambda_\theta D + \eta_\theta V)/(r + \eta_\theta)$, $\forall \theta \in \{0, G, B\}$ denote the total rewards per unit time for a safe, good, and bad arm, respectively.

Theorem 1. *Let Assumption 1 hold. Then, the DM's optimal policy is given by*

$$\alpha_t^*(p_t) = \begin{cases} 0 & \text{if } p_t < p^*, \\ 1 & \text{otherwise,} \end{cases} \quad (6)$$

where

$$p^* \stackrel{\text{def}}{=} \frac{w_B(A_0 - A_B)}{w_B(A_0 - A_B) + w_G(A_G - A_0)}, \quad (7a)$$

$$w_B = \frac{r + \eta_B}{r + \eta_B + \lambda_B} v^*, \quad (7b)$$

$$w_G = \frac{r + \eta_G}{r + \eta_G + \lambda_G} (1 + v^*), \quad (7c)$$

and

$$v^* \stackrel{\text{def}}{=} -\frac{1}{2} + \frac{\sigma^2(\lambda_B + \eta_B - \lambda_G - \eta_G)}{(\mu_G - \mu_B)^4} + \left(\frac{((\mu_G - \mu_B)^4 - 2\sigma^2(\lambda_B + \eta_B - \lambda_G - \eta_G))^2 + 8\sigma^2(r + \eta_B + \lambda_B)(\mu_G - \mu_B)^4}{[2(\mu_G - \mu_B)^4]^{-1}} \right)^{1/2}$$

Theorem 1 confirms that the optimal policy is a threshold policy; in particular, fractional allocations are not needed, and the DM can always select a single arm at each point of time.

Note that the optimal threshold p^* only depends on suitably weighted relative differences of the (per unit time) rewards for each arm type, A_0, A_G, A_B . In particular, p^* depends on the safe arm only through A_0 . It can be readily verified that p^* is increasing in μ_0 , reflecting the intuitive fact that, ceteris paribus, a safe arm with higher instantaneous rewards makes the risky arm less appealing. Furthermore, when $\lambda_G \leq \min(\lambda_B, \lambda_0)$ and $\mu_0 \geq \mu_G$ (as in the case of MS), it can also be verified that p^* is decreasing in λ_0 and in D . This shows that a DM behaving optimally should be more prone

to experimenting with a risky arm when life events under the safe arm become more frequent/likely or when they have more severe consequences. The threshold p^* is also strictly increasing in σ and r , confirming that increased volatility and/or an increasing degree of myopic behavior lead to strictly less experimentation with the risky alternative. Finally, note that p^* does not depend on the prior belief p_0 that the arm is good. This is a useful feature in an optimal policy, since it suggests a certain separation between the (objective) effectiveness of an arm and the (potentially subjective) prior.

We conclude this section with a brief discussion of extensions and implications of the results. We start by noting that, although our approach focuses on two arms/treatments, some of the results generalize. In Online Appendix A, we discuss the important case where several arms with binary (good/bad) types exist. Although we are unable to explicitly characterize the optimal policy, we argue that it remains indexable—involving a single arm used at any point of time—and we use our analytical results above to provide an algorithm that calculates the optimal policy to within an arbitrary precision. Our proposed algorithm only requires an offline solution for a small number of one-dimensional convex optimization problems, and an online updating of the beliefs using Lemmas 1 and 2, making it appealing in settings with many arms or frequent updating.

Second, we note that the “bang-bang” structure of the optimal policy relies on several of our modeling assumptions. The “bang-bang” structure no longer holds, for instance, when the response to the DM's allocation is nonlinear (see our Online Appendix D) or when the lump-sum rewards received from life events depend on the risky arm's type (see Online Appendix B). In such cases, a strictly fractional allocation that trades off the benefits of the safe arm with those of the risky arm turns out to be optimal, and this is true even in cases when the risky arm is exactly known to be good or bad.

In the context of chronic diseases that are consistent with our framework, our results suggest that, given our modeling assumptions, the optimal treatment policy is a discontinuation rule: the patient is given the “risky” treatment as long as the belief that she is responding is above a threshold. Once the belief falls below this threshold, the patient is taken off of treatment, and since no “learning” occurs while on the safe treatment exclusively, the process of experimentation essentially stops. We next illustrate how the findings of our simple analytical model can be potentially used for a disease with more complex and realistic dynamics.

4. Case Study: Multiple Sclerosis

We illustrate our framework with a case study of MS. In MS, affected individuals experience increasing disability to the point of becoming bedridden, as well as

blurred vision, muscle weakness, dizziness, fatigue and various sensory abnormalities (Kremenutzky et al. 2006). No cure currently exists, and treatments are only effective for some patients, with no accurate biomarkers for assessing responsiveness a priori (Romeo et al. 2013). While the search for biomarkers in MS remains an active area of research (Derfuss 2012), physicians currently rely mostly on MRI scans and surveys in which patient-reported symptoms are used to compute the EDSS score, which can be translated into quality-of-life utilities for assessing disease evolution and treatment effectiveness (Prosser et al. 2003).

We focus on the most common form of MS, relapsing–remitting MS (RRMS), which comprises about 80% of cases. The initial stage of the disease, which typically lasts for 10 years on average, is characterized by clearly defined relapses that occur on average once per year (Prosser et al. 2004), from which patients may or may not fully recover. After this stage, patients typically enter the progressive stage of the disease, characterized by gradual worsening of disability (Kremenutzky et al. 2006). Typically, relapse rates decrease over time for all patients regardless of treatment, with rates for responders generally lower than for nonresponders (Horakova et al. 2012). Mortality for MS patients depends on both age and current level of disability (Prosser et al. 2004).

Although MS is incurable, DMTs attempt to slow progression and reduce relapses (National Multiple Sclerosis Society 2017). The most common treatments are injectable DMTs such as interferon- β preparations and glatiramer acetate, and more recently oral DMTs such as dimethyl fumarate (approved for use in the United States in 2013), teriflunomide (approved in 2012), fingolimod (approved in 2010), and natalizumab (approved in 2004) (see Rovaris et al. 2001 and National Multiple Sclerosis Society 2017 for more details). Interferon- β is often the first treatment prescribed, as the newer therapies, especially fingolimod and natalizumab, have been associated with an increased risk of severe side effects, such as potentially fatal infections, tumor development, lowering of cardiac rate, and encephalitis (Cohen et al. 2010). The response profile to interferon has been well documented (Horakova et al. 2012, Romeo et al. 2013), but the long-term effectiveness of oral medications has not been established (Carroll 2010).

Our goal in this section is to build a support tool that can inform medical decision makers about the benefits of administering interferon- β treatment in addition to conducting symptom management without DMT. This decision problem is especially important because patients receiving interferon- β experience a significant decrease in quality of life due to side effects, such as pain at the local injection site, flu-like symptoms, depression, and allergic reactions. Furthermore, interferon- β

generates significant and rapidly escalating healthcare costs, currently amounting to \$60,000 per year for each diagnosed case in the United States, and estimated to increase at rates five to seven times larger than prescription drug inflation (Hartung et al. 2015).

Despite the potential benefits, the problem of determining an optimal policy for administering interferon has not received much attention in the medical literature. For a newly diagnosed patient, current guidelines recommend immediately starting treatment and suggest simple discontinuation rules (Río et al. 2011, Cohen et al. 2004). These rules were not the outcome of a quantitative framework and have not been tested for efficiency.

