

Dynamic Monitoring and Control of Irreversible Chronic Diseases with Application to Glaucoma

Pooyan Kazemian

Medical Practice Evaluation Center, Division of General Internal Medicine, Massachusetts General Hospital, 55 Fruit St, Boston, Massachusetts 02114, USA

Department of Medicine, Harvard Medical School, 25 Shattuck St, Boston, Massachusetts 02115, USA,
pooyan.kazemian@mgh.harvard.edu

Jonathan E. Helm

Operations and Decision Technologies, Kelley School of Business, Indiana University, 1309 E 10th St, Bloomington, Indiana 47405, USA,
helmj@indiana.edu

Mariel S. Lavieri

Department of Industrial and Operations Engineering, University of Michigan, 1205 Beal Ave, Ann Arbor, Michigan 48109, USA,
lavieri@umich.edu

Joshua D. Stein

Department of Ophthalmology and Visual Sciences, University of Michigan Medical School, 1000 Wall St, Ann Arbor, Michigan 48105, USA, jdstein@med.umich.edu

Mark P. Van Oyen*

Department of Industrial and Operations Engineering, University of Michigan, 1205 Beal Ave, Ann Arbor, Michigan 48109, USA,
vanoyen@umich.edu

To manage chronic disease patients effectively, clinicians must know (i) how to monitor each patient (i.e., when to schedule the next visit and which tests to take), and (ii) how to control the disease (i.e., what levels of controllable risk factors will sufficiently slow progression). Our research addresses these questions simultaneously and provides the optimal solution to a novel linear quadratic Gaussian state space model. For the objective of minimizing the relative change in state over time (i.e., disease progression), which is necessary for managing irreversible chronic diseases while also considering the cost of tests and treatment, we show that the classical two-way separation of estimation and control holds. This makes a previously intractable problem solvable by decomposition into two separate, tractable problems while maintaining optimality. The resulting optimization is applied to the management of glaucoma. Based on data from two large randomized clinical trials, we validate our model and demonstrate how our decision support tool can provide actionable insights to the clinician caring for a patient with glaucoma. This methodology can be applied to a broad range of irreversible chronic diseases to devise patient-specific monitoring and treatment plans optimally.

Key words: chronic disease monitoring and treatment; personalized care; medical decision making; glaucoma; Kalman filter
History: Received: February 2016; Accepted: October 2018 by Sergei Savin, after 3 revisions.

1. Introduction

Chronic diseases are the leading cause of both death and disability and affect almost one out of every two adults in the United States (Ward et al. 2014). To manage chronic diseases, patients are tested quantitatively at prescribed intervals using a selected set of testing modalities to assess disease progression and decide whether a change in treatment is warranted. In this context, proper testing and treatment guidance is critical to both cost containment and patient outcomes in the management

of chronic diseases. In this study, we develop a modeling framework for dynamic management of irreversible chronic diseases that enables us to (i) specify the optimal timing of each office visit and the appropriate suite of tests (i.e., the selection of testing modalities) to perform at that visit considering the costs and value of each test and the uncertainty about the patient's disease progression (disease monitoring), and (ii) identify optimal target levels for controllable disease risk factors to slow the rate of disease progression without over-treating or under-treating the patient (treatment control).

To do so, we introduce and solve a new type of objective function for linear quadratic Gaussian (LQG) systems that minimizes the relative change in state (i.e., disease progression) rather than the traditional objective of minimizing the cost of being in each state, while still accounting for the test and treatment costs as in the traditional objective. We extend LQG theory by proving that the classical two-way separation of optimal state estimation and control applies to this new objective. This separation ensures computational tractability for the simultaneous optimization of disease monitoring and treatment control. This innovative modeling of dynamic disease monitoring and treatment control is developed to be applicable to many irreversible chronic diseases. As a proof of concept, we demonstrate the capabilities of this methodology by applying it to glaucoma, a chronic disease causing progressive blindness.

1.1. Scope of the Research

It is important to distinguish the disease monitoring problem from screening for a disease. The goal of disease screening is to determine whether or not a patient has a particular disease. A screening test is taken when the patient is considered to be at some risk of developing a condition but exhibits no symptoms of the illness. For the disease monitoring problem, however, the patient is already known to have the disease, and the goal is to detect the presence of disease progression quickly and identify whether/how to adjust the treatment plan to slow/avert further disease progression.

For the treatment control portion of the problem, the goal is to determine the time-dependent intensity of treatment over a treatment cycle based on dynamically updating information on patient disease state from the monitoring portion.

We emphasize that our model does not suggest a specific intervention. Rather, it provides patient-specific target levels for controllable/modifiable disease risk factors that help guide the doctor in selecting an appropriate treatment plan for the patient. Though one might try to model how each intervention affects the disease progression dynamics, we feel it is best to leave it to the clinician to employ his/her experience and expertise to decide what therapeutic interventions are most likely able to achieve the target levels suggested by our model.

1.2. Main Contributions

- *Theory:* (1) To the best of our knowledge, this is the first research paper to employ measurement adaptive systems theory to the monitoring and control of chronic diseases (or even to any health-care operations research problem), and this new application requires an extension of the underlying theory. We extend the LQG state space

modeling literature by introducing a new objective that minimizes the *relative change* in system state over time (i.e., the difference in estimated state elements between the current period and the previous period), rather than minimizing the cost of *current* state, while still considering the cost of tests and treatment. In prior applications of LQG modeling to other engineering problems, the goal of the controller has been to keep the system state on a *static* desired trajectory using costly control actions by minimizing the deviation of the current system position from the desired trajectory over time. However, in irreversible diseases such as glaucoma, once the disease has progressed, it is biologically impossible to reverse the damage. In this context, the desired trajectory is to maintain the “current disease state position” (i.e., stop the disease from worsening), which the model updates *dynamically* as the disease progresses over time. This necessitates a new structure for the objective function (Equation 3) not yet studied in LQG literature.

(2) For LQG systems theory, the two-way separation of optimal state estimation and control (known as the *separation principle*) has been a critical foundation upon which to optimize estimation and control of probabilistic systems tractably and simultaneously (see Witsenhausen 1971). Our main theoretical results show that the two-way separation of optimal estimation and control extends to this new objective of relative system state change, which involves two correlated state variables from the current and previous time periods. The treatment control can be optimized in closed form as a linear function of the best estimate of the patient’s current disease state (i.e., filtered state mean given by the Kalman filter) via completion of squares. Furthermore, we show that the monitoring problem can be reduced to a continuous-state discrete-action Markov decision process (MDP) model with filtered and smoothed covariance matrices of the state serving as the information state and the Kalman filter and smoother equations acting as the system dynamics. The MDP can be solved via dynamic programming to find the optimal monitoring schedule specific to each individual patient.

(3) A Kalman filter and smoother are built into our modeling framework to extract noise from the raw measurements and to estimate the disease state optimally in each time period based on imperfect/noisy measurements. This is key to accurately identifying genuine disease progression from testing artifacts. The Kalman smoother is a new feature in our model (compared with the

traditional LQG models), and it is essential because of the new objective function we employ. State smoothing means using information gained at time t to update the prior estimate made at $t-1$ of the value of the state at $t-1$. Filtering refers to estimating the current disease state based on new test results.

- *Application:* (1) We develop an integrated, feedback-driven stochastic control model to provide the jointly optimal solution to both the disease monitoring and treatment control problems. Note that the monitoring regime affects the disease control problem. As new tests are performed, the information gained can affect how the doctor controls/slow the progression of the disease. Therefore, it is critical to model and solve the disease monitoring and control problems together to capture the interaction between them. (2) The model explicitly determines which suite of tests to take at each time period. Some tests are significantly easier and cheaper to perform than others. Different tests may provide more or less information about the patient's disease state. Therefore, it is important to be able to differentiate which tests to perform at each time point in terms of improved monitoring and cost containment. (3) We develop a data-driven decision support tool that provides a *menu of options* to the doctor based on how aggressively he/she wants to monitor and control the patient. The doctor can select an appropriate aggressiveness option depending on the patient's life expectancy, severity of disease, and other personal and clinical factors. For each aggressiveness option, the model incorporates new and past test results as well as clinically-believed and data-verified disease dynamics to predict and graph the future disease trajectory and recommend a patient-specific monitoring regime and target level for controllable disease risk factors.
- *Data:* (1) We parametrize and validate our model using data from two landmark randomized clinical trials of patients with glaucoma. Our numerical results indicate that the model achieves low errors in predicting the future disease trajectory. (2) The results suggest the potential to improve both patient outcomes and system cost when applied to patients from the clinical trial already receiving a high level of care. This potential is likely greater for patients being treated by non-glaucoma specialists.

2. Literature Review

Papers relevant to this research are classified into three categories: (i) theoretical papers on

measurement adaptive systems and sensor scheduling, (ii) medical decision making papers on disease screening, diagnosis, and monitoring, and (iii) optimization models on treatment planning and disease control. In this section, we highlight some prominent papers in each category and briefly describe how our research methodologies and objectives are different.

Measurement adaptive systems and sensor scheduling: The closest paper to our work in terms of theory is Meier et al. (1967). This study lays the foundations for measurement adaptive systems in which the controller seeks to keep the system state on a *static* desired trajectory and obtain information about the system state simultaneously with minimum total cost over a finite horizon. They show that, in the special case of discrete-time systems, linear system dynamics, quadratic cost of the *current state*, and Gaussian random noise processes, the problem of finding the optimal measurement policy reduces to the solution of a nonlinear, deterministic control problem. Baron and Kleinman (1969) extend their work to continuous-time measurements and investigate the optimal measurement duration for a human operator. Bansal and Başar (1989) provide an extension of this framework to the infinite-horizon setting with discounted costs. Our work differs in that it deals with a *dynamic* desired trajectory, minimizing the relative change in state in each time period (i.e., disease progression), which is essential for managing irreversible chronic diseases (as discussed in subsection 1.2). For example, experiments using the model provided by Meier et al. (1967) led to results considered clinically incorrect/unbelievable in the experience of our clinical co-author, a glaucoma specialist. We discuss this more in section 3.3.

There is also extensive literature on sensor scheduling problems, in which a set of sensors is used to estimate a stochastic process, but because of cost or design constraints, only one or a subset of them takes measurements at each time point. Athans (1972) considers the problem in which the controller has to select one measurement provided by one sensor out of many available sensors (with different measurement costs) at each time step, such that a weighted combination of prediction accuracy and accumulated observation cost is minimized. Examples of other work in this area include Gupta et al. (2006), Mehra (1976), and Vitus et al. (2012). However, these papers differ from ours in that they do not consider the trade-off between exploration vs. exploitation.

Disease screening, diagnosis, and monitoring: While there is extensive literature on disease screening and diagnosis problems, there is relatively little work on the disease monitoring problem that we defined in section 1.1. Helm et al. (2015) and Schell et al. (2014) provide a heuristic approach for finding the time in between tests based on patient's probability of

progression. Note that their work focuses only on the timing of the next test and does not consider how that test would impact future tests performed nor the type of test performed at each visit. (They assume *all* of the tests are performed at each office visit.) These works also do not incorporate optimization of controllable disease risk factors (i.e., treatment planning), which is a key factor to consider when managing patients with chronic diseases. Ayer et al. (2012) provide a partially observable Markov decision process (POMDP) approach to personalize mammography screening decisions; their work is based on the prior screening history and personal risk characteristics of women. Chhatwal et al. (2010) develop a finite-horizon discrete-time Markov decision process (MDP) model to help radiologists determine the best time for a biopsy based on the initial mammography findings and patient demographics. Their goal is to maximize a patient's total expected quality-adjusted life years. The works of Yang et al. (2013), Mangasarian et al. (1995), Saaty and Vargas (1998), Zhang et al. (2012), Lee et al. (2019), and Erenay et al. (2014) are other examples of disease screening models. These works differ from ours in that they focus on the screening problem where the goal is to detect the presence of a particular disease with minimum delay. They do not provide any insights into how to monitor the patient if the presence of the disease is confirmed and progression trajectory can be monitored over time, nor do they consider treatment planning. Kreke (2007) develops a POMDP model to study the question of when to test for cytokine levels in patients with sepsis using available, costly, and inaccurate tests. Suen et al. (2017) develop a POMDP framework to determine when and how often sputum smear test information should be collected from patients on first-line treatment of tuberculosis and use this information to find the optimal time to administer drug sensitivity testing. These works consider the problem of timing of each test with testing noise and system noise but do not consider simultaneous disease monitoring and treatment decisions.

Treatment planning and disease control: There has been a variety of work considering when to start treatment of a patient when the presence of disease is confirmed (also known as surveillance problems). Lavieri et al. (2012) develop a Kalman filter-based approach to help clinicians decide when to start radiation therapy in patients with prostate cancer based on predictions of the time when the patient's prostate specific antigen (PSA) level reaches its lowest point. Notice that, different from our approach, its actions are discrete (start/do not start radiation therapy) rather than establishing the level of therapy to be given. Furthermore, this work did not consider when the patient should be monitored, assuming the patient would be seen at the

clinic once a month. Shechter et al. (2008) employ Markov decision processes (MDP) to optimize the time to initiate HIV treatment to maximize a patient's quality-adjusted life years. Mason et al. (2014) and Schell et al. (2019) present MDP models to determine the optimal timing of blood pressure and cholesterol medications. All of these papers assume a measurement of the patient's health is taken periodically. Our work differs in that it solves the joint problem of optimal timing of each test and optimal treatment control.

Moreover, in most of the previous research mentioned, the patient's disease dynamics are assumed to be known or are estimated using population-based models. In our model, the population data is integrated with individual patient measurements gathered from sequential testing so that the predictions and decisions made are unique to each patient. Capturing the complex patient disease dynamics requires incorporating several health indices into the state vector. We employ a continuous state space that easily accommodates multivariate states (e.g., nine dimensions in our model for glaucoma) and provide jointly optimal solutions to both disease monitoring and control problems. Employing a continuous state space model is important as many quantitative tests for disease monitoring are continuous. Problems with such a multivariate, continuous state-space often become intractable for MDP-based approaches due to the curse of dimensionality. Discretization of the state space and using approximate dynamic programming (ADP) to mitigate the curse of dimensionality of MDP models is an alternative approach when our modeling framework does not fit. For example, strongly discrete state variables, highly nonlinear disease dynamics, and highly non-Gaussian random noises are features that are difficult for our model to handle. However, it should be noted that discretization of the state space and grouping discrete states together to reduce the model size is shown to reduce the predictive power of the MDP models as a result of lumping error (see Regnier and Shechter 2013). The implication is that if the assumptions in our model are reasonable for a disease, a modeling framework like what we present in this study that can accommodate continuous state space (for a disease with continuous variables) without loss of optimality can benefit from more accurate predictions (specially beyond one period into the future) and can be solved tractably.

3. The Modeling Framework

A continuous state space model is employed at the heart of our modeling framework with two key components: (i) a state transition process to model disease progression dynamics, and (ii) a measurement/testing process to model how the true disease state is

observed. Both processes (Equations 1 and 2) are in the form of first-order (linear) vector difference equations with additive Gaussian white noise (i.e., noise inputs at time t and t' are independent). The objective function seeks to minimize the weighted sum of multiple cost elements with state independent cost parameters/weights. The optimization problem determines both the optimal disease control (how to control modifiable risk factors) and the optimal monitoring (when to take new measurements), both of which are patient-specific and state-dependent. Such a model can be a good fit for many chronic diseases (such as glaucoma) since the measurements are typically on a continuous scale. In the case study of glaucoma (section 5), we will elaborate on how we adjust for some of the model assumptions, for example, we include first and second-order derivatives of key disease state elements in the state vector to capture some degree of nonlinearity in our linear model, and we define various aggressiveness levels that the clinician can choose from, each of which is tied to a set of cost parameters.

