

Bayesian Sequential Detection With Phase-Distributed Change Time and Nonlinear Penalty—A POMDP Lattice Programming Approach

Vikram Krishnamurthy, *Fellow, IEEE*

Abstract—We show that the optimal decision policy for several types of Bayesian sequential detection problems has a threshold switching curve structure on the space of posterior distributions. This is established by using lattice programming and stochastic orders in a partially observed Markov decision process (POMDP) framework. A stochastic gradient algorithm is presented to estimate the optimal linear approximation to this threshold curve. We illustrate these results by first considering quickest time detection with phase-type distributed change time and a variance stopping penalty. Then it is proved that the threshold switching curve also arises in several other Bayesian decision problems such as quickest transient detection, exponential delay (risk-sensitive) penalties, stopping time problems in social learning, and multi-agent scheduling in a changing world. Using Blackwell dominance, it is shown that for dynamic decision making problems, the optimal decision policy is lower bounded by a myopic policy. Finally, it is shown how the achievable cost of the optimal decision policy varies with change time distribution by imposing a partial order on transition matrices.

Index Terms—Blackwell dominance, exponential delay penalty, lattice programming, monotone likelihood ratio ordering, multi-agent decision making, POMDP, quickest time change detection, social learning, stochastic dominance, transient detection, variance penalty.

I. INTRODUCTION

QUIKKEST time change detection has applications in biomedical signal processing, machine monitoring and finance [5], [38]. There are two general formulations for quickest time detection. In the first formulation, the change point τ^0 is an unknown deterministic time, and the goal is to determine a stopping rule such that a certain worst case delay penalty is minimized subject to a constraint on the false alarm frequency (see, e.g., [31], [36], [53]).

The second formulation, which is the formulation we consider in this paper, is the Bayesian approach. The change time τ^0 is a random variable specified by a prior distribution. Consider a sequence of discrete time random measurements $\{y_k, k \geq 1\}$, such that conditioned on the event $\{\tau^0 = t\}$, $y_k, k \leq t$ are i.i.d. random variables with distribution B_1 and $y_k, k > t$ are

i.i.d. random variables with distribution B_2 . The quickest time detection problem involves detecting the change time τ^0 with minimal cost. That is, at each time $k = 1, 2, \dots$, a decision $u_k \in \{\text{continue, stop and announce change}\}$ needs to be made to optimize a tradeoff between false alarm frequency and linear delay penalty.

In classical Bayesian quickest time detection [38], [44], [52], the change time τ^0 is modelled by a geometric distribution. A geometric distributed change time is realized by a two state discrete-time Markov chain, which we denote as x_k . Therefore, in classical quickest time detection, the optimal decision policy at each time k is a function of a 2-D belief state (posterior probability mass function) $\pi_k(i) = P(x_k = i | y_1, \dots, y_k, u_1, \dots, u_{k-1})$, $i = 1, 2$ with $\pi_k(1) + \pi_k(2) = 1$. So it suffices to consider one element, say $\pi_k(2)$, of this probability mass function. Classical quickest time change detection (see for example [38]) says that there exists a threshold point $\pi^* \in [0, 1]$ such that the optimal decision policy is

$$u_k = \begin{cases} \text{continue} & \text{if } \pi_k(2) \geq \pi^* \\ \text{stop and announce change} & \text{if } \pi_k(2) < \pi^*. \end{cases} \quad (1)$$

As a generalization of the Bayesian framework, [51], [52] consider dependent observations from a finite state Markov chain with transition probability matrix affected by the change point.

A. Main Results and Organization of Paper

This paper considers Bayesian quickest time detection with the following generalizations: phase-type distributed change times, a variance stopping penalty, optimal linear threshold policies, and examples in transient detection, nonlinear delay penalty and stopping time problems in social learning. Our goal is to exploit lattice programming techniques to prove the existence of threshold optimal decision policies for a variety of quickest time detection problems. Below is an overview of these results.

- (i) *Phase-type Distributed Change Times*: We consider quickest time detection when the change time τ^0 has a phase-type (PH) distribution [33]. PH-distributions are used widely in modelling discrete event systems. The optimal detection of a PH-distributed change point is useful since the family of PH-distributions forms a dense subset of the set of all distributions, i.e., for any given distribution function F such that $F(0) = 0$, one can find a sequence of PH-distributions $\{F_n, n \geq 1\}$ to

Manuscript received November 13, 2009; revised April 21, 2011; accepted April 28, 2011. Date of current version October 07, 2011. This work was supported in part by NSERC.

V. Krishnamurthy is with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC V6T 1Z4 Canada (e-mail: vikramk@ece.ubc.ca).

Communicated by F. Hlawatsch, Associate Editor for Detection and Estimation.