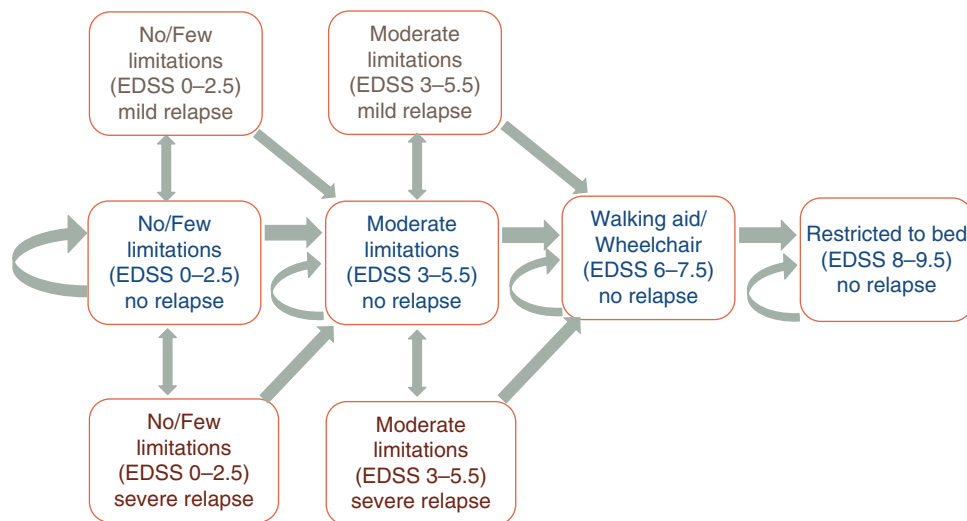
With this motivation, we first describe a detailed disease model from the medical literature, which we use as a basis to design adaptive treatments using our results in Section 3. The goals of our numerical study are (1) to quantify how close existing treatment guidelines are to being optimal and cost-effective, and (2) to understand the potential benefits of using a sophisticated treatment policy that relies on very frequent belief and treatment updating.

4.1. Disease Model

We implement a disease model similar to those in the medical literature (Prosser et al. 2004, Lee et al. 2012). Disease progression is modeled as a Markov chain, with states given by the patient's EDSS score⁸ and whether she is currently experiencing a relapse (see Figure 1 for details).

As in Lee et al. (2012), our simulation follows a hypothetical cohort of 37-year-old RRMS patients with an initial EDSS score of 0–2.5. The cohort includes 10,000 responders and 10,000 nonresponders, consistent with studies documenting the proportion of responders to interferon- β in the population to be around 52% (Horakova et al. 2012). We utilize a one-month time step and simulate patients over a 50-year time horizon.

Each patient can transition from the score of 0–2.5 (no or few limitations) to a score of 3–5.5 (mild to moderate mobility limitations), from there to a score of 6–7.5 (requiring a walking aid), and finally to a score of 8–9.5 (bedridden). While in EDSS states 0–2.5 or 3–5.5, patients can experience relapses, which can be either mild/moderate or severe, and which last for exactly one month, after which they can either remain in their pre-relapse disability level or progress to the next disability level. Once in EDSS state 6–7.5, patients are assumed to have entered the secondary–progressive stage of the disease, characterized by no relapses and gradual destruction of neurons. Consistent with medical studies, we also assume that (a) relapses do not occur in states with EDSS score above 6; (b) the probability that relapses are severe is independent of EDSS state, treatment, or response type; (c) disease progression probabilities are independent of whether

Figure 1. (Color online) MS Disease Model

Notes. All patients start in the lowest disability state (EDSS 0–2.5). In each month, a given patient can (1) remain in the same state, without a relapse; (2) experience a relapse (in the two lowest disability states), which can be either mild or severe, and which lasts exactly one month; or (3) progress to the next level of disability. Values for transition probabilities are shown in Table 1.

the patient is currently experiencing a relapse; and (d) deaths can occur from all states depending on the patient's age, with MS-related deaths only occurring in EDSS state 8–9.5. Furthermore, when simulating patients on treatment, we allow patients in earlier disability states (with EDSS lower than 6) to abandon treatment in any month within the first three years with a fixed probability, consistent with the abandonment rate observed in related medical studies (Cohen et al. 2010, Prosser et al. 2004, O'Rourke and Hutchinson 2005). We assume that once a patient discontinues treatment, she will remain off of treatment for the rest of her life.

Values for the transition probabilities are shown in Table 1 and are consistent with medical studies (Prosser et al. 2004, Horakova et al. 2012, Lee et al. 2012). These probabilities depend on whether the patient is on treat-

ment and on the patient's response type, with successful treatment reducing the disease progression and relapse rates by roughly 50% in responders compared to nonresponders or patients not on treatment. We use age-specific mortality rates published by the U.S. Centers for Disease Control and Prevention (Arias 2014).

Each state in the chain is associated with a mean QALY value capturing a patient's (quality-of-life) utility (Table 2), with a year in perfect health having a utility of one and death having a utility of zero. The realized QALYs in a given state are normally distributed, with a variance that is consistent with quality-of-life surveys (Prosser et al. 2003, Lee et al. 2012). Mean utilities of low-disability states are higher than those of high-disability states, and mean utilities of nonrelapse states are higher than those of relapse states. Furthermore, being on treatment *reduces* the QALYs associated

Table 1. Transition Probabilities for the Disease Model in Figure 1

Parameter	Value	Range	Source
If not in treatment or nonresponder			
Monthly probability of disease progression			
EDSS 0–2.5	0.004438	0.0033–0.0055	(1)
EDSS 3–5.5	0.009189	0.0070–0.0115	(1)
EDSS 6–7.5	0.003583	0.0027–0.0045	(1)
EDSS 9–9.5	0.000952	0.0007–0.0012	(1)
Monthly probability of relapse (EDSS states 0–2.5, 3–5.5)	0.0799	0.0566–0.0944	(1, 2, 3)
Conditional probability of severe relapse	0.23	0.14–0.56	(1)
Responder on treatment			
Probability of progression (relative to nonresponder)	0.5	0.38–1.00	(2)
Probability of relapse (relative to nonresponder)	0.5	0.33–0.90	(2)
Monthly probability of treatment discontinuation	0.0087	0–0.0174	(1)

Sources. (1) Lee et al. (2012), (2) Horakova et al. (2012), (3) Prosser et al. (2004).

Table 2. Utility Values for Each Disease State

Parameter	Value	Range	Source
Utility means (in QALYS per month)			
<i>Baseline utilities by disability level</i>			
EDSS 0–2.5	0.0687	0.0515–0.0833	(1)
EDSS 3–5.5	0.0566	0.0424–0.0708	(1)
EDSS 6–7.5	0.0444	0.0333–0.0555	(1)
EDSS 9–9.5	0.0409	0.0307–0.0512	(1)
Reduction in utility from treatment in first six months	0.0096	0.0038–0.0154	(1)
Reduction in utility from treatment after first six months	0	0–0.0096	(1,2)
<i>Change in utility on treatment, due to response type</i>			
Responder	+0.00058	0–0.002	—
Nonresponder	–0.00058	–0.002–0	—
Reduction in utility from relapse (in QALYS)			
Mild or moderate	0.0076	0.0053–0.0099	(1)
Severe	0.0252	0.0198–0.0305	(1)
Utility standard deviation (in QALYS per month ^a)	0.0087	0.0034–0.0262	(2)

Sources. (1) Lee et al. (2012), (2) Prosser et al. (2003).

^aValues were converted from yearly to monthly values.

with each state for both responders and nonresponders, due to side effects, and this effect is more pronounced during the first six months of treatment. To the best of our knowledge, no study reports a difference in quality of life between responders and nonresponders on treatment. We assume that a responder has a small increase in quality of life compared to a nonresponder (0.0012 per month on average), and we vary this value in sensitivity analysis.

In addition to QALYs, each state also has an associated cost (Table 3), representing the direct and indirect monthly costs—which occur regardless of whether a patient is on treatment—as well as the cost of interferon- β treatment and the cost of managing (severe) relapses.

4.2. Treatment Policies

Consistent with medical practice, we assume that all treatments apply standard care and symptom man-

Table 3. Direct, Indirect, and Treatment Costs for MS

Monthly costs (in USD)	Value	Range
Interferon- β treatment	2,061	1,000–3,828
<i>Direct costs by disability level</i>		
EDSS 0–2.5	536	402–607
EDSS 3–5.5	1,037	778–1,296
EDSS 6–7.5	2,460	1,845–3,075
EDSS 9–9.5	4,327	3,245–5,408
<i>Direct costs per relapse</i>		
Mild or moderate	104	0–200
Severe	5,215	3,911–6,519
<i>Indirect costs by disability level</i>		
EDSS 0–2.5	1,421	1,066–1,776
EDSS 3–5.5	2,964	2,223–3,705
EDSS 6–7.5	3,124	2,343–3,905
EDSS 9–9.5	3,182	2,387–3,978

Sources. Lee et al. (2012), Noyes et al. (2011).

agement throughout the entire lifetime of the patient. Thus, treatment policies choose between two arms, with the “safe” arm corresponding to standard care and the “risky” arm corresponding to prescribing interferon- β in addition to standard care. We consider the following policies:

- *No treatment*: A policy that does not prescribe interferon.
- *Standard*: A policy that immediately starts all patients on interferon- β and only discontinues treatment when patients reach EDSS disability state 6–7.5. This policy is consistent with current recommendations to maintain patients on treatment indefinitely (Río et al. 2011) and has been modeled similarly in previous studies (Lee et al. 2012, Prosser et al. 2004);
- *Consensus criteria*: A policy proposed by Cohen et al. (2004), where all patients are started on interferon- β , but treatment is discontinued if patients experience two or more relapses in a year, or progress to an EDSS state of 6–7.5;
- *Adaptive*: A set of adaptive treatment policies based on our model, which we describe next.