3.1. State Transition Process

The recursive state transition equation for our N -stage time horizon is given by

$$\alpha_{t+1} = T_t \alpha_t + G_t \beta_t + \eta_t, \quad t = 1, \dots, N, \quad (1)$$

where α_t is the random vector representing the state of the disease at time t , β_t is the “disease control” variable administered at time t , η_t is the vector of Gaussian white noise that represents unmodeled disease process noise with $E[\eta_t] = 0$ and $\text{Cov}(\eta_t) = Q_t$, T_t is the state transition matrix governing the underlying disease progression dynamics, and G_t is a vector capturing the effect of disease control variable β_t on the next period state, α_{t+1} . β_t is one of the two optimization variables of the model. It determines how the modifiable disease risk factors should be adjusted at time t to slow the progression of disease optimally. Our model seeks to adjust the modifiable risk factors in the first few time periods, after which β_t converges to zero. By summing the total amount of control β_t over this timeframe, we obtain a target level for each modifiable risk factor. We will illustrate this in section 5.8.2. Having such information will help clinicians select the appropriate treatment plan for the patient.

3.2. Measurement/Testing Process

The measurement equation gives the relationship between the true disease state, α_t , and the noisy raw reading/observation, z_t , as follows.

$$z_t = Z_t \alpha_t + \varepsilon_t, \quad t = 1, \dots, N, \quad (2)$$

where z_t is the observation vector (i.e., the result of test(s) performed on the patient), Z_t is the

observation matrix and determines how components of the true state are observed, and ε_t is the multivariate Gaussian white test noise with $E[\varepsilon_t] = 0$ and $\text{Cov}(\varepsilon_t) = H_t^{(\theta_t)}$. θ_t is the “test/measurement control” variable that determines which subset of tests to take in period t ; it is the other control variable the model optimizes. $H_t^{(\theta_t)}$ models the error associated with the tests and is directly affected by the decision on which test(s) to take at time t (which we highlight by adding (θ_t) to the superscript, that is, $H_t^{(\theta_t)}$). It is worth noting that both α_{t+1} and z_t are Gaussian random vectors since they are a linear combination of independent Gaussian random variables. The initial state, α_0 , is Gaussian with $E[\alpha_0] = \hat{\alpha}_0 = \hat{\alpha}_{1|0}$ and $\text{Cov}(\alpha_0) = \hat{\Sigma}_0 = \hat{\Sigma}_{1|0}$. The random variables α_0 , $\{\eta_t\}$, and $\{\varepsilon_t\}$ are mutually independent. Throughout the study, the notation $\hat{X}_{t|t'}$ means the estimated value of random variable X at time t with information up to time t' .

3.3. Objective Function

The novel objective function (performance criterion) we analyze is given by

$$J = E \left\{ \sum_{t=1}^N [(\alpha_t - \alpha_{t-1})' A_t (\alpha_t - \alpha_{t-1}) + \beta_t' B_t \beta_t + l_t(\theta_t)] + (\alpha_{N+1} - \alpha_N)' A_{N+1} (\alpha_{N+1} - \alpha_N) \right\}, \quad (3)$$

in which A_t is the unit cost matrix of further worsening the disease, B_t is the unit cost of administering disease control (i.e., further adjusting the modifiable disease risk factor), and the scalar $l_t(\theta_t)$ is the cost of taking tests/measurements in period t , which depends on the test control variable, θ_t . The objective function consists of four terms: (i) $(\alpha_t - \alpha_{t-1})' A_t (\alpha_t - \alpha_{t-1})$ is the cost of relative change in the system state random variable (i.e., disease progression) between the previous period $t - 1$ and the current period t (whereas the traditional LQG objective minimizes $\alpha_t' A_t \alpha_t$ as explained in subsection 1.2), (ii) $\beta_t' B_t \beta_t$ is the cost of controlling the disease risk factors including side effects and complications of medical or surgical treatments, (iii) $l_t(\theta_t)$ is the cost of taking tests, and (iv) $(\alpha_{N+1} - \alpha_N)' A_{N+1} (\alpha_{N+1} - \alpha_N)$ is the terminal cost of relative state change at the end of the treatment horizon. The quadratic form of the first part of the objective function ensures that a large disease worsening is penalized more aggressively than a small one. Furthermore, achieving a large adjustment in disease risk factors may require more aggressive treatments (e.g., surgery or laser therapy), which are associated with higher monetary costs as well as more side effects and discomfort than a smaller change in risk factors, which can often be achieved by simpler

treatments such as medications. Hence, the cost associated with a big relative change in a patient's disease risk factors is much higher than a small one, so the quadratic form of the second part of the objective function is a good choice for our application.

One might try to define the first part of the objective function as minimizing the deviation of the current state from the patient's baseline state (i.e., minimizing $E \sum [(\alpha_t - \alpha_0)'A_t(\alpha_t - \alpha_0) + \beta_t'B_t\beta_t + l_t(\theta_t)]$). Note that this is not an appropriate objective for irreversible diseases because, according to our experiments, the model will attempt to reverse the disease progression and send the state back to its baseline (which we know is biologically impossible) by making extreme and often infeasible changes to state variables that are not valid clinically. Moreover, such an objective function assigns quadratically increasing cost to additional disease progression measured as the current deviation from its baseline, which can easily overshadow the other cost elements in the objective function (i.e., treatment and monitoring costs). This will result in recommending the most aggressive treatment and monitoring regimes for all patients when the cost associated with deviation of current state from its baseline overshadows the other two costs in the objective function regardless of how bad the current state is. We believe that good clinical practice is always careful to note the latest disease state, with any worsening being of concern. Our objective function achieves this without introducing the problems noted above.

Another alternative approach is to define the state as $\tilde{\alpha}_t = \begin{bmatrix} \alpha_t \\ \alpha_{t-1} \end{bmatrix}$, and minimize $E \sum_t [\tilde{\alpha}_t' \tilde{A}_t \tilde{\alpha}_t + \beta_t' B_t \beta_t + l_t(\theta_t)]$. While it is possible to develop an alternative formulation equivalent to the one we presented above based on this new definition of state, \tilde{A}_t in the objective function of the alternative formulation is a singular (non-invertible) cost matrix (more specifically, $\tilde{A}_t = \begin{bmatrix} A_t & -A_t \\ -A_t & A_t \end{bmatrix}$; calculation not shown).

However, we would need to invert this matrix to find a closed-form solution for optimal disease control. Hence, the alternative formulation makes it harder to derive a closed-form solution for the optimal disease control. Note that, in the formulation proposed in our study, A_t is always invertible because A is a cost matrix and, by definition, it is diagonal with only non-negative terms on the diagonal.

3.4. Kalman Filter and Kalman Smoother

When the state transition and measurement processes are both in the form of first-order difference equations with Gaussian white noises, the optimal state estimation method that minimizes the mean squared error

of the estimate is given by the Kalman filter (Kalman 1960). The Kalman filter obtains the prediction of state mean and covariance at time t with information up to time $t - 1$, $\hat{\alpha}_{t|t-1}$ and $\hat{\Sigma}_{t|t-1}$ respectively, and the current reading, z_t , as inputs to the algorithm and calculates the filtered state (i.e., optimal estimate of the true state) mean and covariance, $\hat{\alpha}_{t|t}$ and $\hat{\Sigma}_{t|t}$ respectively.

The optimal state mean estimate at time t with information up to time t , $\hat{\alpha}_{t|t}$, is given by

$$\hat{\alpha}_{t|t} = \hat{\alpha}_{t|t-1} + K_t \tilde{y}_t, \quad (4)$$

where $\hat{\alpha}_{t|t-1}$ is the predicted state mean at time t given information up to time $t - 1$ and \tilde{y}_t is the measurement residual (error) given by

$$\hat{\alpha}_{t|t-1} = T_{t-1} \hat{\alpha}_{t-1|t-1} + G_{t-1} \beta_{t-1}, \quad (5)$$

$$\tilde{y}_t = z_t - Z_t \hat{\alpha}_{t|t-1}, \quad (6)$$

and K_t is the Kalman gain given by

$$K_t = \hat{\Sigma}_{t|t-1} Z_t' S_t^{-1}, \quad (7)$$

in which S_t is the predicted covariance around the measurement given by $Z_t \hat{\Sigma}_{t|t-1} Z_t' + H_t^{(\theta_t)}$. The predicted state covariance at time t given the information up to time $t - 1$, $\hat{\Sigma}_{t|t-1}$, and the most recent state covariance estimate at time t with information up to time t , $\hat{\Sigma}_{t|t}$, satisfy

$$\hat{\Sigma}_{t|t} = \hat{\Sigma}_{t|t-1} - \hat{\Sigma}_{t|t-1} Z_t' \left(Z_t \hat{\Sigma}_{t|t-1} Z_t' + H_t^{(\theta_t)} \right)^{-1} \quad (8)$$

$$Z_t \hat{\Sigma}_{t|t-1} = (I - K_t Z_t) \hat{\Sigma}_{t|t-1},$$

$$\hat{\Sigma}_{t|t-1} = T_{t-1} \hat{\Sigma}_{t-1|t-1} T_{t-1}' + Q_{t-1}. \quad (9)$$

The initial state mean and covariance, $\hat{\alpha}_{1|0}$ and $\hat{\Sigma}_{1|0}$ respectively, are calculated based on population data from clinical trials. So, $\tilde{y}_1 = z_1 - Z_1 \hat{\alpha}_{1|0}$ and $S_1 = Z_1 \hat{\Sigma}_{1|0} Z_1' + H_1^{(\theta_1)}$. For more discussion on Kalman filter, see Bertsekas (1995) and Harvey (1990).

Because of the special form of the objective function that minimizes relative state change from time $t - 1$ to time t (i.e., disease progression), we need to refine the estimation of previous state mean and covariance ($\hat{\alpha}_{t-1|t}$ and $\hat{\Sigma}_{t-1|t}$, respectively) after a new measurement is taken at time t . This is called state smoothing and can be done via a fixed-interval Kalman smoother as follows.

$$\hat{\alpha}_{t-1|t} = \hat{\alpha}_{t-1|t-1} + \hat{\Sigma}_{t-1}^* (\hat{\alpha}_{t|t} - \hat{\alpha}_{t|t-1}), \quad (10)$$

$$\hat{\Sigma}_{t-1|t} = \hat{\Sigma}_{t-1|t-1} + \hat{\Sigma}_{t-1}^* (\hat{\Sigma}_{t|t} - \hat{\Sigma}_{t|t-1}) \hat{\Sigma}_{t-1}^{*'}, \quad (11)$$

in which $\hat{\Sigma}_{t-1}^* = \hat{\Sigma}_{t-1|t-1} T_{t-1}' \hat{\Sigma}_{t|t-1}^{-1}$. A derivation of the fixed-interval Kalman smoothing can be found in Ansley and Kohn (1982).

The control system block diagram is depicted in Figure 1. The dashed arrows indicate that the information is carried over from the current period, t , to the next period, $t + 1$. The values in parentheses are not observable. Suppose the patient is in disease state α_t when visiting the doctor’s office. Based on the optimal test control action θ_t^* (already determined in the previous time period), all or a subset of tests are performed on the patient. The noisy observation/reading, z_t , is then sent to the Kalman Filter. Based on the predicted and observed states, the Kalman Filter algorithm calculates the best estimate of the mean and covariance of the patient’s disease state in period t , $\hat{\alpha}_{t|t}$ and $\hat{\Sigma}_{t|t}$ respectively, and sends the filtered values to both the Kalman Smoother and the controller (i.e., the decision support system itself for this analysis). The Kalman Smoother will then modify the best estimates of the state mean and covariance in period $t - 1$, $\hat{\alpha}_{t-1|t}$ and $\hat{\Sigma}_{t-1|t}$ respectively, and send the smoothed values to the controller. Notice that this is a key departure from the traditional methodology. The controller receives both the filtered and smoothed values of the patient’s disease state mean and covariance (the information state for the optimization component of the model) and outputs the optimal treatment and test control actions, β_t^* and θ_{t+1}^* . Finally, the prediction of the state mean and covariance in period $t + 1$, $\hat{\alpha}_{t+1|t}$ and $\hat{\Sigma}_{t+1|t}$, is sent to the Kalman filter and smoother to be used in the following time period.

In section 4, we focus on the controller (decision support tool) and show how the optimal disease and test control actions (β_t^* and θ_{t+1}^* , respectively) can be

calculated given the information state $\wp_t = (\hat{\alpha}_{t|t}, \hat{\Sigma}_{t|t}, \hat{\alpha}_{t-1|t}, \hat{\Sigma}_{t-1|t})$. For convenience, Table 1 provides a list of important notations we used in our modeling.

3.5. Separation of Estimation and Control

One-way separation of estimation and control: A control law is a function selected by the controller from the set of all admissible functions based on all observations available at the time of the decision. This function generates a control action to be applied to the system. The problem is to make an optimal selection of such functions for all time steps that achieves the minimum expected cost (defined by the objective function) for the control horizon of the problem. For the general stochastic control problem with imperfect observations, given all the observations and previous control actions, state estimation from a noisy measurement is always independent of the control law. This is because the conditional density of the state given all the observations and the control actions is independent of the control law. This result is called the separation principle in control theory (see Witsenhausen 1971). The only underlying assumption for the separation principle to hold is to have one controller (i.e., centralized information) with perfect recall (i.e., the information on the previous observations and control actions do not get forgotten).