Digital Object Identifier 10.1109/TIT.2011.2165152

approximate F uniformly over $[0, \infty)$. As described in [33], a PH-distributed change time can be modelled by a multistate Markov chain with an absorbing state. (For a 2-state Markov chain, the PH-distribution specializes to the geometric distribution). So for quickest time detection with PH-distributed change time, the belief states (Bayesian posterior) lie in a multidimensional simplex of probability mass functions.

- (ii) *Variance Penalty*: The second generalization we consider is a stopping penalty comprising of the false alarm and a variance penalty. The variance penalty is essential in stopping problems where one is interested in ultimately estimating the state x . It penalizes stopping too soon if the uncertainty of the state estimate is large.¹ Since the variance is quadratic in the belief state π , it is not possible to reformulate a variance penalty problem as a standard stopping time problem.

Under what conditions does there exist an optimal threshold decision policy for quickest detection with PH-distributed change time and variance penalty? How can the belief states (in a multidimensional simplex) be ordered and compared with a threshold? Section II formulates the quickest time detection problem as a partially observed Markov decision process (POMDP) and characterizes the optimal decision policy as the solution of a stochastic dynamic programming problem. Using lattice programming [48] our main result (Theorems 1 and 2 in Section III) shows that the optimal decision policy is governed by a threshold switching curve on the space of Bayesian distributions (belief states). This result is useful for several reasons: (a) It provides a multidimensional generalization of (1) to PH-distributed change times. (b) Efficient algorithms can be designed to estimate optimal policies that satisfy this threshold structure. (c) The result holds under set-valued constraints on the change time and observation distribution. So there is an inherent robustness since even if the underlying model parameters are not exactly specified, the threshold structure still holds.

Going from a 2 state Markov chain (geometric distributed change time) to multiple states (PH-distributed change time) introduces substantial complications. For 2 state Markov chains, the posterior distribution can be parametrized by a scalar [as in (1)] and therefore can be completely ordered. However, for more than 2 states, comparing posterior distributions requires using stochastic orders which are *partial orders*. In this paper we use the *monotone likelihood ratio (MLR)* stochastic order [22], [24], [32], [41] to prove our structural results. The MLR order is ideally suited for Bayesian problems since it is preserved under conditional expectations. However, determining the optimal policy is nontrivial since the policy can only be characterized on a partially ordered set (more generally a lattice) within the unit simplex.

¹In [3], a continuous time stochastic control problem is formulated with a quadratic stopping cost, and the existence of the solution to the resulting quasi-variational inequality is proved. However, [3] does not deal with structural results.

We modify the MLR stochastic order to operate on line segments within the unit simplex of posterior distributions. Such line segments form chains (totally ordered subsets of a partially ordered set) and permit us to prove that the optimal decision policy has a threshold structure. Theorem 3 shows that for linear delay and false alarm penalties, the stopping region is convex.

- (iii) *Optimal Linear Threshold*: Having established the existence of a threshold curve, Theorem 4 gives necessary and sufficient conditions for the optimal linear hyperplane approximation to this curve. Then a simulation-based stochastic approximation algorithm (Algorithm 1) is presented to compute this optimal linear hyperplane approximation.

The remainder of the paper illustrates the above structural results in several examples.

- (iv) *Example 2: Quickest Transient Detection*: Section IV considers quickest transient detection. We refer the reader to [39] for a nice description of the quickest transient detection problem and various cost functions. In quickest transient detection, a Markov chain state jumps from a starting state to a transient state at a geometric distributed time, and then jumps out of the state to an absorbing state at another geometric distributed time. We show in Theorem 5 that a similar structural result to quickest time detection holds (i.e., existence of a threshold switching curve and convexity of the stopping region).

- (v) *Example 3: Quickest Time Detection with Exponential Delay Penalty*: In Section V, we generalize the results of Poor [36] to PH-distributed change times. Poor [36] considers a novel variation of the quickest time detection problem where the time delay in the detection is penalized exponentially. By converting the resulting problem into a standard stopping problem [12], it is shown in [36] that the optimal decision policy is a threshold policy under mild conditions.

The exponential delay penalty cost function in [36] is a special case of *risk sensitive stochastic control* with geometric change times. Assuming more general PH-distributed change times, Theorem 6 shows that the optimal detection policy is characterized by a threshold switching curve and the stopping region is convex in the risk-sensitive belief state.

Risk sensitive stochastic control is widely used in mathematical finance, see [6], [14], [17] for comprehensive treatments in discrete and continuous time. In simple terms, quickest time detection seeks to optimize the objective $\mathbb{E}\{J^0\}$ where J^0 is the accumulated sample path cost until some stopping time τ . In risk sensitive control, one seeks to optimize $J = \mathbb{E}\{\exp(\epsilon J^0)\}$. Note that ϵJ can be written as $\epsilon J = \epsilon + \epsilon^2 \mathbb{E}\{J^0\} +$ higher order terms. It therefore follows that for $\epsilon > 0$, the scaled cost ϵJ and hence, J is robust and penalizes heavily large sample path costs due to the presence of second order moments. This is termed a risk-averse control and is of significant importance in mathematical finance, see [6]. Risk sensitive control provides a nice formalization of the exponential penalty delay cost and allows us to

generalize the results in Poor [36] to phase-distributed change times by applying lattice programming.