We henceforth refer to the first three policies above (no treatment, standard, and consensus criteria) as *treatment guidelines*.

4.2.1. Implementation of Adaptive Treatments. Our adaptive policies are derived by applying the analytical results in Section 3 to the Markov model in Section 4.1. Since this procedure is required whenever implementing our framework for a complex disease model, we elaborate on the critical steps below.

We identify “life events” with relapses and the “stopping event” with disease progression. We apply our model on a *disease state* basis, and in *age-dependent* fashion. Consistent with practice, we assume that interferon- β is not prescribed for patients with EDSS scores exceeding 6 and thus design treatment only for patients

with EDSS scores of 0–2.5 or 3–5.5, which we henceforth refer to as s_1 and s_2 , respectively. Since mortality is age-dependent, we allow treatment decisions to depend on the patient’s age when transitioning into a particular EDSS state. Applying our model thus generates policies in the form of belief thresholds $p^*(s, x)$ for both EDSS scores $s \in \{s_1, s_2\}$ and for various initial age values x . This requires fitting our model parameters—the instantaneous reward rates $\mu_{0,G,B}$, the standard deviation σ , the relapse rates $\lambda_{0,G,B}$, the disutility associated with a relapse D , the stopping event rates $\eta_{0,G,B}$, the terminal lump-sum reward V , and the discount rate r —for each relevant state, as we discuss next.

Objective. We follow the typical approach in health-care economics to combine the two objectives of QALYs and costs into a single objective, using the net monetary benefit (NMB) conversion (Hunink et al. 2014, Drummond 2005). This requires predefining a WTP threshold, which reflects how much policymakers are willing to spend to gain one additional QALY. Then, the NMB of an intervention that achieves additional QALYs of Q at cost C is calculated as $\text{NMB} = Q \times \text{WTP} - C$. Policies can then be compared in terms of their associated NMB, with a higher WTP value shifting the weight from costs to QALYs. We adopt this objective and solve the model for various values of the WTP parameter; this allows us to recover an entire frontier of cost-effective treatments.

Instantaneous rewards (μ, σ). A month in each disease state $s \in \{s_1, s_2\}$ is associated with an average QALY value and an average cost. We therefore have

$$\mu(s) = (\text{mean monthly QALY in } s) \\ \times \text{WTP} - (\text{monthly costs in } s).$$

For instance, for a patient in EDSS state 3–5.5 who is not on treatment, the reward $\mu_0(s_2)$ is determined by a baseline utility of 0.0566 QALYs/month (per Table 2) and a baseline cost of \$4,001/month (\$1,037 + \$2,964, per Table 3). For all patients on treatment, we use the instantaneous reward rates based on the QALYs after the first six months. For instance, a responder on treatment in EDSS state 3–5.5 has a reward rate $\mu_B(s_2)$ determined by a mean QALY of $0.0566 - 0.001 + 0.00058$ (per Table 2) and a monthly cost of \$4,001 + \$2,061 (per Table 3). We set σ equal to the standard deviation of utility in Table 2.

Disutility from relapse (D). Each relapse causes a decrement in the utility and an increment in the cost for the month in which it occurs, by the amounts listed in Table 2. Therefore,

$$D = (\text{decrement in monthly QALY}) \\ \times \text{WTP} + (\text{increment in monthly cost}).$$

For instance, if a severe relapse occurs while in EDSS 0–2.5, the mean QALY per month decreases by 0.0252, and the cost increases by \$5,215.

Relapse and progression rates (λ, η). Table 1 provides the values for the monthly probabilities of relapse and progression in the disease Markov model. Using our assumption that relapse and progression events occur according to Poisson processes, these monthly probabilities can be converted into monthly rates using the relationship $p_{\text{monthly}} = 1 - e^{-\text{rate}_{\text{monthly}}}$. We note that the rates for responders on treatment are lower, according to Table 1.

Discount rate (r). We take a societal perspective, aggregating costs and QALYs across all patients, and discount at an annual rate of 3% (Gold 1996). We incorporate patient death by viewing it as an exponentially distributed event with a terminal reward equal to zero. Since the mortality rate is the same for both responders and nonresponders, we can easily account for the death event by directly adding the mortality rate to the discount rate (this can be verified formally in the context of Theorem 1). For simplicity, since mortality rates increase with age, we set the mortality for a patient of age x equal to the average mortality rate over years $x + 1, x + 2, \dots, x + \tau$, where τ is the average time spent in a state by responders. The discount rates used in each state $s \in \{s_1, s_2\}$ thus depend on the patient’s age x on initially transitioning into state s .

Terminal reward (V). The lump-sum terminal reward in any state corresponds to the expected NMB on transitioning from that state to the next disease state. Our implementation requires such rewards for any state $s \in \{s_1, s_2\}$ and for every initial patient age x on entering state s , i.e., we need to specify $V(s, x)$. Rewards may also depend on the patient’s response type θ when the patient is subjected to treatment. To account for this, we calculate the terminal rewards separately for responders and nonresponders (i.e., assuming perfect identification), and then weight these by the probability of the patient being a responder.

To determine $V(s_2, x)$, we first simulate the Markov model separately for each type θ , calculating the expected remaining QALYs $Q_{6-7.5}(x + \tau_\theta)$ and costs $C_{6-7.5}(x + \tau_\theta)$ from the random time τ_θ when the patient transitions into the next disease stage (with EDSS 6–7.5) until her death. We then set $V(s_2, x) = \mathbb{E}_{p_\theta} [Q_{6-7.5}(x + \tau_\theta) \times \text{WTP} - C_{6-7.5}(x + \tau_\theta)]$, where the weights are taken with respect to the prior probability that the patient is a responder. For $V(s_1, x)$, we proceed similarly, calculating rewards from the random transition into state s_2 onwards.

Initial prior probability of the patient being a responder.

We start every patient in state s_1 with a prior of 0.52 of being a responder, in accordance with the distribution of responders and nonresponders in the population. When simulating our adaptive policies, we update this belief as long as the patient is on treatment. The updates are done monthly, depending on the

observed quality-of-life utility and whether a relapse occurred during the month, using the results in Lemmas 2(i) and 1(i), respectively.

We note that, although our implementation captures some of the features of the complex MS model in Figure 1, such as disease progression and age-dependent mortality, it nonetheless remains an approximation. For instance, it ignores the different magnitudes of side effects in the first six months of treatment, the different variances in QALYs in a relapse month, and the patient abandonment rates, and does not rigorously account for mortality rates. However, our simulation accounts for all of these features, making the performance assessment for all policies under consideration considerably more accurate.

4.3. Results

We start by determining optimal adaptive treatments based on our implementation, for different values of the WTP parameter. This generates an efficient frontier of policies that systematically trade off QALYs and costs, and is useful for benchmarking potential alternatives. We consider WTP values from \$50,000/QALY to \$800,000/QALY, consistent with MS studies that report costs in excess of \$500,000/QALY gained (Noyes et al. 2011). For each WTP, we find the optimal belief thresholds at which treatment should be discontinued for each relevant EDSS state (0–2.5 and 3–5.5), and for every patient age.

Table 4 shows these thresholds for a typical 37-year-old patient. As expected, we find that the propensity to recommend treatment increases with WTP and with age. Note that for a typical patient starting with an EDSS of 0–2.5 and a 52% prior probability of being a responder, the optimal adaptive treatment would prescribe interferon only for a WTP above \$200,000/QALY. Since interventions are generally considered cost-effective when the cost per QALY is less than three times the country’s per-capita GDP (Drummond 2005, Hunink et al. 2014), this suggests that interferon treatment might not be considered cost-effective, and that a no-treatment policy is optimal if the WTP is no more than \$150,000/QALY.