In general, the control law depends on the estimate of the system state; however, the estimate at time t is independent of all control laws given all observations up to time t and all the control actions up to time $t - 1$. This is also known as the one-way separation

Figure 1 Control System Block Diagram

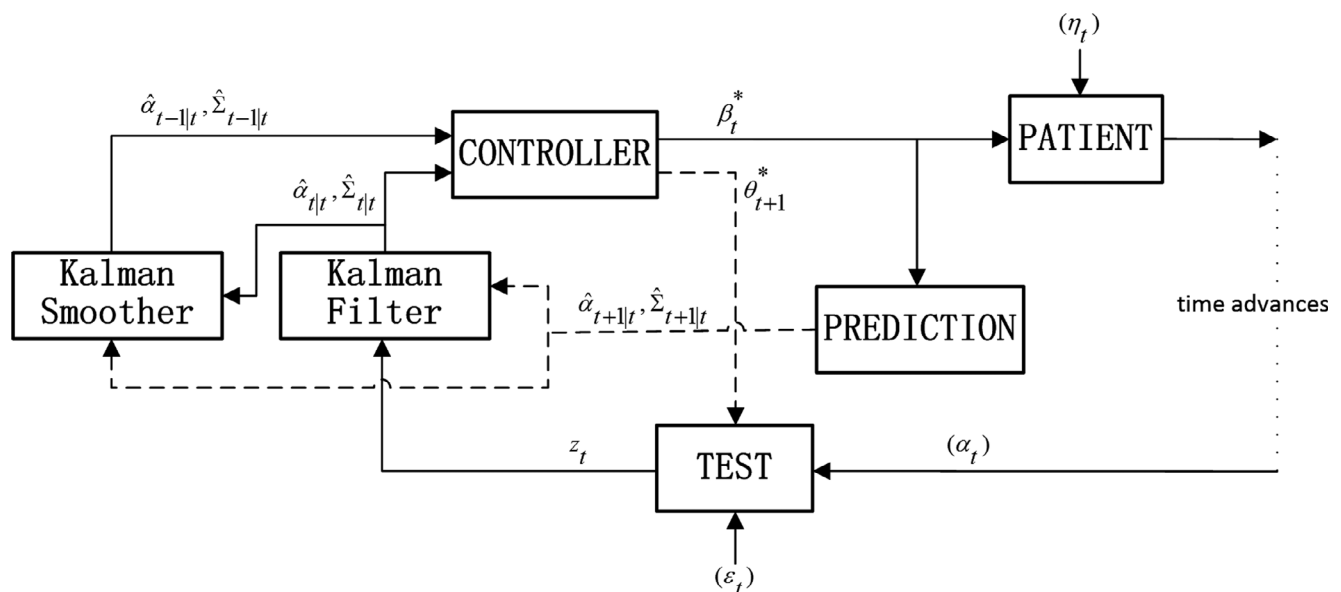


Table 1 List of Important Model Notation

Notation	General type [size in case study]	Description
α_t	Vector [9 × 1]	Disease state at time t
T_t	Square matrix [9 × 9]	State transition matrix governing the underlying disease progression dynamics at time t
β_t	Vector [1 × 1]*	Disease control action—how much to reduce IOP at time t
G_t	Square matrix [9 × 1]*	Coefficient capturing the effect of disease control action at time t on the next period state
η_t	Vector [9 × 1]	Disease process noise at time t
Q_t	Square matrix [9 × 9]	Covariance of the disease process noise at time t
Z_t	Vector [9 × 1]	Observation/measurement vector at time t
Z_t	Square matrix [9 × 9]	Observation/measurement matrix that determines how components of the true state are observed at time t
ε_t	Vector [9 × 1]	Measurement process noise at time t
θ_t	Discrete choices [3 choices]	Test/measurement control action—which subset of tests to take at time t , e.g., no test, only IOP test, or both IOP and VF tests
$H_t^{(\theta_t)}$	Square matrix [9 × 9]	Covariance of the measurement process noise at time t that depends on which tests are taken at time t
α_0	Vector [9 × 1]	Initial state
$\hat{\alpha}_0$ or $\hat{\alpha}_{1 0}$	Vector [9 × 1]	Estimate of the initial state mean
$\hat{\Sigma}_0$ or $\hat{\Sigma}_{1 0}$	Square matrix [9 × 9]	Estimate of the initial state covariance
$\hat{\alpha}_{t t'}$	Vector [9 × 1]	Estimate of state mean at time t with information up to time t' **
$\hat{\Sigma}_{t t'}$	Square matrix [9 × 9]	Estimate of state covariance at time t with information up to time t' **
A_t	Square matrix [9 × 9]	Cost of further worsening of the disease at time t —e.g., cost of vision loss due to glaucoma
B_t	Square matrix [1 × 1]*	Cost of further adjusting the modifiable disease risk factor at time t —e.g., cost of further lowering the IOP
$l_t(\theta_t)$	Set of scalars [3 scalars]	Cost of taking tests/measurements at time t that depends on which tests are taken at time t
J	Scalar	Expected total cost during the entire control horizon
K_t	Square matrix [9 × 9]	Kalman gain at time t
\tilde{y}_t	Vector [9 × 1]	Measurement residual at time t —the difference between observed and predicted measurement at time t
S_t	Square matrix [9 × 9]	Predicted covariance around the measurement at time t
$\hat{\Sigma}_t^*$	Square matrix [9 × 9]	Smoothing adjustment factor at time t
φ_t	Collection of two vectors [9 × 1] and two matrices [9 × 9]	Information state at time t , which includes filtered and smoothed disease state mean and covariance
$V_t(\varphi_t)$	Scalar	Value function at time t —the minimum expected cost from period t to N , the end of the control horizon, given the information state φ_t
U_t	Square matrix [9 × 9]	Control law at time t , given by Equation (31)

Notes: Auxiliary notation used in section 4 to prove Theorems 1–3 are not shown in this table; they are defined upon introduction in section 4. *Note that since glaucoma has only one controllable risk factor (i.e., IOP), in our case study of glaucoma, β_t and B_t become scalars and G_t becomes a vector. **Smoothing, filtering, and prediction refer to the case in which $t' > t$, $t' = t$, and $t' < t$, respectively.

of estimation and control. Since our LQG model is a special case of the general stochastic control problem with centralized information and perfect recall, the one-way separation principle holds. As seen in section 3.4, the optimal state estimation at time t , described by $\hat{\alpha}_{t|t}$ and $\hat{\Sigma}_{t|t}$, is given by the Kalman filter (Equations 4–9) and is independent of the control law given all the previous observations and control actions.

Two-way separation of estimation and control: For LQG stochastic systems in which (i) the transition and measurement equations are linear in state and control action, (ii) the objective function penalizes the quadratic cost of *current* state, and (iii) the state and measurement noises are Gaussian, it has been shown that the control law is also independent of the state estimation (Meier et al. 1967). Therefore, for this traditional form of LQG models, we have two-way separation of the estimation and

the control; namely, the estimation is independent of the control law and the control law is independent of the estimation. In section 4, we show for the new objective of minimizing the *relative change* in state, which involves two correlated state variables of current and previous time periods and requires smoothing in addition to filtering and prediction, that the optimal control law is still independent of the state estimate. Thus, in this new and more complex environment, the two-way separation still holds. Furthermore, the optimal control action is linear in the state estimate. This is extremely desirable for application because the control law is data independent and can be calculated offline (which greatly reduces the computational burden). The two-way separation of estimation and control for this special case of LQG models is a fundamental finding, which is critical to solution tractability.

4. Derivation of Optimal Disease and Test Controls

In this section, we derive the optimal disease and test control actions given the information state at time t , \wp_t , which is defined as the filtered state mean and covariance at time t and the smoothed state mean and covariance at time $t - 1$ with information up to time t , that is, $\wp_t = (\hat{\alpha}_{t|t}, \hat{\Sigma}_{t|t}, \hat{\alpha}_{t-1|t}, \hat{\Sigma}_{t-1|t})$.

In terms of \wp_t , a dynamic programming algorithm can be derived to find the optimum disease and test controls. The value function, $V_t(\wp_t)$, can be found recursively as follows.

$$V_t(\wp_t) = \min_{\beta_t, \theta_{t+1}} \left\{ L_t(\wp_t, \beta_t, \theta_{t+1}) + E_{z_{t+1}} [V_{t+1}(\wp_{t+1})] \right\},$$

$$t = 1, \dots, N - 1, \quad (12)$$

where $V_t(\wp_t)$ is the minimum expected cost from period t to N , the end of the control horizon, given the information state \wp_t , and $L_t(\wp_t, \beta_t, \theta_{t+1})$. Here $L_t(\wp_t, \beta_t, \theta_{t+1})$ is the expected instantaneous (one-stage) cost incurred in period t given that the information state is \wp_t and the control actions β_t and θ_{t+1} are chosen, and it is computed as

$$L_t(\wp_t, \beta_t, \theta_{t+1}) = E[(\alpha_t - \alpha_{t-1})' A_t (\alpha_t - \alpha_{t-1}) | \wp_t] + \beta_t' B_t \beta_t + l_{t+1}(\theta_{t+1}). \quad (13)$$

The boundary condition is given by

$$V_N(\wp_N) = \min_{\beta_N} \{ \tilde{L}_N(\wp_N, \beta_N) \}, \quad (14)$$

where $\tilde{L}_N(\wp_N, \beta_N)$ is the expected cost incurred in the final period N if the information state is \wp_N and the disease control action β_N is chosen. The minimum cost during the entire control horizon can, therefore, be obtained by

$$J^* = l_1(\theta_1) + V_1(\wp_1), \quad (15)$$

in which $l_1(\theta_1)$ is the cost of initial tests during the patient's first office visit and $V_1(\wp_1)$ is the minimum cost to go from period 1 to the end of control horizon obtained recursively via Equations 12 and 14. We assume all diagnostic tests are taken during the first visit, so a baseline of all readings is established.

In the remainder of this section we use an induction argument to prove the following theorems.

THEOREM 1. *For arbitrary time t ($t = 1, \dots, N$), the control law is independent of the state estimate (i.e., we have two-way separation of optimal estimation and control). Moreover, the optimal disease control, β_t^* , is linear in the filtered state mean, $\hat{\alpha}_{t|t}$.*

THEOREM 2. *At an arbitrary time t ($t = 1, \dots, N$), the optimal monitoring schedule, θ_{t+1}^* , θ_{t+2}^* , \dots , can be found by solving a continuous-state discrete-action MDP model with filtered and smoothed covariance matrices of the state serving as the information state and the Kalman filter and smoother equations acting as the system dynamics.*

THEOREM 3. *For arbitrary time t ($t = 1, \dots, N$), the value function with information up to time t has the following form.*

$$V_t(\wp_t) = (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) + \text{tr}[A_t \hat{\Sigma}_{t|t}] + \hat{\alpha}_{t|t}' P_t \hat{\alpha}_{t|t} + \text{tr}[P_t \hat{\Sigma}_{t|t}] + V_t^0(\wp_t^0) + b_t, \quad (16)$$

in which $V_t^0(\wp_t^0)$ represents the recursive terms that only depend on measurement control actions, that is, when to take tests and which test(s) to take. They do not depend on observations or on disease control actions. Therefore, the measurement control problem can be solved separately from the treatment control problem. \wp_t^0 represents those elements of information state that are only affected by measurement control actions (i.e., $\wp_t^0 = (\hat{\Sigma}_{t|t}, \hat{\Sigma}_{t-1|t})$), and b_t is a constant. $V_t^0(\wp_t^0)$ and b_t will be obtained later in the proof.

PROOF BY INDUCTION. In Appendix A we prove that the value function in the final period is given by

$$V_N(\wp_N) = (\hat{\alpha}_{N|N} - \hat{\alpha}_{N-1|N})' A_N (\hat{\alpha}_{N|N} - \hat{\alpha}_{N-1|N}) + \text{tr}[A_N \hat{\Sigma}_{N|N}] + \hat{\alpha}_{N|N}' P_N \hat{\alpha}_{N|N} + \text{tr}[P_N \hat{\Sigma}_{N|N}] + \text{tr}[A_N (\hat{\Sigma}_{N-1|N} - T_{N-1} \hat{\Sigma}_{N-1|N} - \hat{\Sigma}_{N-1|N} T_{N-1}')] + \text{tr}[\tilde{P}_{N+1} \hat{\Sigma}_{N|N}] + \text{tr}[A_{N+1} Q_N], \quad (17)$$

where tr represents the trace of the matrix. By comparing Equations 16 and 17, we obtain our basis for induction:

$$V_N^0(\wp_N^0) = \text{tr}[A_N (\hat{\Sigma}_{N-1|N} - T_{N-1} \hat{\Sigma}_{N-1|N} - \hat{\Sigma}_{N-1|N} T_{N-1}')] + \text{tr}[\tilde{P}_{N+1} \hat{\Sigma}_{N|N}], \quad (18)$$

$$b_N = \text{tr}[A_{N+1} Q_N]. \quad (19)$$

Assume the induction hypothesis that

$$V_{t+1}(\wp_{t+1}) = (\hat{\alpha}_{t+1|t+1} - \hat{\alpha}_{t|t+1})' A_{t+1} (\hat{\alpha}_{t+1|t+1} - \hat{\alpha}_{t|t+1}) + \text{tr}[A_{t+1} \hat{\Sigma}_{t+1|t+1}] + \hat{\alpha}_{t+1|t+1}' P_{t+1} \hat{\alpha}_{t+1|t+1} + \text{tr}[P_{t+1} \hat{\Sigma}_{t+1|t+1}] + V_{t+1}^0(\wp_{t+1}^0) + b_{t+1}. \quad (20)$$

We show that $V_t(\wp_t)$ follows the form given in Equation 16 to complete the induction step. From Equation 12 we know the general form of value function is

$$V_t(\wp_t) = \min_{\beta_t, \theta_{t+1}} \left\{ L_t(\wp_t, \beta_t, \theta_{t+1}) + E_{z_{t+1}} [V_{t+1}(\wp_{t+1})] \right\}, \quad (21)$$

in which the information state at time $t + 1$, \wp_{t+1} , is a function of \wp_t , β_t , θ_{t+1} , and z_{t+1} . The expected instantaneous cost in period t , $L_t(\wp_t, \beta_t, \theta_{t+1})$, is given in Equation 13. Application of Lemma 3 to the expectation in Equation 13 results in

$$\begin{aligned} L_t(\wp_t, \beta_t, \theta_{t+1}) &= (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \\ &\quad + \text{tr} [A_t (\hat{\Sigma}_{t|t} + \hat{\Sigma}_{t-1|t} \\ &\quad - T_{t-1} \hat{\Sigma}_{t-1|t} - \hat{\Sigma}_{t-1|t} T_{t-1}')] \\ &\quad + \beta_t' B_t \beta_t + l_{t+1}(\theta_{t+1}). \end{aligned} \quad (22)$$

Replacing $L_t(\wp_t, \beta_t, \theta_{t+1})$ and $V_{t+1}(\wp_{t+1})$ in Equation 21 by the values given by Equations 22 and 20 respectively, yields

$$\begin{aligned} V_t(\wp_t) &= \min_{\beta_t, \theta_{t+1}} \left\{ (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \right. \\ &\quad + \text{tr} [A_t (\hat{\Sigma}_{t|t} + \hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} \\ &\quad - \hat{\Sigma}_{t-1|t} T_{t-1}')] + \beta_t' B_t \beta_t + l_{t+1}(\theta_{t+1}) \\ &\quad + E_{z_{t+1}} \left[(\hat{\alpha}_{t+1|t+1} - \hat{\alpha}_{t|t+1})' A_{t+1} (\hat{\alpha}_{t+1|t+1} - \hat{\alpha}_{t|t+1}) \right. \\ &\quad + \hat{\alpha}_{t+1|t+1}' P_{t+1} \hat{\alpha}_{t+1|t+1} \left. \right] + \text{tr} [A_{t+1} \hat{\Sigma}_{t+1|t+1}] \\ &\quad \left. + \text{tr} [P_{t+1} \hat{\Sigma}_{t+1|t+1}] + V_{t+1}^\theta(\wp_{t+1}^\theta) + b_{t+1} \right\}, \end{aligned} \quad (23)$$

Replacing $E_{z_{t+1}} [(\hat{\alpha}_{t+1|t+1} - \hat{\alpha}_{t|t+1})' A_{t+1} (\hat{\alpha}_{t+1|t+1} - \hat{\alpha}_{t|t+1})]$ and $E_{z_{t+1}} [\hat{\alpha}_{t+1|t+1}' P_{t+1} \hat{\alpha}_{t+1|t+1}]$ using Lemmas 7 and 8 in Appendix B, respectively, yields

$$\begin{aligned} V_t(\wp_t) &= \min_{\beta_t, \theta_{t+1}} \left\{ (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \right. \\ &\quad + \text{tr} [A_t (\hat{\Sigma}_{t|t} + \hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} \\ &\quad - \hat{\Sigma}_{t-1|t} T_{t-1}')] + \beta_t' B_t \beta_t + l_{t+1}(\theta_{t+1}) \\ &\quad + ((T_t - I) \hat{\alpha}_{t|t} + G_t \beta_t)' A_{t+1} ((T_t - I) \hat{\alpha}_{t|t} + G_t \beta_t) \\ &\quad + \text{tr} [A_{t+1} (T_t \hat{\Sigma}_{t|t} T_t' + Q_t - \hat{\Sigma}_{t+1|t+1} \\ &\quad - \hat{\Sigma}_{t+1|t+1} \hat{\Sigma}_t^* - \hat{\Sigma}_t^* \hat{\Sigma}_{t+1|t+1} + \hat{\Sigma}_{t+1|t+1})] \\ &\quad + (T_t \hat{\alpha}_{t|t} + G_t \beta_t)' P_{t+1} (T_t \hat{\alpha}_{t|t} + G_t \beta_t) \\ &\quad + \text{tr} [P_{t+1} (T_t \hat{\Sigma}_{t|t} T_t' + Q_t - \hat{\Sigma}_{t+1|t+1})] \\ &\quad + \text{tr} [A_{t+1} \hat{\Sigma}_{t+1|t+1}] + \text{tr} [P_{t+1} \hat{\Sigma}_{t+1|t+1}] \\ &\quad \left. + V_{t+1}^\theta(\wp_{t+1}^\theta) + b_{t+1} \right\}. \end{aligned} \quad (24)$$