Table 4. Optimal Discontinuation Thresholds for a 37-Year-Old Patient in State s_1 (EDSS Score 0–2.5) and State s_2 (EDSS Score 3–5.5), for Various WTP Values

WTP (\$/QALY) (K)	$p^*(s_1, 37)$	$p^*(s_2, 37)$
50	1	1
100	1	1
150	0.77	1
200	0.48	0.67
250	0.36	0.45
300	0.28	0.34
500	0.16	0.18
800	0.10	0.11

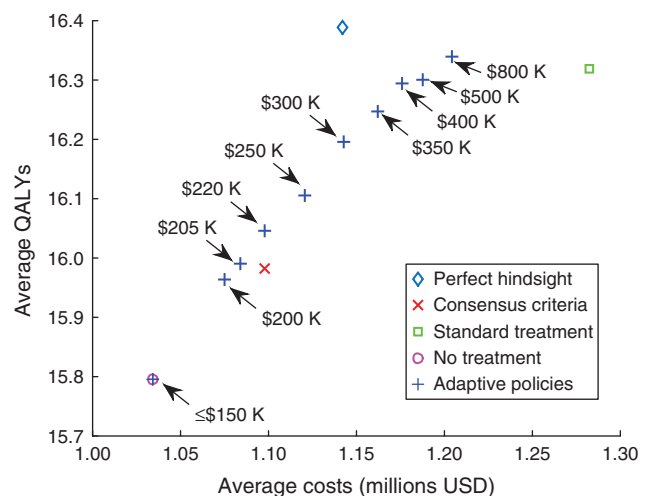
We next simulate all of the policies under consideration—no-treatment, standard, consensus, and all of our adaptive policies—using the detailed Markov disease model. It is important to emphasize that under this simulation, all of the policies and information sets/beliefs are updated on a *monthly* basis, so that all of the results correspond to the realized performance under this frequency. In particular, although all of our adaptive policies were calculated under the assumption of a continuous-time model, they are implemented and assessed under discrete-time updates.

A visual summary of the results is shown in Figure 2, which displays the costs and QALYs per patient averaged over responders and nonresponders, assuming a 52% fraction of responders in the population. To put the results into perspective, the figure also displays the performance of a *perfect hindsight* policy, which correctly classifies all patients a priori and only prescribes interferon to responders. In our simulation, this policy would result in 16.389 QALYs and costs of \$1,143,031.

The no-treatment policy yields the smallest number of QALYs on average (15.794) but is also the least expensive (\$1,036,656). The standard policy yields an average of 16.333 QALYs, an increase of 3.4% compared to no-treatment, but is also the most expensive of all policies, with an average cost of \$1,281,692 per patient. The consensus criteria policy falls in between, achieving 15.984 QALYs at a cost of \$1,096,950.

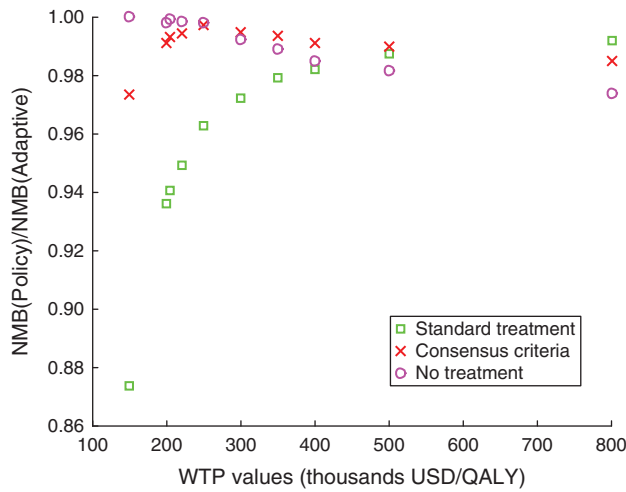
As expected, the adaptive policies form an efficient frontier that dominates all policies except the perfect-hindsight one. At WTP values below \$150,000/QALY, a no-treatment policy is equivalent to adaptive policies. The consensus policy is strictly dominated by adaptive policies for a WTP value of \$205,000–\$220,000/QALY, although the differences are not very substantial (with

Figure 2. (Color online) Population-Averaged Costs and QALYs for All Policies



Note. The arrows indicate the values of the WTP parameter under which the adaptive policies were obtained.

Figure 3. (Color online) Net Monetary Benefits (NMB) Achieved by Each Policy Relative to the Optimal Adaptive Policy



QALYs increased by 0.4% or costs reduced by 1.2%). The standard policy is dominated by an adaptive policy with a WTP of \$800,000/QALY, which increases QALYs slightly (by 0.006) but significantly reduces costs (by \$77,000, or 6%).

Figure 3 compares the policies in terms of their achieved NMB, expressed as a percentage of the NMB of the best adaptive policy. Consistent with our prior observations, note that the highest NMB is achieved by the no-treatment policy at low WTP, by the consensus policy at intermediate WTP, and by the standard policy at high WTP. Furthermore, all policies except standard achieve a high NMB uniformly (i.e., for any WTP value).

These results can be used by policy makers to quantify the benefits of interferon treatment and weigh them against the corresponding costs. Our findings suggest that gains from interferon treatment are not large in absolute terms and come at steep costs: even the best adaptive policy can increase QALYs by only 3.45% relative to the no-treatment alternative, while increasing costs by 16.2%.⁹ We find that the WTP required for such improvements exceeds \$800,000/QALY; this

confirms earlier studies reporting costs larger than \$500,000/QALY for interferon (Noyes et al. 2011), showing that this persists even when considering optimal adaptive treatments instead of heuristic treatment policies. This reinforces our earlier observation that interferon treatment is not necessarily cost-effective and suggests that even in environments with larger WTP (e.g., above \$150,000/QALY), not prescribing interferon may be the optimal action. However, this recommendation should be interpreted with caution—even though improvements in QALYs may not be large in absolute terms, they may nonetheless be significant in relative terms, and particularly for chronic diseases as debilitating as MS. Furthermore, patients (and policy makers alike) may not easily accept the cost-benefit analysis inherent in such a no-treatment recommendation.

Second, our results provide empirical validation of the consensus criteria proposed by Cohen et al. (2004). We find that the resulting policy is close to being efficient at intermediate values of WTP and achieves net monetary benefits close to a fully adaptive policy. Thus, these simple discontinuation rules may represent a viable alternative to implementing a complex optimal adaptive policy, particularly at intermediate values of WTP.

Finally, the results suggest that none of the treatment guidelines is satisfactory at very large WTP: the consensus and no-treatment policies generate low QALYs, while the standard policy is inefficient, dominated in both QALYs and costs by an adaptive policy. To better illustrate the differences in performance, and to understand how the benefits are distributed among responders and nonresponders, we examine this case in more detail.

4.3.1. An Optimal Policy at High WTP. We compare the three treatment guidelines with the adaptive policy calculated for a WTP of \$800,000/QALY (Online Appendix E provides a detailed description of this policy). The simulation results by response type are summarized in Table 5. As can be seen, the optimal adaptive policy achieves QALYs for responders that are close to those of the standard policy (which is opti-

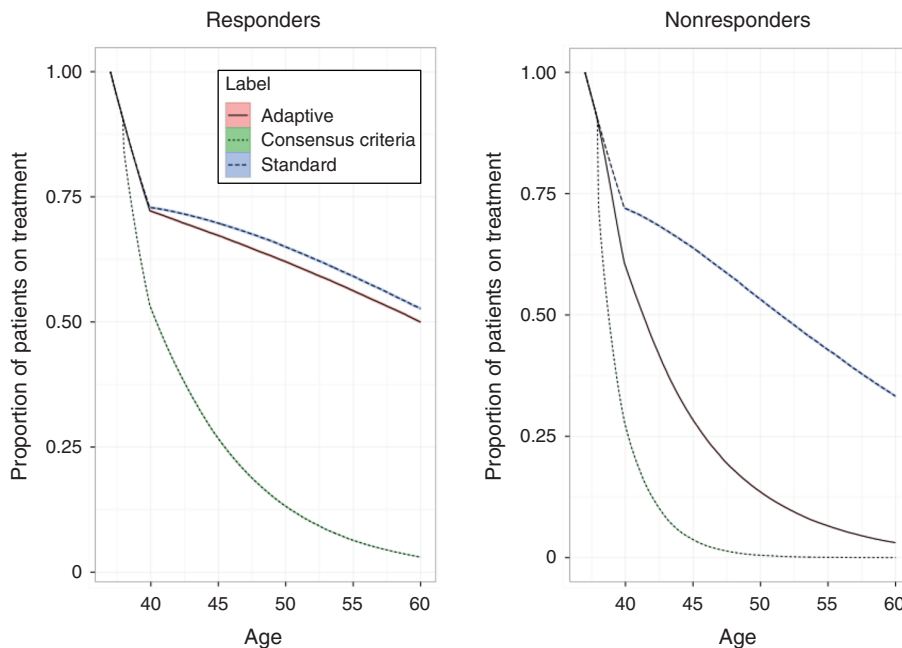
Table 5. QALYs and Costs for Each Policy, by Patient Type

	Standard		No-treatment		Consensus		Adaptive (800,000)	
	Resp	Nonresp	Resp	Nonresp	Resp	Nonresp	Resp	Nonresp
QALY means	16.939	15.677	15.794	15.794	16.225	15.724	16.917	15.712
QALY stderr	0.037	0.035	0.035	0.035	0.035	0.035	0.036	0.036
Cost means (\$)	1,241,224	1,325,533	1,036,656	1,036,656	1,099,763	1,150,516	1,236,178	1,170,188
Cost stderr (\$)	3,778	4,095	4,214	4,214	3,963	4,227	3,912	4,208
NMB means (\$)	12,309,976	11,216,067	11,598,544	11,598,544	11,880,237	11,428,684	12,297,348	11,399,254
NMB stderr (\$)	15,673	15,311	17,026	17,026	16,346	17,162	16,159	17,210
NMB means (\$)	11,784,899		11,598,544		11,663,491		11,866,262	
NMB stderr (\$)	13,377		17,026		14,615		14,586	

Notes. “(Non)Resp” denotes a (non)responder to interferon- β . The adaptive policy and all NMBs are calculated for a WTP of \$800,000.

Downloaded from informs.org by [171.67.216.21] on 15 November 2017, at 08:00. For personal use only, all rights reserved.

Figure 4. (Color online) Proportion of Patients on Treatment Until Age 60 Under the Consensus, Standard and Adaptive Policy (for $WTP = \$800,000/QALY$), for Responders and Nonresponders



mal for this type), and considerably exceeds the no-treatment and consensus policies, by 6.6% and 4.1%, respectively. At the same time, the adaptive policy also achieves higher QALYs than the standard policy for nonresponders, by identifying them and removing them from treatment earlier; no-treatment and consensus, which are both more aggressive in removing patients from treatment, are only marginally superior to the adaptive policy for nonresponders, with QALY improvements of less than 0.6%.

The outcomes of the policies are primarily driven by the number of months each patient spends in treatment. The standard policy incurs the most costs, as it keeps patients on treatment for the longest time on average (289 months for responders, and 208 months for nonresponders). In contrast, the adaptive policy keeps responders on treatment for 279 months on average and nonresponders for 78 months, achieving the highest overall gain in QALYs. Perhaps the best illustration of the effectiveness of the adaptive policy is Figure 4, which plots the proportion of patients on treatment over time, by response type. The proportion of nonresponders on treatment converges to zero for the adaptive policy, whereas the proportion of responders on treatment remains relatively high. The optimal policy thus attains a good balance between the optimal treatment for responders (standard) and for nonresponders (no-treatment).

These results can be observed consistently for each year in our simulation, as displayed in Figures 5 and 6. Note that both costs and QALYs decrease over time under both the standard and the adaptive policy and

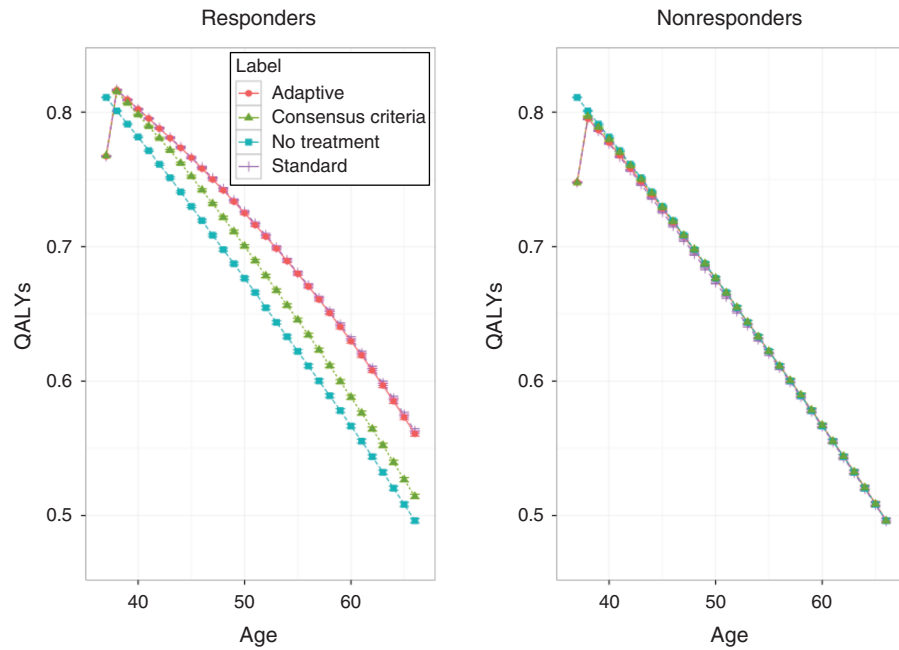
under both response types; this is due to disease progression and the fact that all patients are taken off of treatment once they reach EDSS state 6–7.5. Also, consistent with reality, the disutility incurred by patients due to side effects is higher for the first six months on treatment, which is why both groups display a non-monotonic pattern in the first year.

To test the robustness of our findings, we performed a probabilistic sensitivity analysis in which we randomly generated 1,000 problem instances. For each instance, each parameter was randomly sampled from a triangle distribution with mode given by the base case value, and the lowest and highest values corresponding to the ranges in Tables 1–3. The results of the analysis are shown in Figure 7. In this scatter plot, each point represents the result of a simulation with 10,000 responders and 10,000 nonresponders for a given set of parameters. As can be seen, all statistically significant differences in mean QALYs (and costs) between the adaptive and standard policies were positive (respectively, negative) for nonresponders, indicating that our adaptive policy is especially cost-effective for nonresponders compared to the standard policy.

5. Conclusions, Limitations, and Future Directions

We have introduced a quantitative framework that can inform treatment policies for chronic diseases sharing the following features: (1) there is a priori uncertainty about whether a patient will respond to an available

Figure 5. (Color online) QALYs Experienced Under All Treatment Policies for Responders and Nonresponders: Yearly Means and 95% Confidence Intervals for the Means



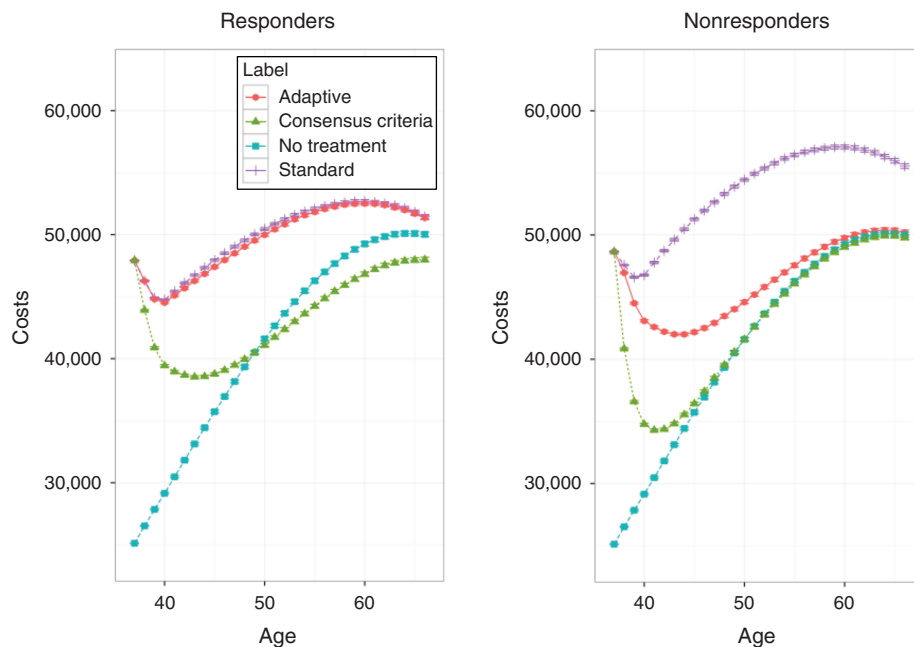
treatment; (2) observations of the effectiveness of treatment are noisy; and (3) learning about treatment effectiveness occurs not only from monitoring day-to-day disease progression, but also from observing the timing and severity of less frequent, major health events.