Canceling terms results in

$$\begin{aligned} V_t(\wp_t) &= \min_{\beta_t, \theta_{t+1}} \left\{ (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \right. \\ &\quad + \text{tr} [A_t (\hat{\Sigma}_{t|t} + \hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} \\ &\quad - \hat{\Sigma}_{t-1|t} T_{t-1}')] + \beta_t' B_t \beta_t + l_{t+1}(\theta_{t+1}) \\ &\quad + ((T_t - I) \hat{\alpha}_{t|t} + G_t \beta_t)' A_{t+1} ((T_t - I) \hat{\alpha}_{t|t} + G_t \beta_t) \\ &\quad + (T_t \hat{\alpha}_{t|t} + G_t \beta_t)' P_{t+1} (T_t \hat{\alpha}_{t|t} + G_t \beta_t) \\ &\quad + \text{tr} [(A_{t+1} + P_{t+1}) (T_t \hat{\Sigma}_{t|t} T_t')] + \text{tr} [(A_{t+1} + P_{t+1}) Q_t] \\ &\quad + \text{tr} [A_{t+1} (\hat{\Sigma}_{t+1|t+1} - \hat{\Sigma}_t^* \hat{\Sigma}_{t+1|t+1} - \hat{\Sigma}_{t+1|t+1} \hat{\Sigma}_t^*)] \\ &\quad \left. + V_{t+1}^\theta(\wp_{t+1}^\theta) + b_{t+1} \right\}. \end{aligned} \quad (25)$$

The terms in Equation 25 can be separated into three types: (i) those terms whose values are known at time t with information up to time t , (ii) those that depend only on disease control action, β_t , and (iii) those that depend only on test control action, θ_{t+1} . Hence, the minimization over β_t and θ_{t+1} can be separated as $V_t(\wp_t) = V_t^{(1)}(\wp_t) + V_t^{(2)}(\wp_t) + V_t^{(3)}(\wp_t)$ in which

$$\begin{aligned} V_t^{(1)}(\wp_t) &= (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \\ &\quad + \text{tr} [A_t (\hat{\Sigma}_{t|t} + \hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} - \hat{\Sigma}_{t-1|t} T_{t-1}')] \\ &\quad + \hat{\alpha}_{t|t}' ((T_t - I)' A_{t+1} (T_t - I) + T_t' P_{t+1} T_t) \hat{\alpha}_{t|t}, \end{aligned} \quad (26)$$

$$\begin{aligned} V_t^{(2)}(\wp_t) &= \min_{\beta_t} \{ \beta_t' (B_t + G_t' (A_{t+1} + P_{t+1}) G_t) \beta_t \\ &\quad + (\hat{\alpha}_{t|t}' ((T_t - I)' A_{t+1} + T_t' P_{t+1}) G_t) \beta_t \\ &\quad + \beta_t' (G_t' (A_{t+1} (T_t - I) + P_{t+1} T_t) \hat{\alpha}_{t|t}) \}, \end{aligned} \quad (27)$$

$$\begin{aligned} V_t^{(3)}(\wp_t) &= \text{tr} [(A_{t+1} + P_{t+1}) (T_t \hat{\Sigma}_{t|t} T_t')] \\ &\quad + \text{tr} [(A_{t+1} + P_{t+1}) Q_t] + b_{t+1} \\ &\quad + \min_{\theta_{t+1}} \{ l_{t+1}(\theta_{t+1}) + \text{tr} [A_{t+1} (\hat{\Sigma}_{t+1|t+1} - \hat{\Sigma}_t^* \hat{\Sigma}_{t+1|t+1} \\ &\quad - \hat{\Sigma}_{t+1|t+1} \hat{\Sigma}_t^*)] + V_{t+1}^\theta(\wp_{t+1}^\theta) \}. \end{aligned} \quad (28)$$

As before, the minimization over β_t is denoted by \tilde{J}_t , so

$$\begin{aligned} \tilde{J}_t &= \min_{\beta_t} \{ \beta_t' (B_t + G_t' (A_{t+1} + P_{t+1}) G_t) \beta_t \\ &\quad + (\hat{\alpha}_{t|t}' ((T_t - I)' A_{t+1} + T_t' P_{t+1}) G_t) \beta_t \\ &\quad + \beta_t' (G_t' (A_{t+1} (T_t - I) + P_{t+1} T_t) \hat{\alpha}_{t|t}) \} \end{aligned} \quad (29)$$

The minimization over β_t can be performed by completion of squares (similar to what is done in Lemma 5 for the minimization over β_N) to yield Equations 30–33. The optimal disease control at time t is given by

$$\beta_t^* = -U_t \hat{\alpha}_{t|t}, \tag{30}$$

in which the control law, U_t , is given by

$$U_t = (B_t + G_t'(A_{t+1} + P_{t+1})G_t)^{-1} (G_t' A_{t+1}(T_t - I) + G_t' P_{t+1} T_t). \tag{31}$$

Moreover, the result of minimization over β_t is given by

$$\tilde{J}_t = -\hat{\alpha}'_{t|t} \tilde{P}_{t+1} \hat{\alpha}_{t|t}, \tag{32}$$

in which

$$\begin{aligned} \tilde{P}_{t+1} = & ((T_t - I)' A_{t+1} G_t + T_t' P_{t+1} G_t) \\ & (B_t + G_t'(A_{t+1} + P_{t+1})G_t)^{-1} \\ & (G_t' P_{t+1} T_t + G_t' A_{t+1}(T_t - I)). \end{aligned} \tag{33}$$

As seen in Equation 30, the optimal disease control β_t^* is a linear function of the filtered state mean $\hat{\alpha}_{t|t}$. It is worth noting that this function (more precisely, the control law U_t) depends only on parameters of the system dynamics and the objective function cost inputs. Hence, the control law is data independent and can be calculated offline prior to solving the measurement and control problems. As seen in Equations 4–11, the optimal state estimation is independent of the control law. We also just showed that the optimal control law is independent of the state estimation. This completes proof of Theorem 1. \square

Replacing the minimization over β_t in Equation 27 by its value given by Equation 32 results in

$$\begin{aligned} V_t(\varphi_t) = & (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \\ & + \text{tr}[A_t (\hat{\Sigma}_{t|t} + \hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} - \hat{\Sigma}_{t-1|t} T_{t-1}')] \\ & + \hat{\alpha}'_{t|t} ((T_t - I)' A_{t+1} (T_t - I) + T_t' P_{t+1} T_t - \tilde{P}_{t+1}) \hat{\alpha}_{t|t} \\ & + \text{tr}[(A_{t+1} + P_{t+1})(T_t \hat{\Sigma}_{t|t} T_t')] \\ & + \min_{\theta_{t+1}} \{l_{t+1}(\theta_{t+1}) + \text{tr}[A_{t+1} (\hat{\Sigma}_{t|t+1} - \hat{\Sigma}_t^* \hat{\Sigma}_{t+1|t+1} \\ & - \hat{\Sigma}_{t+1|t+1} \hat{\Sigma}_t^{*'})] + V_{t+1}^\theta(\varphi_{t+1}^\theta)\} \\ & + \text{tr}[(A_{t+1} + P_{t+1})Q_t] + b_{t+1}. \end{aligned} \tag{34}$$

Letting

$$P_t = (T_t - I)' A_{t+1} (T_t - I) + T_t' P_{t+1} T_t - \tilde{P}_{t+1}, \tag{35}$$

and replacing $\text{tr}[(A_{t+1} + P_{t+1})(T_t \hat{\Sigma}_{t|t} T_t')]$ in Equation 34 by its other form given by Lemma 9, Equation 34 can be written as follows to match the form of value function we claimed in Equation 16.

$$\begin{aligned} V_t(\varphi_t) = & (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t})' A_t (\hat{\alpha}_{t|t} - \hat{\alpha}_{t-1|t}) \\ & + \text{tr}[A_t \hat{\Sigma}_{t|t}] + \hat{\alpha}'_{t|t} P_t \hat{\alpha}_{t|t} + \text{tr}[P_t \hat{\Sigma}_{t|t}] \\ & + \text{tr}[A_t (\hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} - \hat{\Sigma}_{t-1|t} T_{t-1}')] \\ & + \text{tr}[(\tilde{P}_{t+1} + A_{t+1} T_t + T_t' A_{t+1} - I) \hat{\Sigma}_{t|t}] \\ & + \min_{\theta_{t+1}} \{l_{t+1}(\theta_{t+1}) + \text{tr}[A_{t+1} (\hat{\Sigma}_{t|t+1} - \hat{\Sigma}_t^* \hat{\Sigma}_{t+1|t+1} \\ & - \hat{\Sigma}_{t+1|t+1} \hat{\Sigma}_t^{*'})] + V_{t+1}^\theta(\varphi_{t+1}^\theta)\} \\ & + \text{tr}[(A_{t+1} + P_{t+1})Q_t] + b_{t+1}. \end{aligned} \tag{36}$$

Now, by comparing Equations 36 and 16, it can be easily seen that for $t = 1, \dots, N - 1$

$$\begin{aligned} V_t^\theta(\varphi_t^\theta) = & \text{tr}[A_t (\hat{\Sigma}_{t-1|t} - T_{t-1} \hat{\Sigma}_{t-1|t} - \hat{\Sigma}_{t-1|t} T_{t-1}')] \\ & + \text{tr}[(\tilde{P}_{t+1} + A_{t+1} T_t + T_t' A_{t+1} - I) \hat{\Sigma}_{t|t}] \\ & + \min_{\theta_{t+1}} \{l_{t+1}(\theta_{t+1}) + \text{tr}[A_{t+1} (\hat{\Sigma}_{t|t+1} - \hat{\Sigma}_t^* \hat{\Sigma}_{t+1|t+1} \\ & - \hat{\Sigma}_{t+1|t+1} \hat{\Sigma}_t^{*'})] + V_{t+1}^\theta(\varphi_{t+1}^\theta)\}, \end{aligned} \tag{37}$$

and

$$b_t = \text{tr}[(A_{t+1} + P_{t+1})Q_t] + b_{t+1}, \tag{38}$$

while from Equations 18 and 19 we know for $t = N$

$$\begin{aligned} V_N^\theta(\varphi_N^\theta) = & \text{tr}[A_N (\hat{\Sigma}_{N-1|N} - T_{N-1} \hat{\Sigma}_{N-1|N} \\ & - \hat{\Sigma}_{N-1|N} T_{N-1}')] + \text{tr}[\tilde{P}_{N+1} \hat{\Sigma}_{N|N}], \end{aligned} \tag{39}$$

and

$$b_N = \text{tr}[A_{N+1} Q_N]. \tag{40}$$

Hence, the proof of Theorem 3 (i.e., the value function we claimed in Equation 16) is complete.

Note that the dynamic program defined by value function $V_t^\theta(\varphi_t^\theta)$ for $t = 1, \dots, N$ can be solved to find the optimal monitoring schedule. We succinctly highlight the elements of this dynamic program. The information state of the model (which can be thought as summarizing the sufficient statistics) is $\varphi_t^\theta = (\hat{\Sigma}_{t|t}, \hat{\Sigma}_{t-1|t})$. The action space is $\theta \in \Theta$, where Θ is the set of all available tests for the corresponding disease. Equations 8, 9, and 11 state the system dynamics of the model, while Equation 37 specifies the optimality equation (or Bellman equation). The boundary conditions are given in Equation 39. This completes the proof of Theorem 2.

5. Case Study of Glaucoma

Thus far we have presented the modeling framework in its general form and derived the optimal disease and test control actions. In this section, we provide a

proof of concept by applying our approach to glaucoma and demonstrating how it can help guide clinicians in tailoring disease monitoring and treatment control.

Glaucoma is a major public health problem affecting almost 3 million patients in the United States (Vajaranant et al. 2012) and over 60 million patients worldwide (Tham et al. 2014). Glaucoma is the second leading cause of blindness in the US and a leading cause of visual impairment among Americans (Stein et al. 2011). In this section, we show how the modeling framework and solution approaches described in sections 3 and 4 can be applied to help clinicians in caring for patients with glaucoma. Furthermore, we elaborate on additional features of our approach designed specifically for glaucoma. Numerical results presented in this section are based on data from patients with glaucoma who were enrolled in two large clinical trials.

5.1. Glaucoma

Glaucoma is a progressive eye disease which can cause irreversible vision loss and blindness if not adequately monitored and treated. From a societal perspective, the direct medical costs of managing glaucoma are estimated to total over 2.86 billion USD annually (Rein et al. 2006). Furthermore, on a per patient basis, costs more than quadruple when patients progress from early to advanced glaucoma (Lee et al. 2006). Key risk factors associated with glaucoma development and its progression include: non-white race, older age, elevated intraocular pressure (IOP), genetics, and family history (Tielsch et al. 1990). It is worth noting that the patient's IOP (i.e., the pressure inside the eye) is the only known controllable/modifiable glaucoma risk factor. Therefore, the current management of glaucoma focuses on lowering the eye pressure by establishing a "target IOP," which is the level of IOP that the clinician feels is low enough to slow disease progression sufficiently (Jampel 1997).

Patients with glaucoma are monitored for disease progression using several quantitative tests. Two primary methods to monitor a patient are: (i) tonometry (or measuring the IOP), and (ii) perimetry (or visual field (VF) testing). Tonometry measures the patient's IOP, is relatively easy to perform, and is part of a standard eye examination. In most patients, vision loss from glaucoma occurs because elevated IOP damages the optic nerve, the structure that carries visual information to the brain for interpretation and processing (Sommer et al. 1991). Vision loss caused by glaucoma cannot be reversed but treatment can help slow or halt future vision loss (Weinreb and Khaw 2004). Lowering IOP has been shown to reduce the risk of disease progression and is the primary treatment goal

in glaucoma management (Hyman et al. 2010, Leske et al. 2003, Maier et al. 2005). With glaucoma, patients often progressively lose peripheral vision and eventually as the disease worsens, central vision. The VF test quantifiably measures the extent and rate of peripheral vision loss by examining the sensitivity of the eye to light stimuli of varying intensities. It is more time-consuming than checking IOP but provides important information on the status of the disease. VF testing can be anxiety-inducing and challenging for patients as it requires patient attention and cooperation. Two key global performance measures from VF testing include Mean Deviation (MD) and Pattern Standard Deviation (PSD), which estimate the extent deviation of peripheral vision from a reference population who do not have glaucoma (Choplin and Edwards 1995). MD is usually a negative number; higher values of MD (i.e., values closer to zero) correspond to better vision quality (less vision loss). PSD is usually a positive number. MD, PSD, and IOP are all measured on a continuous scale.

It is well established from prior work that both IOP and VF tests can be associated with noise. For example, patient MD and PSD performance on an automated VF test can fluctuate considerably from one test to the next (Choplin and Edwards 1995). Likewise, IOP can fluctuate from hour to hour and day to day (Wilensky et al. 1993). To take such noise into consideration in deciding how to monitor the patient optimally and determine an appropriate target IOP, we harness the Kalman filter method (Kalman 1960) to remove noise from the raw measurements and provide a dynamic model for the state over time.

There are a number of treatments available to lower the IOP for a patient with glaucoma. Different eye drops, laser therapies, and incisional surgery can reduce the IOP to any number above 6 mmHg. However, glaucoma medications can be expensive and can have serious side effects including stinging, blurred vision, eye redness, itching, burning, low blood pressure, reduced pulse rate, fatigue, shortness of breath, headache, and depression. Therefore, the goal is to find the optimal target IOP for each patient, which corresponds to an IOP that is low enough to sufficiently slow disease progression but not overly low so as to subject the patient to unnecessary treatments. The appropriate target IOP varies from one patient to another as some patients (eyes) experience progression at a certain IOP level while others do not.

In current practice it is common to use fixed-interval monitoring regimes to test for disease progression. Furthermore, for each patient (eye), the eye care professional must regularly make a gestalt-based estimate of a reasonable target IOP that considers the patient's individual risk of disease progression and the side effects and costs associated with lowering the

IOP. Prior studies in glaucoma management have employed various machine learning algorithms to predict progression of glaucoma progression (Bowd et al. 2012, Goldbaum et al. 2012, Sharpsten et al. 2014, Yousefi et al. 2014, Yousefi et al. 2016). Simulation models of glaucoma progression have been developed to evaluate the cost-effectiveness of frequent testing in patients with glaucoma (Boodhna and Crabb 2016) as well as to evaluate different treatment strategies (Van Gestel et al. 2010). However, to the best of our knowledge, no optimization-based approach presently exists to jointly determine how best to monitor a patient with glaucoma and how best to control the disease. Our approach considers the history of the patient (prior test performances) and her unique disease dynamics to provide clinicians with (i) a personalized monitoring regime to achieve an accurate assessment of whether there is disease progression (exploration), and (ii) in a menu format, how the glaucoma is likely to progress for different target IOP levels; the doctor can leverage these to devise an individualized treatment plan for the particular patient (exploitation). As a feature of our model, the clinician is able to select the desired aggressiveness level to monitor and treat the patient based on the unique characteristics/circumstances of the specific patient. We will elaborate on this menu of options in subsection 5.5.

Figure 2 depicts a high-level overview of our dynamic monitoring and control decision support framework for patients with glaucoma. It serves as a schematic table of contents for what follows in the remainder of this section. It shows how different parts of the modeling framework are linked together and provides the corresponding subsection/equations for each part. At each office visit, IOP and/or VF test(s) are performed. The raw measurements from these tests (which are known to be noisy) are fed into a

Kalman filter to obtain an optimal estimation of the current disease state for a particular eye. Then, the estimate of the previous state is refined via a Kalman smoother given the new information acquired during the visit. Each patient’s label/type (fast-, slow-, or non-progressor) determines the estimate of how quickly the patient’s glaucoma is likely to progress in the future. The decision support tool provides an optimal monitoring schedule (i.e., the timing of the next exam/test and which tests the patient should take) and a personalized target IOP for the patient for different aggressiveness levels/options (super-high, high, moderate, low, or super-low). Finally, the clinician chooses an aggressiveness option from the menu of choices that is appropriate for the individual patient.

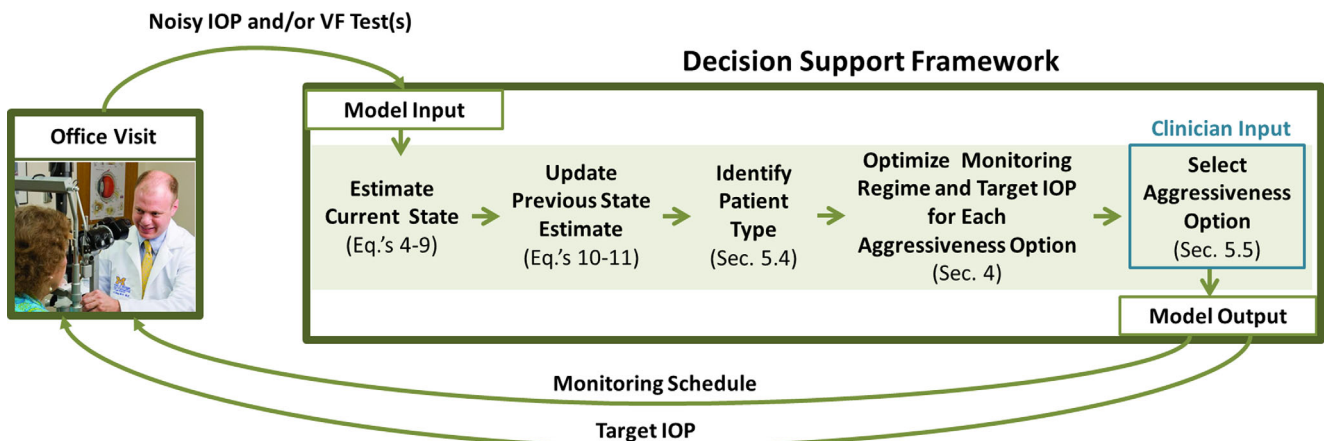
5.2. Patient Disease State

We use a nine-dimensional state vector, α_t , to model the patient’s disease state. The elements of the state vector include MD, PSD, and IOP together with their discrete time first and second derivatives (i.e., velocity and acceleration, respectively);

$$\alpha_t = \begin{bmatrix} MD & MD' & MD'' & PSD & PSD' & PSD'' \\ IOP & IOP' & IOP'' \end{bmatrix}. \tag{41}$$

The nonlinear behavior of disease dynamics is captured by including the velocity and acceleration of key disease state elements in the state vector. This is known to be an effective way to linearize a nonlinear model of state evolution (see Bertsekas 1995). MD , PSD' , and IOP' are the slope of a linear regression of the latest three MD, PSD, and IOP measurements, respectively. MD'' , PSD'' , and IOP'' are the difference of the latest two MD' , PSD' , and IOP' values divided by the time interval between them. It is

Figure 2 Illustration of the Decision Support Framework Showing the Model Inputs and Outputs As Well As the Sequence of Main Steps of the Disease Monitoring and Control Algorithm [Color figure can be viewed at wileyonlinelibrary.com]



worth noting that, in the application of our modeling framework to glaucoma, the disease control action (i.e., β_t in Equation 1), is defined as the amount in *mmHg* to reduce the patient's IOP at time t . This value is determined by the optimization model, not by the user/clinician.

5.3. Data

To parameterize and validate our model, we use data from two multi-center randomized clinical trials, the Collaborative Initial Glaucoma Treatment Study (CIGTS) (Musch et al. 1999) and the Advanced Glaucoma Intervention Study (AGIS) (Ederer et al. 1994). These datasets are chosen because they include structured tonometry and perimetry data (IOP and VF readings) of glaucoma patients taken every 6 months during the course of the trials. We match the time step of our LQG model with these datasets to avoid the need for data interpolation (i.e., there is a 6-month time interval between periods t and $t + 1$ in our model).

CIGTS studied newly diagnosed glaucoma patients with mild to moderate disease who were randomized to medical or surgical therapy and were followed for up to 11 years with IOP and VF tests taken every 6 months to assess disease progression. In AGIS, patients with advanced glaucoma were randomized to laser therapy or incisional surgery and followed for up to 11 years with IOP and VF readings taken every 6 months.

For the purpose of this case study, we excluded patients from these trials with fewer than five readings. We also restricted our focus to the patients who received either medical or laser therapy. We excluded glaucoma patients who received incisional surgical interventions because incisional surgery can abruptly change disease progression dynamics in a manner that may be challenging to accurately model. We randomly divided all eligible participants from the trials (571 participants) into two sets of equal size: (i) a training set, and (ii) a testing set. Both sets have approximately the same number of mild, moderate, and advanced glaucoma patients, with similar numbers of white and black patients, and equal numbers of patients contributing data from each trial. The training set is used for parametrization and calibration of our state space model and the testing set is used to evaluate the performance of our approach.

5.4. Patient Types (Fast/Slow/Non-Progressor)

A fast-progressing patient is someone whose glaucoma is worsening rapidly and is part of the subset of patients at the greatest risk of blindness. Although there is presently no gold standard for defining glaucoma fast-progressors, prior literature considers a loss of MD >1 dB per year as a reasonable identifying

feature of patients who are exhibiting fast-progression of glaucoma (see Gardiner and Crabb 2002, Heijl et al. 2013). We built our algorithms based on this definition of fast-progressors. To classify each patient, we calculated the slope obtained from a linear regression of their entire set of MD readings and labeled them as a:

- *fast-progressor* if the MD slope is declining by ≥ 1 dB/year,
- *slow-progressor* if the MD slope is declining between 0 and 1 dB/year, and
- *non-progressor* if the MD slope is not declining.

5.5. Aggressiveness Levels/Policies/Options/Regimes

In clinical practice, the goals of care must be tailored to each individual patient's needs. Rather than proposing one solution for all patients, a more powerful and useful approach is to provide the clinician with a range of options for how much effort (both from provider and the patient) will be put into monitoring and how aggressively IOP should be lowered such that future progression can be slowed. For instance, clinicians will likely see the need to monitor and treat a young patient who only has sight in one eye more aggressively than an older patient with mild glaucoma in one eye and no glaucoma in the other eye. All else being equal, less aggressiveness may be appropriate for a patient with multiple systemic medical comorbidities who is likely to die before they go blind from glaucoma. As a useful and not overly complex approach, our clinical collaborator suggested creating optimization models tailored to three regimes, or "options," for monitoring and treatment to facilitate adoption of this system into clinical practice. We refer to these three options as *low*, *moderate*, and *high levels of aggressiveness* to represent the level of intensity in care and monitoring. We also define two extreme levels of aggressiveness: *super-high* and *super-low*. Note that we choose these terms only for convenience in presenting the five options and to make it easier for the reader to remember the order of them in terms of how aggressively they test and treat patients; they are not meant to correspond to any existing terms or approaches currently used in clinical practice. These five options are useful not only for sensitivity analysis but also suggest an effective way to implement a decision support system so that clinicians can pursue monitoring and treatment with the level of intensity that they, together with the patient, determine to be the most appropriate for each individual. Appendix C provides insight into selecting an appropriate aggressiveness option. It is important to differentiate between the aggressiveness level and the optimal control (target IOP and monitoring schedule).

In the treatment planning cycle, an aggressiveness level is chosen by the clinician based on the tradeoff between the costs of treatment (e.g., side-effects and financial costs) and the benefits for the patient being treated (e.g., preventing future vision loss from glaucoma). This provides an aggressiveness regime under which each patient's monitoring and treatment plans can be optimized individually with respect to their particular disease trajectory and history of readings. Thus, the optimal control is personalized in terms of monitoring and treatment decisions driven by their personal data and governed by the model for the particular aggressiveness regime selected. Note that the regime itself does not predetermine the monitoring and treatment plans for the individual. For the same aggressiveness regime, two individuals will have different histories of test results and, therefore, potentially different optimal decisions of when to take the next test and which tests to take as well as different target IOPs over time. Different aggressiveness levels translate into different cost matrices A , B , and $l_t(\theta_t)$ in our objective function. Essentially, the more aggressive the policy the larger the cost of disease progression (A), which subsequently encourages the optimization problem to suggest a larger reduction in IOP (i.e., larger β_t^*) and more frequent testing. In summary, the aggressiveness level is chosen by the clinician in consultation with the patient based on treatment goals, which then sets the cost parameters of the model, whereas the actual target IOP and monitoring schedule are chosen by the optimization model unique to the patient's quantitative disease measurements over time. The five aggressiveness levels/options follow:

1. *Super-high aggressiveness option*, which drops the IOP immediately to 6 mmHg (an ideal level of IOP for patients with any severity of glaucoma, but one that may be impractical for many patients due to limitations with the effectiveness of existing interventions, side effects, and/or complications associated with attaining such a low IOP),
2. *High aggressiveness option*, which tends to lower the IOP by 40%–60% compared to the patient's treated level of IOP that was achieved in the CIGTS/AGIS clinical trials after the initial intervention was given,
3. *Moderate aggressiveness option*, which tends to lower the IOP by 20%–40% compared to the patient's treated level of IOP that was achieved in the CIGTS/AGIS clinical trials after the initial intervention was given,
4. *Low aggressiveness option*, which corresponds to the IOP achieved under no additional

interventions beyond those employed in the CIGTS/AGIS trials, and

5. *Super-low aggressiveness option*, which attempts to estimate progression of an untreated patient with glaucoma by removing the effect of existing interventions that were employed in CIGTS/AGIS on the patient's IOP.

It should be noted that the exact amount of IOP control suggested by high, moderate, and low aggressiveness policies is patient-specific and is optimized to yield the minimum total cost as defined by the objective function. In contrast, the super-high and super-low aggressiveness policies are static policies that do not take the objective function into account. They are mainly added for purposes of analysis and comparison, but they can still provide valuable insight in a clinical setting by presenting the clinician the "best" and "worst" case options and their forecasted impact on disease progression dynamics. These five options/levels also provide sensitivity analysis on the model cost parameters.

5.6. System Model Parameterization

The expectation maximization (EM) algorithm was employed to parameterize the state space model. EM is an iterative algorithm for finding the maximum likelihood estimate of model parameters in the presence of missing or unobserved data. The EM algorithm alternates between the expectation step (E-step) and the maximization step (M-step). In the E-step, raw, noisy readings are filtered and missing data is estimated based on the observed data and the most recent estimate of system parameters. In the M-step, the log-likelihood function of all data points is maximized assuming the missing data is given by the estimates from the E-step. For more information about the EM algorithm, please see Dempster et al. (1977), Digalakis et al. (1993), Ghahramani and Hinton (1996). While the model was presented in its general setting in section 3, for the purpose of this case study we assume the model parameters are time-invariant. The output of the EM algorithm is the best estimate of system matrices $T_t = T$, $Q_t = Q$, $Z_t = Z$, $H_t^{(\theta_t)} = H^{(\theta_t)}$ for $t = 1, \dots, N$, and initial state mean and covariance, $\hat{\alpha}_0$ and $\hat{\Sigma}_0$ (see subsections 3.1 and 3.2 for the definition of these parameters). We assume $A_t = A$, $B_t = B$, and $l_t(\theta_t) = l(\theta_t)$ are time-invariant. We further assume that $G_t = G = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0]^T$ for $t = 1, \dots, N$, because the control variable β_t is designed to control the patient's IOP only, which is the only controllable glaucoma risk factor. Furthermore, it is known that an intervention started or employed at time t has instantaneous effect on lowering the patient's IOP. For example, if patient's IOP

is 20 mmHg at time period 7 and the control $\beta_7^* = -3$ mmHg, the expected value of IOP right after time period 7 is 17 mmHg. This IOP reduction affects other state elements and progression dynamics in the following time period through the transition equation. Also, note that, in our case study of glaucoma, T , Q , Z , $H^{(\theta_i)}$, A , and \sum_0 are 9×9 matrices, G and $\hat{\alpha}_0$ are 9×1 vectors, and β_t , B , and $l(\theta_t)$ are scalars. None of the square matrices are assumed to be diagonal. This is important because, for example, the components of T represent how each component in the current state affects the future states. A non-diagonal component would represent the fact that a higher IOP will affect the change in MD in the next period. Hence, the current disease state (including the position, velocity, and acceleration of key glaucoma variables) directly affects disease progression.

We use the EM algorithm to obtain four sets of system parameters. These sets of parameters are obtained from (i) all patients in the training set, (ii) only fast-progressors, (iii) only slow-progressors, and (iv) only non-progressors, to enable stratified analyses, such as in Figure 4.