We showed that the problem of choosing between two treatments with linear dose–response can be analyzed in closed form, resulting in intuitive optimal

policies that take the form of discontinuation rules. We also discussed how our analytical results can be used to optimally select among several treatments by solving a small number of one-dimensional, convex optimization problems, and provided conditions when the optimal treatment is no longer a simple discontinuation rule.

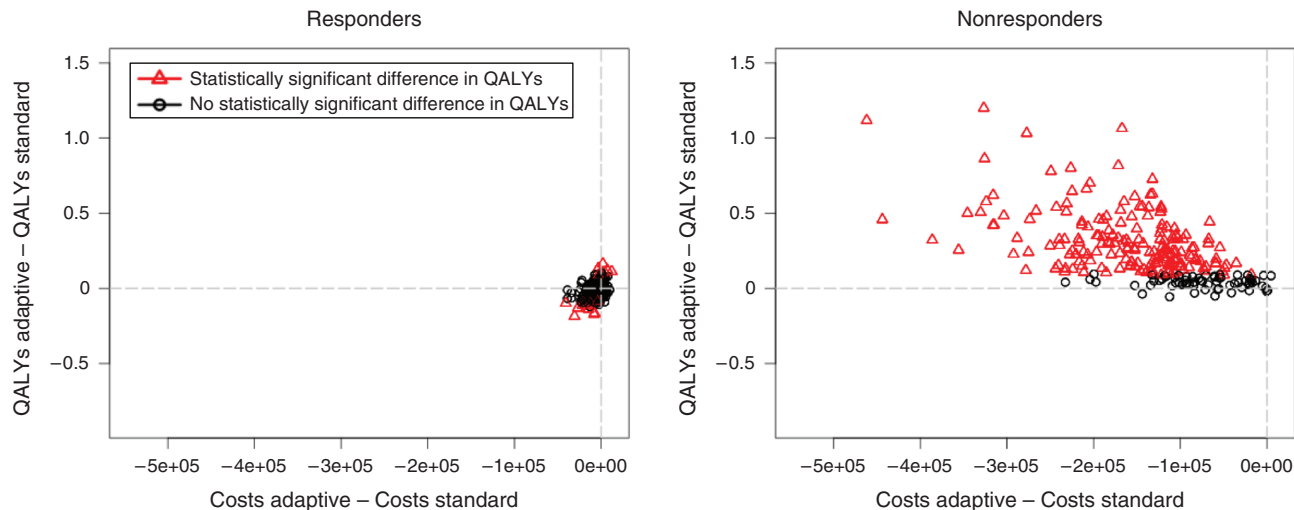
Finally, we used our framework to develop a set of treatment policies for administering interferon to

Figure 6. (Color online) Costs Incurred Under All Treatment Policies for Responders and Nonresponders: Yearly Means and 95% Confidence Intervals for the Means



Downloaded from informs.org by [171.67.216.21] on 15 November 2017, at 08:00 . For personal use only, all rights reserved.

Figure 7. (Color online) Results of Probabilistic Sensitivity Analysis: Incremental Costs and QALYs for Responders and Nonresponders (Adaptive Treatment Policy Compared to Standard Treatment Policy)



patients suffering from MS. Our policies explicitly traded off treatment benefits and costs through a parameter capturing the policy makers' WTP for every QALY gained. Using these policies as benchmarks, we then assessed several treatment guidelines used in practice for administering interferon, which lead to three conclusions that can inform policy makers and medical practitioners:

1. At WTP values below \$150,000/QALY, we found that a no-treatment policy is optimal.

2. At WTP values between \$150,000/QALY and \$500,000/QALY, we found that a policy based on the consensus criteria discussed in Cohen et al. (2004) delivers a good balance between QALYs and costs, and is almost efficient. Considering its simplicity relative to our adaptive policies, it thus emerges as the preferred treatment guideline at intermediate WTP values.

3. At WTP values above \$500,000/QALY, none of the treatment guidelines considered deliver adequate performance; an adaptive policy derived from our framework under a WTP of \$800,000/QALY attained a better balance between administering sufficient treatment to responders and identifying nonresponders early.

Several next steps can bring our research and findings closer to a treatment recommendation. First, our MS case study could be generalized to allow choosing among multiple drugs for symptom management and multiple disease-modifying agents (National Multiple Sclerosis Society 2017). Online Appendix A provides a computationally tractable procedure for a model with multiple risky treatments/arms, which could be a building block in this direction. To enable this approach, one first requires a thorough understanding of the patient response to each drug, so as to calibrate the corresponding reward, relapse, and progression rates. These could be obtained from clinical trials,

such as those conducted as part of the drug approval process.

Second, one could extend our model by relaxing some of the assumptions made for analytical tractability. Of key importance here are the assumptions concerning the linear dose response and the ability to continuously measure rewards and update beliefs, which we discuss next.

Linear dose response. Although exact response curves are the subject of active research, response curves reported in the literature for many drugs tend to be S-shaped, exhibiting diminishing marginal returns at high dosage values. The assumption of linearity may nonetheless remain reasonable within a certain dosage range. For instance, clinical trials with interferon β -1a for MS suggest an approximately linear reduction in relapse rate when the dosage is below 66 micrograms per week, and a decreasing rate for higher dosage (OWIMS 1999).¹⁰ In Online Appendix D, we discuss in detail the impact of the linearity assumption on optimal policies and performance when the underlying dose-response is S-shaped. We find that optimal policies are no longer “bang-bang,” and that a strictly fractional treatment allocation may be optimal even when the patient is known to be a (non)responder. The optimality loss varies from 0% to 16%, depending on the degree of “nonlinearity,” which suggests that embedding nonlinear response curves without sacrificing tractability may be a practically (and theoretically) meaningful future direction.

Continuous updates. Our framework allows for continuously measuring rewards and conducting belief and treatment updates. This is reasonable when the policies generated from our results are interpreted as upper bounds, which are then used either to suggest

or otherwise benchmark simpler treatments with less frequent updates. Depending on the disease and treatment in question, these assumptions may also be (come) realistic. For instance, in MS, the use of wearable devices has shown to have great potential for the collection and relaying of real-time patient information (McIninch et al. 2015).¹¹ Combined with research aimed at understanding how disease progression and treatment response are related to observed mobility,¹² such developments could potentially make a near-continuous-time treatment policy feasible in the future (provided, of course, that the benefits outweigh the costs). Despite these examples, however, assuming continuous evaluations and treatment updates may not be reasonable when extensive medical exams are required (e.g., involving doctor visits, MRI scans, etc.). Our model could be extended to allow belief updates only at particular points in time, provided that the information between these points can be suitably aggregated. In Online Appendix C, we discuss the impact of monitoring frequency in more detail, and provide several theoretical and computational results that characterize the losses under less frequent updating. For our MS case study, we find that the loss from a monthly monitoring policy is less than 8%. However, we also find that as treatments for MS become more efficient at reducing the frequency of relapses in responders, these losses are likely to increase, prompting the need for more research that explicitly captures the costs of more frequent belief and treatment updating.

In addition to these, one other assumption worth relaxing would be the requirement that a risky (treatment) arm have exactly two types. In practice, more types may exist (e.g., corresponding to a patient fully, partially, or not responding to treatment). Our results would readily apply if the optimal treatment for each patient type still involved a binary choice between the same two alternatives, since then the various types could be aggregated into two “macro-types.” When different patient types require different dosages or treatment options, our model would have to be extended to explicitly allow learning for all types simultaneously. This requires a multidimensional state that tracks the probability for each type, which considerably complicates the analysis.