The model cost parameters were estimated based on input from our glaucoma specialist collaborator so that the model outputs are reasonable from a clinical perspective. Note that it is the relative costs (rather than each absolute cost) that play a key role in our analysis. It is beyond our scope to obtain definitive cost parameters; however, significant sensitivity analysis was performed around those estimates to understand the model's behavior better and to ensure that the model, in its entirety, provides credible decision support. In order to construct reasonable tradeoffs between cost parameters for each aggressiveness option in our model, we employed a human-in-the-loop methodology. We tested several parameter sets, showed the resulting optimal IOP controls and monitoring policies to our glaucoma specialist collaborator, and the clinician helped us fine-tune and calibrate cost parameters. We confirmed with our glaucoma specialist collaborator that the model generates clinically reasonable target IOPs and monitoring schedules for each patient under the desired aggressiveness level. For instance, under the high aggressiveness option, the model for fast-progressing patients (i.e., the most aggressive combination) should suggest taking both IOP and VF tests every 6 months and a target IOP of around 6–9 mmHg. Under the low aggressiveness option, the model for non-progressing patients (i.e., the least aggressive combination) should suggest no further IOP reduction and taking IOP and VF tests every 2 years. This is in line with

recommendations put forth by the American Academy of Ophthalmology Clinical Practice Guidelines (2015). For all other combinations of aggressiveness level and patient severity type, the model cost parameters were fine-tuned so it suggests a monitoring regime and target IOP level that are reasonable according to expert clinical opinion and are in between the two extreme combinations. The behavior of optimal policies is discussed more thoroughly in section 5.8.2.

5.7. Model Usage for a Glaucoma Patient

For a patient who is newly diagnosed with glaucoma with no prior history of IOP and VF readings, both tests are taken in every period (i.e., every 6 months) for the first five periods. Gardiner and Crabb (2002) found that five initial VF test results are a reasonable predictor of future glaucoma progression dynamics in most patients. We used five initial VF and IOP readings to (i) obtain baseline values for key disease state elements (i.e., MD, PSD, and IOP), (ii) calculate the velocities and accelerations of key state elements, (iii) warmup the Kalman filter and smoother, (iv) reduce the initial uncertainty surrounding a given patient's disease state, (v) calculate the initial 5-period rate/slope of MD progression to label the patient as a fast-, slow-, or non-progressor, and (vi) differentiate the patient from the population mean and tailor the disease transition model to the specific patient.

The system parameters obtained from all training patients are used in the Kalman filter and the Kalman smoother during the warmup period. At the end of the warmup period (i.e., after five readings), the patient is labeled as a fast-progressor, a slow-progressor, or a non-progressor based on her MD progression rate. Thereafter, the Kalman filter and smoother with type-specific system parameters (i.e., fast, slow, or non-progressor set of parameters) are used. Each time a subsequent test is taken, the MD progression slope is recalculated. We always consider the latest five filtered MD values to update the MD slope (i.e., a sliding window of length 5). Whenever the patient's latest MD slope indicates a label upgrade (i.e., the patient moves from non-progressor to slow/fast-progressor, or from slow-progressor to fast-progressor), the model (i) calls for a follow-up visit to take IOP and VF testing in the following time period, and (ii) labels the patient as a suspect of the higher label/category (e.g., slow-progressor suspect or fast-progressor suspect). If the label change is confirmed at the next follow-up visit, the higher label is assigned to the patient. Otherwise, the patient is returned to the previous lower label status. Note that our analyses take a conservative approach and do not allow any label

downgrading (on the recommendation of our glaucoma specialist collaborator). Once the label is upgraded for a patient, the model will recommend applying more IOP control (i.e., greater intensity of interventions) to slow glaucoma progression. Therefore, it can be expected that the MD will tend to decline less rapidly once the amount of IOP control is increased, resulting in a lower classification/label at some point. However, if we were to downgrade the patient’s label and let the model decrease the amount of IOP control, the patient would be at risk to start losing vision at the same rate as earlier in the disease course, which is not desirable. Therefore, we do not allow label downgrading for any patient once the label upgrade is confirmed. If a patient suspected of belonging to a higher category does not get a confirmatory result at the very next follow-up visit, then, the patient remains at the

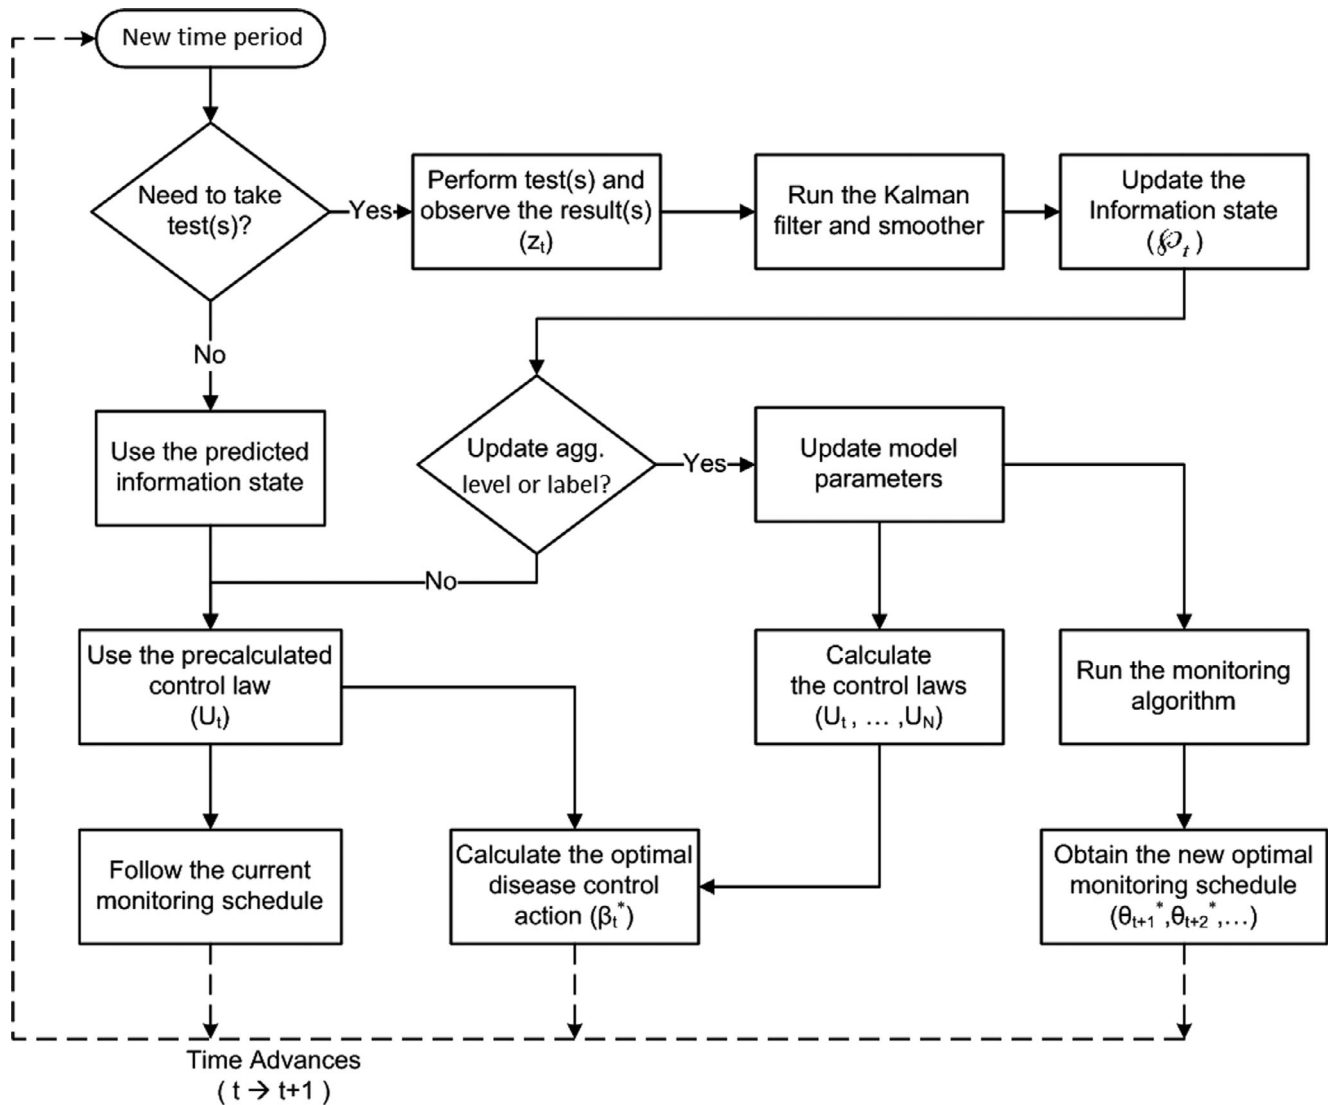
original/lower label he/she was previously at. The glaucoma monitoring and treatment control algorithm steps are illustrated in Figure 3.

5.8. Numerical Results

In this subsection, we test the performance of our dynamic disease monitoring and control algorithm on glaucoma patients from the CIGTS and AGIS clinical trials. We first validate our prediction model on the testing dataset, using the training dataset for parameterization. Then, we provide numerical results and examples on how the optimal policies behave. Lastly, we provide further results of the impact of optimal policies on patients with glaucoma.

5.8.1. Validation. We first validate that the model is good at forecasting future disease progression trajectory and then validate that the results are

Figure 3 Glaucoma Monitoring and Treatment Control Flow Diagram



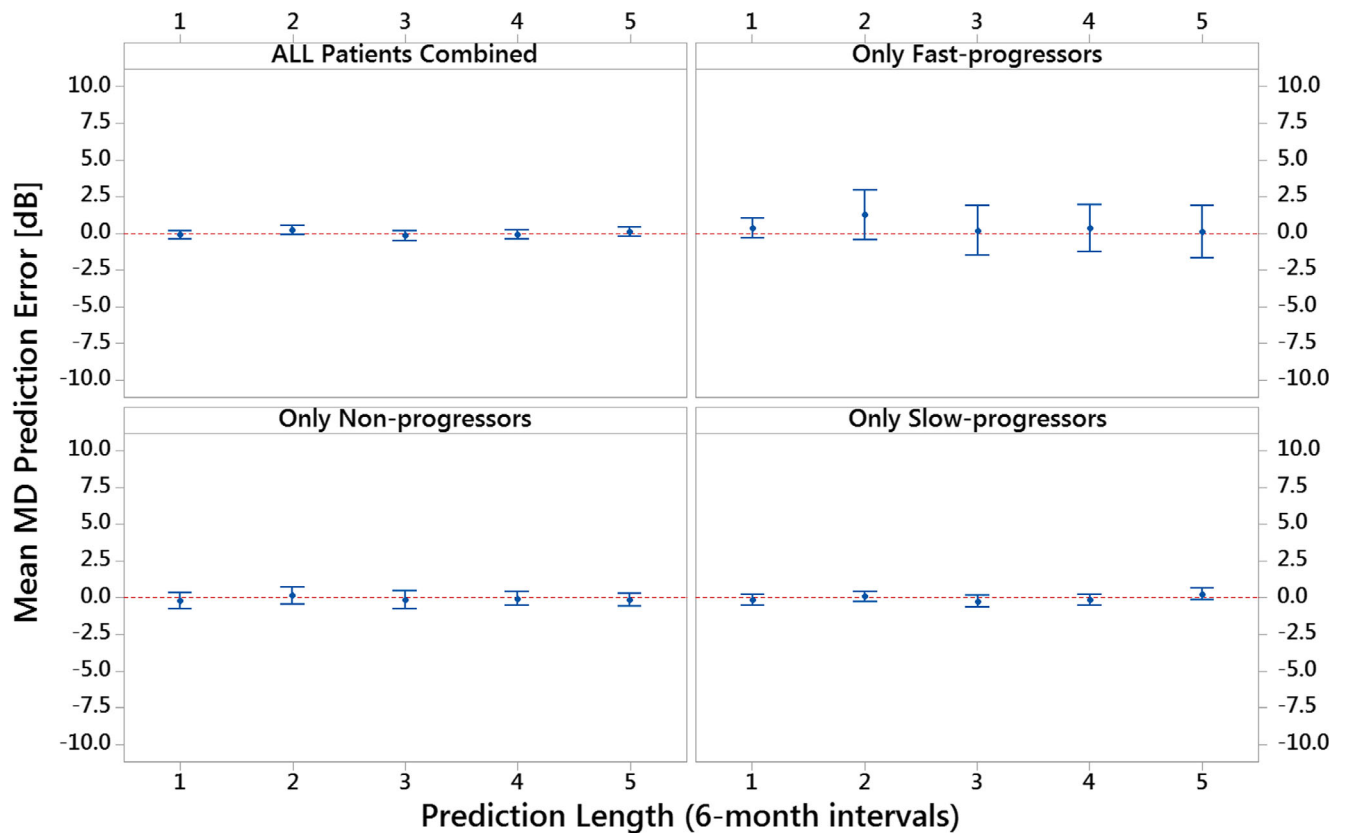
consistent with clinical expectations. Our modeling approach efficiently captures the system and measurement noises using a set of stochastic first-order vector difference equations. To evaluate the performance of our prediction model, calibrated using the EM algorithm on the entire training set, we used the first five data points of each patient in the testing dataset to warm up the Kalman filter and determine the patient type. Then, we predicted MD values for five periods into the future for each patient type (fast-progressor, slow-progressor, and non-progressor) and calculated the prediction error (i.e., the predicted state mean minus the actual reading as obtained from the patients during their follow-up in the trial). The predicted state mean was calculated using Kalman filter with $\beta_t = 0 \forall t$ (i.e., no additional IOP reduction beyond that in the clinical trial). Figure 4 shows the mean MD prediction error for up to five 6-month periods (2.5 years) into the future. The dots correspond to the mean error, and the bars represent the 95% confidence interval for the mean. These interval plots confirm that our prediction model has very little error in predicting the glaucoma state progression. One also sees that the fast-progressors (as defined in subsection

5.4) vary the most in the datasets, and this is reflected in greater uncertainty and error. However, 25 rounds of cross-validation were performed to confirm the robustness of our Kalman filter model to patient data. Here, we present the results of one of the cross-validation rounds for MD, as it is the most clinically useful state variable. Similar results were obtained for other state elements (PSD and IOP).

5.8.2. Evaluation of the Optimal Policies.


Having confirmed that the model accurately forecasts future MD values, we next test that the output provides clinically reasonable results as confirmed by our clinical collaborator. The structure of optimal IOP control generated by our model under the moderate or high aggressiveness policy is key to identifying the target IOP for each patient. We applied the high and moderate aggressiveness policies to all fast- and slow-progressing patients in the testing dataset to achieve a statistical characterization of how each policy behaves. For each group and each aggressiveness policy, we used the first five data points to warmup our model. We then recorded the amount of optimal IOP control suggested by our model in the next 20 time

Figure 4 Interval Plot of Mean MD Prediction Error for Different Prediction Lengths [Color figure can be viewed at wileyonlinelibrary.com]



Note: The dots correspond to mean error, and the bars represent 95% confidence interval for the mean.