Lastly, an important step in making the results implementable is a clinical trial testing the performance of our adaptive policy against other guidelines. To that end, Online Appendix E provides an implementation-driven description of our proposed policy, which could guide such a design in conjunction with an appropriate selection of a cohort of patients.

To conclude, although we illustrated our framework with a case study on MS and interferon- β , we believe that the ideas could be used to inform the treatment of other chronic diseases, such as celiac disease, rheumatoid arthritis, Crohn’s disease, or depression.

Acknowledgments

The authors are grateful to the department editor, the associate editor, and three anonymous referees, whose comments and suggestions have considerably improved the paper. The authors are particularly thankful to Allie Dunworth Leeper and Stephen Chick for helpful comments. All remaining errors are the authors’ own.

Endnotes

¹ For instance, in an editorial paper, Murphy and Collins (2007, p. S2) state that “[d]espite the activity in evaluating adaptive treatment strategies, the development of data collection and analytic methods that directly inform the construction of adaptive treatment strategies lags behind.”

² Biomarkers exist for some chronic diseases (e.g., breast cancer). Our focus is on diseases for which the existing biomarkers do not perfectly classify patients as responders and nonresponders, and thus where there is scope for experimentation with the available treatments. As we mention below, one such disease is multiple sclerosis. It is worthwhile to note that the discovery of biomarkers for MS is a very active field of research (e.g., Derfuss (2012)).

³ “The effects of current therapies on attack rates and MRI measures of newly accumulated lesion burdens [...] are the events that are most readily available to the clinician when considering treatment failure or suboptimal response in an individual patient” (National Multiple Sclerosis Society 2004, p. 2).

⁴ The National Institute of Health in the United Kingdom launched an innovative risk-sharing scheme in 2002, according to which patients would be closely monitored to evaluate the cost-effectiveness of the drugs used in standard treatment, with an agreement that prices would be reduced if overall patient outcomes were worse than predicted. The scheme became controversial when reports from observational cohorts suggested that the outcomes were far below expectations—implying that treatment was generally not cost-effective—yet the drug providers did not reduce their prices as per the agreement (Raftery 2010, Sudlow and Counsell 2003, Boggild et al. 2009). It is worth noting that personalized discontinuation rules for patients were not considered, though such rules might have reduced total costs and also improved patient outcomes.

⁵ An alternative formulation could have considered the stopping event as a “special instance” of a life event. Since splitting a Poisson process would yield an exponentially distributed time for the stopping event, this would be equivalent to our current model.

⁶ This effectively means that the expectations of quantities at time t that depend on θ should be taken with respect to a corresponding two-point distribution given by p_t . For instance, $\mathbb{E}[\mu_\theta] = p_t \mu_G + (1 - p_t) \mu_B$.

⁷ Formally, \mathcal{F}_t is the sigma-algebra generated by the allocations, rewards, events, and lump-sum rewards up to time t (i.e., $\mathcal{F}_t \stackrel{\text{def}}{=} \sigma(\{a_\tau, d\pi^0(\tau), d\pi^1(\tau), N_\tau, L_\tau\}_{0 \leq \tau < t})$).

⁸ We choose to focus on EDSS instead of MRI in our study for several pragmatic reasons. First, EDSS is considerably more widespread, and there is no consensus in the medical community concerning the use of MRI for monitoring therapeutic response in MS (see Cohen et al. 2004). Second, MRI scans may not be available for a large subset of the population or may be difficult or costly to administer frequently. Third, there is insufficient data in medical studies concerning the difference in MRI scans between responders and nonresponders.

⁹ Similar observations can be made with respect to the perfect hindsight policy, which increases QALYs by 3.7% and costs by 10.2% relative to no-treatment.

¹⁰ The Once Weekly Interferon for MS Study Group reports a reduction in relapse rates of 9.6%, 19%, 33%, and 37% for respective weekly dosages of 30 μg , 44 μg , 66 μg , and 132 μg (OWIMS 1999).

¹¹In a recent study conducted by the nonprofit PatientsLikeMe and Biogen Idec, 248 FitBit One™ devices were distributed to patients suffering from MS, and the personal mobility data of all the patients was collected and sent to centralized data servers. The results of the study were reported in the 67th American Academy of Neurology's Annual Meeting (April, 2015), revealing "a high degree of patient interest and perceived value in using activity tracking devices to help patients manage their MS" (McIninch et al. 2015).

¹²This research endeavor has recently been taken up in a collaboration by Biogen, Google X, and Cleveland Clinic (Chen 2015).

References

- Ahuja V, Birge J (2016) Response-adaptive designs for clinical trials: Simultaneous learning from multiple patients. *Eur. J. Oper. Res.* 248(2):619–633.
- Almiral D, Compton S, Gunlicks-Stoessel M, Duan N, Murphy S (2012) Designing a pilot sequential multiple assignment randomized trial for developing an adaptive treatment strategy. *Statist. Medicine* 31(17):1887–1902.
- Arias E (2014) United States life tables, 2010. *Natl. Vital Statist. Rep.* 63(7), http://www.cdc.gov/nchs/data/nvsr/nvsr63/nvsr63_07.pdf.
- Bank P, Küchler C (2007) On Gittins' index theorem in continuous time. *Stochastic Processes Their Appl.* 117(9):1357–1371.
- Berry D (1978) Modified two-armed bandit strategies for certain clinical trials. *J. Amer. Statist. Assoc.* 73(362):339–345.
- Berry D, Fristedt B (1985) *Bandit Problems: Sequential Allocation of Experiments* (Chapman and Hall, London).
- Berry D, Pearson L (1985) Optimal designs for clinical trials with dichotomous responses. *Statist. Medicine* 4(4):497–450.
- Bertsimas D, O'Hair A, Relyea S, Silberholz J (2016) An analytics approach to designing combination chemotherapy regimens for cancer. *Management Sci.* 62(5):1511–1531.
- Boggild M, Palace J, Barton P, Ben-Shlomo Y, Bregenzer T, Dobson C, Gray R (2009) Multiple sclerosis risk sharing scheme: Two year results of clinical cohort study with historical comparator. *BMJ* 339:b4677.
- Bolton P, Harris C (1999) Strategic experimentation. *Econometrica* 67(2):349–374.
- Carroll W (2010) Oral therapy for multiple sclerosis—Sea change or incremental step? *New England J. Medicine* 362(5):456–458.
- Chen C (2015) Google, Biogen seek reasons for advance of multiple sclerosis. *Bloomberg* (January 27), <http://www.bloomberg.com/news/articles/2015-01-27/google-biogen-look-reasons-for-advance-of-multiple-sclerosis>.
- Cohen A, Solan E (2013) Bandit problems with Lévy processes. *Math. Oper. Res.* 38(1):92–107.
- Cohen BA, Khan O, Jeffery DR, Bashir K, Rizvi SA, Fox EJ, Agius M, et al. (2004) Identifying and treating patients with suboptimal responses. *Neurology* 63(12, Suppl. 6):S33–S40.
- Cohen J, Barkhof F, Comi G, Hartung H, Khatir B, Montalban X, Pelletier J, et al. (2010) Oral fingolimod or intramuscular interferon for relapsing multiple sclerosis. *New England J. Medicine* 362(5):402–415.
- Denton B, Kurt M, Shah N, Bryant S, Smith S (2009) Optimizing the start time of statin therapy for patients with diabetes. *Medical Decision Making* 29(3):351–367.
- Derfuss T (2012) Personalized medicine in multiple sclerosis: Hope or reality? *BMC Medicine* 10(1):116.
- Drummond M (2005) *Methods for the Economic Evaluation of Health Care Programmes* (Oxford University Press, Oxford, UK).
- El Karoui N, Karatzas I (1994) Dynamic allocation problems in continuous time. *Ann. Appl. Probab.* 4(2):255–286.
- Gold M (1996) *Cost-Effectiveness in Health and Medicine* (Oxford University Press, New York).
- Harrison JM, Sunar N (2015) Investment timing with incomplete information and multiple means of learning. *Oper. Res.* 62(2):442–457.
- Hartung DM, Bourdette DN, Ahmed SM, Whitham RH (2015) The cost of multiple sclerosis drugs in the us and the pharmaceutical industry too big to fail? *Neurology* 84(21):2185–2192.
- Helm JE, Laviere MS, Van Oyen MP, Stein JD, Musch DC (2015) Dynamic forecasting and control algorithms of glaucoma progression for clinician decision support. *Oper. Res.* 63(5):979–999.
- Horakova D, Kalincik T, Dolezal O, Krasensky J, Vaneckova M, Seidl Z, Havrdova E (2012) Early predictors of non-response to interferon in multiple sclerosis. *Acta Neurologica Scandinavica* 126(6):390–397.
- Hunink M, Weinstein M, Wittenberg E, Drummond M, Pliskin J, Wong J, Glasziou P (2014) *Decision Making in Health and Medicine: Integrating Evidence and Values* (Cambridge University Press, Cambridge, UK).
- Keller G, Rady S (2010) Strategic experimentation with Poisson bandits. *Theor. Econom.* 5(2):275–311.
- Keller G, Rady S (2015) Breakdowns. *Theor. Econom.* 10(1):175–202.
- Keller G, Rady S, Cripps M (2005) Strategic experimentation with exponential bandits. *Econometrica* 73(1):39–68.
- Kremenchtzky M, Rice GPA, Baskerville J, Wingerchuk DM, Ebers GC (2006) The natural history of multiple sclerosis: A geographically based study—9: Observations on the progressive phase of the disease. *Brain* 129(3):584–594.
- Lee S, Baxter D, Limone B, Roberts M, Coleman C (2012) Cost-effectiveness of fingolimod versus interferon beta-1a for relapsing remitting multiple sclerosis in the United States. *J. Medical Econom.* 15(6):1088–1096.
- Mandelbaum A (1987) Continuous multi-armed bandits and multi-parameter processes. *Ann. Probab.* 15(4):1527–1556.
- Mariette X, Matucci-Cerinic M, Pavelka K, Taylor P, van Vollenhoven R, Heatley R, Walsh C, Lawson R, Reynolds A, Emery P (2011) Malignancies associated with tumour necrosis factor inhibitors in registries and prospective observational studies: A systematic review and meta-analysis. *Ann. Rheumatic Diseases* 70(11):1895–1904.
- Mason J, Denton B, Shah N, Smith S (2014) Optimizing the simultaneous management of blood pressure and cholesterol for type 2 diabetes patients. *Eur. J. Oper. Res.* 233(3):727–738.
- McIninch J, Datta S, DasMahapatra P, Chiauzzi E, Bhalerao R, Spector A, Goldstein S, Morgan L, Relton J (2015) Remote tracking of walking activity in MS patients in a real-world setting. *Neurology* 84(14, Suppl. P3.209), https://patientslikeme_posters.s3.amazonaws.com/2015_Remote%20Tracking%20of%20Walking%20Activity%20in%20MS%20Patients%20in%20Real%20World.pdf.
- Murphy S (2005) An experimental design for the development of adaptive treatment strategies. *Statist. Medicine* 24(10):1455–1481.
- Murphy SA (2003) Optimal dynamic treatment regimes. *J. Roy. Statist. Soc. Ser. B Statist. Methodol.* 65(2):331–355.
- Murphy SA, Collins LM, Rush AJ (2007) Customizing treatment to the patient: Adaptive treatment strategies. *Drug Alcohol Dependence* 88(Suppl. 2):S1–S3.
- National Multiple Sclerosis Society (2004) Changing therapy in relapsing multiple sclerosis: Considerations and recommendations of a task force of the National Multiple Sclerosis Society. Expert opinion paper, http://www.nationalmssociety.org/NationalMSSociety/media/MSNationalFiles/Brochures/Clinical_Bulletin_Changing-Therapy-in-Relapsing-MS.pdf.
- National Multiple Sclerosis Society (2008) Disease management consensus statement. Expert opinion paper, http://www.nationalmssociety.org/NationalMSSociety/media/MSNationalFiles/Brochures/ExpOp_Consensus.pdf.
- National Multiple Sclerosis Society (2017) Disease-modifying therapies for MS. <http://stage.nationalmssociety.org/NationalMSSociety/media/MSNationalFiles/Brochures/Brochure-The-MS-Disease-Modifying-Medications.pdf>.
- Noyes K, Bajorska A, Chappel A, Schwid S, Mehta L, Weinstock-Guttman B, Holloway R, Dick A (2011) Cost-effectiveness of disease-modifying therapy for multiple sclerosis: A population-based study. *Neurology* 77(4):355–363.

- Once Weekly Interferon for MS Study Group (OWIMS) (1999) Evidence of interferon β -1a dose response in relapsing-remitting MS: The OWIMS study. *Neurology* 53(4):679–686.
- O'Rourke KE, Hutchinson M (2005) Stopping beta-interferon therapy in multiple sclerosis: An analysis of stopping patterns. *Multiple Sclerosis* 11(1):46–50.
- Phillips CJ (2004) The cost of multiple sclerosis and the cost effectiveness of disease-modifying agents in its treatment. *CNS Drugs* 18(9):561–574.
- Pincus T, Callahan L, Sale W, Brooks A, Payne L, Vaughn W (1984) Severe functional declines, work disability, and increased mortality in seventy-five rheumatoid arthritis patients studied over nine years. *Arthritis Rheumatism* 27(8):864–872.
- Pineau J, Bellemare MG, Rush AJ, Ghizaru A, Murphy SA (2007) Constructing evidence-based treatment strategies using methods from computer science. *Drug Alcohol Dependence* 88(Suppl. 2):S52–S60.
- Powell WB, Ryzhov IO (2012) *Optimal Learning* (John Wiley & Sons, Hoboken, NJ).
- Prosser L, Kuntz K, Bar-Or A, Weinstein M (2003) Patient and community preferences for treatments and health states in multiple sclerosis. *Multiple Sclerosis* 9(3):311–319.
- Prosser L, Kuntz K, Bar-Or A, Weinstein M (2004) Cost-effectiveness of interferon beta-1a, interferon beta-1b, and glatiramer acetate in newly diagnosed non-primary progressive multiple sclerosis. *Value Health* 7(5):554–568.
- Raftery J (2010) Multiple sclerosis risk sharing scheme: A costly failure. *BMJ* 340:c1672.
- Río J, Comabella M, Montalban X (2011) Multiple sclerosis: Current treatment algorithms. *Current Opinion Neurol.* 24(3):230–237.
- Romeo M, Martinelli-Boneschi F, Rodegher M, Esposito F, Martinelli V, Comi G, San Raffaele Multiple Sclerosis Clinical Group (2013) Clinical and MRI predictors of response to interferon-beta and glatiramer acetate in relapsing-remitting multiple sclerosis patients. *Eur. J. Neurol.* 20(7):1060–1067.
- Rovaris M, Comi G, Rocca M, Wolinsky J, Filippi M, European/Canadian Glatiramer Acetate Study Group (2001) Short-term brain volume change in relapsing-remitting multiple sclerosis: Effect of glatiramer acetate and implications. *Brain* 124(9):1803–1812.
- Rudick R, Stuart W, Calabresi P, Confavreux C, Galetta S, Radue E, Lublin F, et al. (2006) Natalizumab plus interferon beta-1a for relapsing multiple sclerosis. *New England J. Medicine* 354(9):911–923.
- Scalfari A, Neuhaus A, Degenhardt A, Rice G, Muraro P, Daumer M, Ebers G (2010) The natural history of multiple sclerosis, a geographically based study 10: Relapses and long-term disability. *Brain* 133(7):1914–1929.
- Sudlow C, Counsell C (2003) Problems with UK government's risk sharing scheme for assessing drugs for multiple sclerosis. *British Medical J.* 326(7385):388–392.
- Tappenden P, McCabe C, Chilcott J, Simpson E, Nixon R, Madan J, Fisk JD, Brown M (2009) Cost-effectiveness of disease-modifying therapies in the management of multiple sclerosis for the Medicare population. *Value Health* 12(5):657–665.
- Young P, Olsen L (2010) *The Healthcare Imperative: Lowering Costs and Improving Outcomes: Workshop Series Summary* (National Academies Press, Washington, DC).
- Zhang J, Denton B, Balasubramanian H, Shah N, Inman B (2012) Optimization of prostate biopsy referral decisions. *Manufacturing Service Oper. Management* 14(4):529–547.