Figure 5 Optimal IOP Controls Suggested by Our Model for Fast- and Slow-Progressing Patients Under the High and Moderate Aggressiveness Policies Over 10 Year [Color figure can be viewed at wileyonlinelibrary.com]



Note: Period 0 is the current time period (i.e., the period at which the IOP control starts).

periods (i.e., the following 10 years). Figure 5 depicts the results over all the patients in the testing dataset as box plots of optimal additional IOP control (β_t^*) applied in the current (i.e., period 0) and the following time periods. An IOP control of $-x$ mmHg corresponds to lowering the patient’s IOP by x mmHg more than what was achieved in AGIS/CIGTS. The bottom and top of each box are the first and third quartiles, respectively. The lower and upper whiskers extend to the minimum and maximum data points within 1.5 box heights from the bottom and top of the box, respectively. In our model, the optimization problem finds the optimal value for beta in each period (i.e., how much to reduce IOP in each time period). As seen in Figure 5, our feedback-driven model recommends further lowering the IOP within the first few time periods. Afterwards, the optimal additional IOP control is close to zero. This results in the patient’s IOP converging to “a number” over time. We call this number the “target IOP,” the common term used in the glaucoma medical community. As one would expect, the group of non-progressing patients do not get any additional benefit from further lowering their IOP since they exhibit no signs of progression at the IOP levels they maintain from the

treatments already employed in the trials. Therefore, they are not included in the graph.

We applied the five aggressiveness options to fast- and slow-progressing patients in the test dataset and obtained the following IOP-related metrics for each combination of patient type and aggressiveness policy: (i) target IOP (mmHg), (ii) additional IOP control (mmHg) applied in 10 years (representing the amount/intensity of treatment), and (iii) percentage of IOP change after 10 years. For each metric, we report the median and interquartile range (IQR), which are robust measures of location and scale, respectively. The IQR is the difference between the upper and lower quartiles and provides a range that contains 50% of the data. As seen in Figure 5, the optimal additional IOP reduction is almost entirely applied during the first six periods (3 years) of employing the control policy. Hence, we evaluate the patient’s IOP after 10 years, which is a sufficient length of time for the IOP to become stable under treatment. We tested our IOP control model with longer time horizons and obtained similar results.

Table 2 summarizes the IOP-related results. For instance, applying the high aggressiveness policy to

Table 2 Comparison of the Effect of Different Aggressiveness Options on Patient’s IOP for Fast and Slow-Progressing Patients in CIGTS and AGIS

	Target IOP [mmHg] (10-year predicted IOP)		Cumulative/Additional IOP control [mmHg] applied over 10 years		% of IOP change after 10 years	
	Median	IQR	Median	IQR	Median	IQR
<i>Fast-progressors</i>						
Super-high aggressiveness policy	6	0	−10.57	4.97	−64.17	9.50
High aggressiveness policy	7.17	1.70	−9.36	4.17	−55.24	8.76
Moderate aggressiveness policy	9.62	2.66	−7.43	3.91	−41.19	10.28
Low aggressiveness policy	18.45	5.69	0	0	+3.47	29.32
Super-low aggressiveness policy	22.98	6.37	+3.91	1.21	+26.53	39.23
<i>Slow-progressors</i>						
Super-high aggressiveness policy	6	0	−12.00	3.09	−65.46	10.61
High aggressiveness policy	9.52	2.07	−6.38	4.01	−44.11	8.51
Moderate aggressiveness policy	13.08	2.02	−2.13	1.52	−22.69	23.29
Low aggressiveness policy	14.95	2.66	0	0	−11.44	34.37
Super-low aggressiveness policy	18.66	4.24	+4.93	1.27	+11.88	43.01

Notes: Target IOP: IOP value in period 25 (i.e., 10 years into the future). Cumulative/Additional IOP control applied over 10 years: $\sum_{t=5}^{24} \beta_t$. % of IOP change after 10 years: $100 \times (\text{Target IOP} - \text{Baseline IOP}) / (\text{Baseline IOP})$.

fast-progressors results in a median target IOP of 7.17 mmHg. This can be achieved by administering a median additional 9.36 mmHg IOP reduction from the baseline level of IOP attained in the trials. Such a target IOP is, on average, 55.24% lower than the baseline IOP (the IOP at the beginning of the 10-year prediction period) of fast-progressing patients in the trials. Since target IOP is an important metric that helps guide the clinician in selecting the appropriate treatment plan for the patient, the distribution of target IOPs is also given in Figure S1 in Appendix D.

Table 3 summarizes the optimal monitoring regime for different combinations of patient type and aggressiveness level. For example, under the moderate aggressiveness level, the model for slow-progressing patients recommends measuring IOP every 6 months and checking the visual field every 12 months. It is worth noting that these protocols remain optimal as long as (i) the patient follows the monitoring schedule (i.e., does not miss a test), (ii) the patient type/label remains unchanged, and (iii) the doctor does not change the aggressiveness level. If any of the three criteria are not met, the model modifies the monitoring schedule to account for the missing information or the change in patient label/aggressiveness level. The monitoring regimes presented in Table 3 and the range and mean of target IOPs presented in Figure S1


are clinically appropriate in the professional opinion of our glaucoma specialist collaborator. Appendix E estimates costs of the optimal monitoring policies.

5.8.3. Menu of Options. Now that we have validated our model and elaborated on the structure of the optimal policies, we provide an example of how our decision support tool can help guide clinicians in managing a patient with glaucoma. Figure 6 depicts the glaucoma progression trajectory (change in MD over time) for a randomly chosen fast-progressing patient from the AGIS trial. The figure depicts a sample output of the decision support tool (with regard to disease control) that compares how this patient is likely to progress over the following 10 periods (5 years) under different aggressiveness options defined in subsection 5.5. As demonstrated in Figure 6, the patient progresses much slower and would have better MD values (i.e., vision quality) 5 years into the future as the aggressiveness of IOP control is increased. The figure provides a menu of options related to how aggressively the doctor wants to treat the patient, depicts the future disease progression trajectory, and provides the optimal target IOP and monitoring schedule for each aggressiveness option. The doctor is then able to select the right aggressiveness option based on evolving needs of the patient,

Table 3 Optimal Monitoring Regime for Different Combinations of Patient Type and Aggressiveness Level

	Fast-progressors	Slow-progressors	Non-progressors
High aggressiveness	IOP + VF every 6 months	IOP + VF every 6 months	IOP + VF every 12 months
Moderate aggressiveness	IOP + VF every 6 months	IOP every 6 months, VF every 12 months	IOP every 12 months, VF every 24 months
Low aggressiveness	IOP every 6 months, VF every 12 months	IOP 3× in 24 months, VF every 24 months	IOP + VF every 24 months

Figure 6 An Example of the Trajectory of Glaucomatous Progression as Captured by Changes to MD over Time by Employing Each of the Five Different Aggressiveness Policies for a Sample Fast-Progressing Patient from the AGIS Study [Color figure can be viewed at wileyonlinelibrary.com]



Note: Higher MD correlates with better vision quality.


adherence, health status, and other personal or clinical factors.

5.8.4. Insights into Treatment Effectiveness by Patient Type. Figure 7 graphs the average MD loss per year against the total IOP reduction applied under different aggressiveness policies for all fast- and slow-progressing patients in the testing set of CIGTS and AGIS trials. This graph provides three important insights for managing patients with glaucoma: (i) The curve for fast-progressors has a steeper slope, which indicates that this group of glaucoma patients benefits the most from further lowering of IOP from levels attained in CIGTS/AGIS. (ii) It can be deduced from Figure 7 that the low aggressiveness policy (point B), which, roughly corresponds to treatment using eye drops or laser therapy, works well enough for most slow-progressors since the curve is fairly flat around point B on the slow-progressors curve. In other words, increasing the aggressiveness level from low to moderate/high has only a minimal advantage at preserving sight for this group of patients. This highlights the importance of differentiating patients by progression type. Treating all patients similarly risks over-treating for little gain for some patients or irreversible vision loss from disease progression due to under-treating others. It can also be seen that slow-

progressors gain long-term benefits if treated under the super-high aggressiveness policy (point E), which often requires incisional surgery. Therefore, for slow-progressing patients, the doctor may choose either the low or super-high aggressiveness policy, depending on the individual’s life expectancy, severity of glaucoma, other personal and medical factors, and the patient’s preferences. (iii) The steep slope of the fast-progressors’ curve around point B indicates that vision loss could be significantly averted (even in the short-term) by further reduction of their IOP. Hence, moderate, high, or super-high aggressiveness policies (points C, D, and E on the graph, respectively) may be more suitable for most fast-progressing glaucoma patients and these patients often require incisional surgery. The same result is also verified by plotting the MD loss averted over 10 years by following the IOP controls suggested by our model as shown in Figure S2 in Appendix D. Moreover, Table S1 in Appendix D provides a comparison of the performance of different IOP control options against the low aggressiveness option for patients in our testing dataset.

5.8.5. Sample Application of the Model in Practice. In this subsection, we provide an illustration of how our modeling framework may be used to

Figure 7 Average MD Loss Per Year Can Be Reduced by Applying More IOP Control [Color figure can be viewed at wileyonlinelibrary.com]




Note: Fast-progressors get more benefit from lowering their eye pressure.

guide monitoring and control of a patient with glaucoma. Figure 8 portrays the disease trajectory of a sample patient. After the warmup period (i.e., the first 5 periods) the patient is initially identified as a non-progressor. In our example, the clinician chooses the low aggressiveness option to monitor and control the patient. Subsequently, the model suggests taking an IOP reading every year and a VF test every other year. We assume the clinician and the patient follow this protocol. The patient remains a non-progressor up to period 13, when she becomes classified as a slow-progressor suspect. This slow-progressor status is confirmed after obtaining IOP and VF testing at a follow-up visit in time period 14. When the patient becomes a confirmed slow-progressor, for the sake of this example, assume the doctor decides to increase the aggressiveness level to the moderate aggressiveness policy. Under this policy, the model recommends (i) lowering the IOP from 24 to about 21 mmHg, (ii) measuring the IOP every 6 months, and (iii) taking a VF test every year. After 1.5 years (i.e., at time period 17) the doctor and patient decide to increase the aggressiveness level further and continue care under

the high aggressiveness policy. This policy suggests taking both IOP and VF tests every 6 months and recommends additional IOP reduction. Figure 8 also illustrates how the patient’s glaucoma would likely progress after period 14 if the doctor and patient had maintained the low aggressiveness IOP control policy during periods 5–14.

While this example relates to managing a single patient, a few aspects should be highlighted. (i) As described in subsection 5.7, glaucoma patients do not always maintain the same progression rate over time. Recall that each time a test is taken, the model updates the MD slope estimate; hence, it is possible that a patient moves from non-progressor status to slow- or fast-progressor, or from slow-progressor to fast-progressor status. (ii) As described in subsection 5.8.2, whenever the patient’s label is changed or the clinician decides to change the aggressiveness level, the model modifies the monitoring regime subsequently. (iii) As described in subsection 5.8.4, there is little gain (in terms of preventing vision loss) in increasing the aggressiveness level from low to high for slow-progressing patients. Note the big gap in the

Figure 8 A Non-Progressing Patient Becomes a Slow-Progressor; The Clinician Tailors Care by Increasing the Aggressiveness Level [Color figure can be viewed at wileyonlinelibrary.com]



Notes: The first 5 periods are warmup time. From period 5 to 13, the patient is a non-progressor, and the doctor selects the low aggressiveness policy. In period 13, the patient becomes a slow-progressor suspect, and this label upgrade is confirmed in period 14. The doctor treats the patient under moderate aggressiveness policy from period 14 to 17. In period 17, the doctor increases the aggressiveness policy to high in order to slow the progression rate further. Periods 14–25 show forecasted values.

optimal IOP under the low and high aggressiveness policies at time period 25 in Figure 8. However, this gap results in a very small difference in the patient’s MD values. Benefiting from this type of insight in a busy clinic can significantly enhance the ability of ophthalmologists and optometrists to take care of patients with glaucoma appropriately.

6. Model Limitations and Future Work

While our modeling framework has shown great potential in improving the monitoring and control of patients with glaucoma, it comes with a few limitations and areas for improvement. First, we only considered the IOP and VF tests in developing a monitoring schedule for the patient, whereas there are additional tests that can also be used to monitor glaucomatous progression in practice. For instance, in well-equipped clinics, optical coherence tomography (OCT) is a non-invasive imaging test that measures the thickness of retinal nerve fiber layer (RNFL) (see Schuman et al. 2004). This newer testing modality was not commercially available at the time of the

CIGTS and AGIS clinical trials on which our analysis is based. Fortunately, the decision framework we have developed is scalable and can easily accommodate quantitative measurements of tests such as OCT. Should newer modalities for quantitatively assessing the status of a patient’s glaucoma arise, measurements from such modalities can be incorporated into the model as well. In the future, we hope to acquire access to other data sources which contain OCT data and expand our state vector to accommodate data from this testing modality.

Second, we used 6-month spaced time intervals because CIGTS and AGIS datasets contain readings of patient’s IOP and VF every 6 months. One can leverage the same modeling framework for data that is collected more or less frequently (e.g., monthly or every 3 months) without loss of generality. However, expanding our algorithm to handle unequally spaced data automatically is another potential path for future research.

Third, it is the case that our current approach uses one system dynamics model for each of the three distinct patient types (i.e., fast-, slow-, and non-

progressors). In other words, the dynamics that govern state transition are not patient specific and are not dynamically updated with each new measurement of MD, PSD, and IOP. It is beyond the current state of the art to integrate Kalman filtering models with models of uncertain dynamic changes in patient type in an optimal manner. Fourth, our state transition models are not designed to prevent any improvement in disease state completely; in fact, there is a non-zero probability that MD may increase due to the Gaussian noise assumption (noise can take positive and negative values). In our model, the state is just an estimate of the patient's true state, so the state estimate may increase or decrease (due to measurement noise) even while the true state of the disease is continuously declining. Nevertheless, our models still perform well in describing an irreversible disease such as glaucoma. This is in part because measurement noise indicates that the previous state may have been mis-measured. Thus, occasional improvement of state is expected to be observed in the short term, but in our model every patient will exhibit a non-increasing trend in the long term.

Finally, future research can focus on further studying the model cost parameters. In this research, we relied on expert opinion and developed a set of cost parameters for each aggressiveness policy that performed well on our extant data. Future research, however, can try to develop algorithms to optimize the balance between the cost of losing vision over time from glaucoma, cost of purchasing medications/undergoing surgery, cost of office visit and diagnostic testing, anxiety and stress of undergoing glaucoma tests, and side effects and complications of medical and surgical therapy to lower IOP further. Also, notice that we only considered fixed cost parameters in this study (i.e., the effect and cost of the control is independent of the current state variable). While the difficulty/cost of modifying a risk factor may vary based on the current level of the factor, we believe our state-invariant cost assumptions to be a valid approximation of the true costs. This is confirmed by the fact that applying the model to clinical trial patients creates control policies that are considered well within the realm of practical treatment schemes by our glaucoma specialist co-author. Tractable methods that include state-dependent cost matrices would be a welcome extension of the model.

7. Conclusions

In this study, we developed a dynamic personalized modeling paradigm for simultaneous monitoring and control of irreversible chronic diseases (e.g., glaucoma). Our model incorporates each patient's past and present readings in a feedback-driven control

model to provide the jointly optimal solution to two critical questions facing clinicians: (i) when to schedule office visits and which suite of tests to perform to monitor for disease progression (exploration); and (ii) what levels of controllable disease risk factors should be targeted to slow the rate of disease progression (exploitation). Kalman filtering methodology is built into our modeling framework to extract noise from the raw measurements and to estimate the disease state in each time period optimally based on imperfect observations. This is key to accurately identifying genuine disease progression from testing artifacts. We developed a multivariate continuous state space model of disease progression and model the state transition and the testing processes as first-order vector difference equations with multivariate Gaussian random noises. For the new objective of minimizing the relative change in state (i.e., disease progression), which is imperative for managing irreversible chronic diseases, we proved the two-way separation of optimal estimation and control. This is a fundamental finding upon which solution tractability depends. To demonstrate the effectiveness of our approach, we harnessed data from two landmark glaucoma randomized clinical trials to parametrize and validate our model. We showed that our Kalman filter-based model has low error in predicting the future disease progression trajectory. Further, we designed our decision support tool to provide a menu of options for the clinician based on how aggressively the doctor wants to manage the patient's disease. For each aggressiveness option, the model provides for each glaucoma patient (i) future disease progression trajectory, (ii) optimal monitoring schedule, and (iii) optimal target IOP. The doctor has the choice to select an appropriate aggressiveness level depending on the patient's life expectancy, severity of glaucoma, and other personal and clinical factors. Our decision support tool can, however, provide insight into such a selection as discussed in Appendix C. Our numerical results demonstrated that following the recommendations of our model not only results in patients with better vision quality over the treatment horizon but also achieves significantly slower glaucoma progression rate, which means patients will keep their sight longer.

Acknowledgments

This research was supported in part by NSF grants CMMI-1161439, CMMI-1552545, and CMMI-1548201, NIH grant R01EY026641, Glaucoma Research Foundation Shaffer Grant, and Seth Bonder Foundation Bonder Scholarship for Applied Operations Research in Health Services. The sponsors had no role in the design, conduct, or writing of this research.

References

- American Academy of Ophthalmology Clinical Practice Guidelines. 2015. *Preferred Practice Pattern Guidelines: Primary Open-Angle Glaucoma*. American Academy of Ophthalmology, San Francisco, CA. Available at <https://www.aao.org/preferred-practice-pattern/primary-open-angle-glaucoma-ppp-2015> (accessed date November 28, 2018).
- Ansley, C. F., R. Kohn. 1982. A geometrical derivation of the fixed interval smoothing algorithm. *Biometrika* **69**(2): 486–487.
- Athans, M. 1972. On the determination of optimal costly measurement strategies for linear stochastic systems. *Automatica* **8**(4): 397–412.
- Ayer, T., O. Alagoz, N. K. Stout. 2012. A POMDP approach to personalize mammography screening decisions. *Oper. Res.* **60**(5): 1019–1034.
- Bansal, R., T. Başar. 1989. Simultaneous design of measurement and control strategies for stochastic systems with feedback. *Automatica* **25**(5): 679–694.
- Baron, S., D. L. Kleinman. 1969. The human as an optimal controller and information processor. *IEEE Trans. Man-Machine Syst.* **10**(1): 9–17.
- Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control*, Vol. 1. Athena Scientific, Belmont, MA.
- Boodhna, T., D. P. Crabb. 2016. More frequent, more costly? Health economic modelling aspects of monitoring glaucoma patients in England. *BMC Health Serv. Res.* **16**(1): 611.
- Bowd, C., I. Lee, M. H. Goldbaum, M. Balasubramanian, F. A. Medeiros, L. M. Zangwill, C. A. Girkin, J. M. Liebmann, R. N. Weinreb. 2012. Predicting glaucomatous progression in glaucoma suspect eyes using relevance vector machine classifiers for combined structural and functional measurements. *Invest. Ophthalmol. Vis. Sci.* **53**(4): 2382–2389.
- Chhatwal, J., O. Alagoz, E. S. Burnside. 2010. Optimal breast biopsy decision-making based on mammographic features and demographic factors. *Oper. Res.* **58**(6): 1577–1591.
- Choplin, N. T., R. P. Edwards. 1995. *Visual Field Testing with the Humphrey Field Analyzer*. Slack Inc, Thorofare, NJ.
- Dempster, A. P., N. M. Laird, D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Series B Methodol.* **39**(1): 1–38.
- Digalakis, V., J. R. Rohlicek, M. Ostendorf. 1993. ML estimation of a stochastic linear system with the EM algorithm and its application to speech recognition. *IEEE Trans. Speech Audio Process.* **1**(4): 431–442.
- Ederer, F., D. Gaasterland, E. Sullivan, A. I. A. Investigators. 1994. The Advanced Glaucoma Intervention Study (AGIS): 1. Study design and methods and baseline characteristics of study patients. *Control. Clin. Trials* **15**(4): 299–325.
- Erenay, F. S., O. Alagoz, A. Said. 2014. Optimizing colonoscopy screening for colorectal cancer prevention and surveillance. *Manuf. Serv. Oper. Manag.* **16**(3): 381–400.
- Gardiner, S. K., D. P. Crabb. 2002. Examination of different pointwise linear regression methods for determining visual field progression. *Invest. Ophthalmol. Vis. Sci.* **43**(5): 1400–1407.
- Ghahramani, Z., G. E. Hinton. 1996. Parameter estimation for linear dynamical systems. Technical report, Department of Computer Science, University of Toronto.
- Goldbaum, M. H., I. Lee, G. Jang, M. Balasubramanian, P. A. Sample, R. N. Weinreb, J. M. Liebmann, C. A. Girkin, D. R. Anderson, L. M. Zangwill, M.-J. Fredette, T.-P. Jung, F. A. Medeiros, C. Bowd. 2012. Progression of patterns (POP): A machine classifier algorithm to identify glaucoma progression in visual fields. *Invest. Ophthalmol. Vis. Sci.* **53**(10): 6557–6567.
- Gupta, V., T. H. Chung, B. Hassibi, R. M. Murray. 2006. On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage. *Automatica* **42**(2): 251–260.
- Harvey, A. C. 1990. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge.
- Heijl, A., P. Buchholz, G. Norrgrén, B. Bengtsson. 2013. Rates of visual field progression in clinical glaucoma care. *Acta Ophthalmol.* **91**(5): 406–412.
- Helm, J. E., M. S. Lavieri, M. P. Van Oyen, J. D. Stein, D. C. Musch. 2015. Dynamic forecasting and control algorithms of glaucoma progression for clinician decision support. *Oper. Res.* **63**(5): 979–999.
- Hyman, L., A. Heijl, M. C. Leske, B. Bengtsson, Z. Yang. 2010. Natural history of intraocular pressure in the Early Manifest Glaucoma Trial: A 6-year follow-up. *Arch. Ophthalmol.* **128**(5): 601–607.
- Jampel, H. D. 1997. Glaucoma care update: Target pressure in glaucoma therapy. *J. Glaucoma* **6**(2): 133–138.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *J. Fluids Eng.* **82**(1): 35–45.
- Kreke, J. E. 2007. Modeling disease management decisions for patients with pneumonia-related sepsis. PhD thesis, University of Pittsburgh.
- Lavieri, M. S., M. L. Puterman, S. Tyldesley, W. J. Morris. 2012. When to treat prostate cancer patients based on their PSA dynamics. *IIE Trans. Healthc. Syst. Eng.* **2**(1): 62–77.
- Lee, P. P., J. G. Walt, J. J. Doyle, S. V. Kotak, S. J. Evans, D. L. Budenz, P. P. Chen, A. L. Coleman, R. M. Feldman, H. D. Jampel, L. Jay Katz, R. P. Mills, J. S. Myers, R. J. Noecker, J. R. Piltz-Seymour, R. R. Ritch, P. N. Schacknow, J. B. Serle, G. L. Trick. 2006. A multicenter, retrospective pilot study of resource use and costs associated with severity of disease in glaucoma. *Arch. Ophthalmol.* **124**(1): 12–19.
- Lee, E., M. S. Lavieri, M. Volk. 2019. Optimal screening for hepatocellular carcinoma: A restless bandit model. *Manuf. Serv. Oper. Manag.* **21**(1): 198–212. <https://doi.org/10.1287/msom.2017.0697>
- Leske, M. C., A. Heijl, M. Hussein, B. Bengtsson, L. Hyman, E. Komaroff. 2003. Factors for glaucoma progression and the effect of treatment: The Early Manifest Glaucoma Trial. *Arch. Ophthalmol.* **121**(1): 48–56.
- Maier, P. C., J. Funk, G. Schwarzer, G. Antes, Y. T. Falck-Ytter. 2005. Treatment of ocular hypertension and open angle glaucoma: Meta-analysis of randomised controlled trials. *BMJ* **331**(7509): 134.
- Mangasarian, O. L., W. N. Street, W. H. Wolberg. 1995. Breast cancer diagnosis and prognosis via linear programming. *Oper. Res.* **43**(4): 570–577.
- Mason, J. E., B. T. Denton, N. D. Shah, S. Smith. 2014. Optimizing the simultaneous management of blood pressure and cholesterol for type 2 diabetes patients. *Eur. J. Oper. Res.* **233**(3): 727–738.
- Mehra, R. K. 1976. Optimization of measurement schedules and sensor designs for linear dynamic systems. *IEEE Trans. Autom. Control* **21**(1): 55–64.
- Meier, L., J. Peschon, R. M. Dressler. 1967. Optimal control of measurement subsystems. *IEEE Trans. Autom. Control* **12**(5): 528–536.
- Musch, D. C., P. R. Lichter, K. E. Guire, C. L. Standardi. 1999. The Collaborative Initial Glaucoma Treatment Study: Study design, methods, and baseline characteristics of enrolled patients. *Ophthalmology* **106**(4): 653–662.
- Regnier, E. D., S. M. Shechter. 2013. State-space size considerations for disease-progression models. *Stat. Med.* **32**(22): 3862–3880.
- Rein, D. B., P. Zhang, K. E. Wirth, P. P. Lee, T. J. Hoerger, N. McCall, R. Klein, J. M. Tielsch, S. Vijan, J. Saaddine. 2006. The

- economic burden of major adult visual disorders in the United States. *Arch. Ophthalmol.* **124**(12): 1754–1760.
- Saaty, T. L., L. G. Vargas. 1998. Diagnosis with dependent symptoms: Bayes theorem and the analytic hierarchy process. *Oper. Res.* **46**(4): 491–502.
- Schell, G. J., M. S. Lavieri, J. E. Helm, X. Liu, D. C. Musch, M. P. Van Oyen, J. D. Stein. 2014. Using filtered forecasting techniques to determine personalized monitoring schedules for patients with open-angle glaucoma. *Ophthalmology* **121**(8): 1539–1546.
- Schell, G. J., G. P. Garcia, M. S. Lavieri, J. Sussman, R. Hayward. 2019. Optimal coinsurance rates for a heterogeneous patient population under inequality and resource constraints. *IIEE Trans.* **51**(1): 74–91. <https://doi.org/10.1080/24725854.2018.1499053>
- Schuman, J. S., C. A. Puliafito, J. G. Fujimoto, J. S. Duker. 2004. *Optical Coherence Tomography of Ocular Diseases*. Slack Inc, Thorofare, NJ.
- Sharpsten, L., J. Fan, S. Demirel, X. Su, C. A. Johnson, R. A. Levine. 2014. Predicting progressive glaucomatous optic neuropathy using random forests based on longitudinally collected standard automated perimetry data. *Invest. Ophthalmol. Vis. Sci.* **55**(13): 5619–5619.
- Shechter, S. M., M. D. Bailey, A. J. Schaefer, M. S. Roberts. 2008. The optimal time to initiate HIV therapy under ordered health states. *Oper. Res.* **56**(1): 20–33.
- Sommer, A., J. M. Tielsch, J. Katz, H. A. Quigley, J. D. Gottsch, J. Javitt, K. Singh. 1991. Relationship between intraocular pressure and primary open angle glaucoma among white and black Americans: The Baltimore Eye Survey. *Arch. Ophthalmol.* **109**(8): 1090–1095.
- Stein, J. D., D. S. Kim, L. M. Niziol, N. Talwar, B. Nan, D. C. Musch, J. E. Richards. 2011. Differences in rates of glaucoma among Asian Americans and other racial groups, and among various Asian ethnic groups. *Ophthalmology* **118**(6): 1031–1037.
- Suen, S. C., M. L. Brandeau, J. D. Goldhaber-Fiebert. 2017. Optimal timing of drug sensitivity testing for patients on first-line tuberculosis treatment. *Health Care Manage. Sci.* **21**(4): 632–646.
- Tham, Y.-C., X. Li, T. Y. Wong, H. A. Quigley, T. Aung, C.-Y. Cheng. 2014. Global prevalence of glaucoma and projections of glaucoma burden through 2040: A systematic review and meta-analysis. *Ophthalmology* **121**(11): 2081–2090.
- Tielsch, J. M., A. Sommer, K. Witt, J. Katz, R. M. Royall. 1990. Blindness and visual impairment in an American urban population: The Baltimore Eye Survey. *Arch. Ophthalmol.* **108**(2): 286–290.
- Vajaranant, T. S., S. Wu, M. Torres, R. Varma. 2012. The changing face of primary open-angle glaucoma in the United States: Demographic and geographic changes from 2011 to 2050. *Am. J. Ophthalmol.* **154**(2): 303–314.
- Van Gestel, A., J. L. Severens, C. A. Webers, H. J. Beckers, N. M. Jansonius, J. S. Schouten. 2010. Modeling complex treatment strategies: Construction and validation of a discrete event simulation model for glaucoma. *Value Health* **13**(4): 358–367.
- Vitus, M. P., W. Zhang, A. Abate, J. Hu, C. J. Tomlin. 2012. On efficient sensor scheduling for linear dynamical systems. *Automatica* **48**(10): 2482–2493.
- Ward, B. W., J. S. Schiller, R. A. Goodman. 2014. Multiple chronic conditions among U.S. adults: A 2012 update. *Prev. Chronic Dis.* **11**: 1–4. <https://doi.org/10.5888/pcd11.130389>
- Weinreb, R. N., P. T. Khaw. 2004. Primary open-angle glaucoma. *Lancet* **363**(9422): 1711–1720.
- Wilensky, J. T., D. K. Gieser, M. L. Dietsche, M. T. Mori, R. Zeimer. 1993. Individual variability in the diurnal intraocular pressure curve. *Ophthalmology* **100**(6): 940–944.
- Witsenhausen, H. S. 1971. Separation of estimation and control for discrete time systems. *Proc. IEEE* **59**(11): 1557–1566.
- Yang, Y., J. D. Goldhaber-Fiebert, L. M. Wein. 2013. Analyzing screening policies for childhood obesity. *Manage. Sci.* **59**(4): 782–795.
- Yousefi, S., M. H. Goldbaum, M. Balasubramanian, F. A. Medeiros, L. M. Zangwill, J. M. Liebmann, C. A. Girkin, R. N. Weinreb, C. Bowd. 2014. Learning from data: Recognizing glaucomatous defect patterns and detecting progression from visual field measurements. *IEEE Trans. Biomed. Eng.* **61**(7): 2112–2124.
- Yousefi, S., M. Balasubramanian, M. H. Goldbaum, F. A. Medeiros, L. M. Zangwill, R. N. Weinreb, J. M. Liebmann, C. A. Girkin, C. Bowd. 2016. Unsupervised Gaussian mixture-model with expectation maximization for detecting glaucomatous progression in standard automated perimetry visual fields. *Transl. Vis. Sci. Technol.* **5**(3): 2–2.
- Zhang, J., B. T. Denton, H. Balasubramanian, N. D. Shah, B. A. Inman. 2012. Optimization of prostate biopsy referral decisions. *Manuf. Serv. Oper. Manag.* **14**(4): 529–547.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix A. Optimization of the Final Period Disease Control Action.

Appendix B. Lemmas.

Appendix C. Insight into Selecting an Aggressiveness Option.

Appendix D. Results on Target IOP and MD Loss Averted.

Appendix E. Monitoring Cost of the Optimal Policies.

Figure S1. Histogram of Target IOPs for CIGTS and AGIS Patients under Different Aggressiveness Policies.

Figure S2. MD Loss Averted for Fast- and Slow-Progressing Patients under the High and Moderate Aggressiveness Policies Compared with the Low Aggressiveness Policy.

Table S1. Comparison of the Performance of Different IOP Control Policies with the Low Aggressiveness Policy.