- (vi) *Example 4 and 5: Stopping Time Problems in Multi-agent Social Learning:* Section IV presents two examples of stopping time problems involving social learning amongst multiple agents. We consider: How do local decisions in social learning affect the global decision in a stopping time problem? Social learning has been used in economics [2], [8], [11], for example to model behavior in financial markets; see also [27], [46]. In social learning, each agent optimizes its local utility selfishly and then broadcasts its action. Subsequent agents then use their private observation together with the actions of previous agents to learn an underlying state.

Our first result (Example 4) deals with a multi-agent Bayesian stopping time problem where agents perform greedy social learning and reveal their local actions to subsequent agents. How can the multi-agent system make a global decision when to stop? Such problems arise in automated decision systems (e.g., sensor networks) where agents make local decisions and reveal these local decisions to subsequent agents. Theorem 7 shows that the optimal decision policy of the stopping time problem has multiple thresholds. This is unusual: if it is optimal to declare state 1 based on a Bayesian belief, it may not be optimal to declare state 1 when the belief about state 1 is stronger. We also give an explicit example of an optimal double threshold policy. The result shows that making global decisions based on local decisions involves nonmonotone policies.

Our second result (Example 5) deals with “constrained optimal” social learning. A key result in social learning is that rational agents eventually herd, that is, they pick the same action irrespective of their private observation and social learning stops. To enhance social learning, Chamley [11] (see also [45] for related work) formulated constrained social learning as a stopping time problem where agents either reveal their observations or they herd (which is equivalent to stopping in a sequential decision problem). When should a multi-agent system make the global decision to stop (herd)? Intuitively, the decision to stop should be made when the state estimate is sufficiently accurate so that revealing private observations is no longer required. Theorem 9 in Section VI shows that the constrained optimal social learning proposed by Chamley [11] has a threshold switching curve in the space of public belief states. Thus, the global decision to stop in [11] can be implemented efficiently in a multi-agent system.

- (vii) *Example 6: Multi-agent Scheduling in a Changing World:* In Section VII we examine: How can the optimal decision policy be bounded in terms of a myopic policy? How does the achievable cost of the optimal policy vary with transition probabilities (and therefore change time distribution)? We answer these two questions in a general setting where optimal decisions need to be made when the underlying state x evolves according to a finite state Markov chain without necessarily having an absorbing

state. The problem is no longer a stopping problem; it is a more general partially-observed stochastic control problem.

To formulate these results, Section VII considers a multi-agent scheduling problem. Using Blackwell dominance, Theorem 10 shows that the optimal policy is lower bounded by a myopic policy. The myopic policy can be computed efficiently and is a rigorous lower bound to the computationally intractable optimal policy. Finally, Theorem 11 examines how the optimal expected cost varies with transition matrix for a stopping time problem (e.g., quickest detection problem) and more general dynamic decision problem in a changing world. The theorem shows that for the underlying Markovian state, the larger the transition matrix (according to an order defined in Section VII), the cheaper the optimal expected cost.

II. PARTIALLY OBSERVED STOCHASTIC CONTROL FORMULATION

In this section we present a partially observed stochastic control formulation that allows us to tackle the various stopping time problems considered in subsequent sections.

A. Stopping-Time Stochastic Control Model

The model comprises of the following ingredients

Absorbing-State Markov Chain and Phase-Type Distribution Change Time: We model the change point τ^0 by a *phase type (PH) distribution*. The family of all PH-distributions forms a dense subset for the set of all distributions [33] and hence can be used to approximate change points with an arbitrary distribution. This is done by constructing a multistate Markov chain as follows: Let $k = 0, 1, \dots$ denote discrete time. Assume the state of nature x_k evolves as a Markov chain on the finite state space

$$\{e_1, \dots, e_X\} \text{ where } e_i \text{ is the } X\text{-dimensional unit vector with 1 in the } i\text{-th position.} \quad (2)$$

Here state “1” (corresponding to e_1) is an absorbing state and denotes the state after the jump change. The states $2, \dots, X$ (corresponding to e_2, \dots, e_X) can be viewed as a single composite state that x resides in before the jump. Denote

$$\mathbb{X} = \{1, 2, \dots, X\}. \quad (3)$$

We assume that the change occurs after at least one measurement. So the initial distribution

$$\begin{aligned} \pi_0 &= (\pi_0(i), i \in \mathbb{X}) \\ \pi_0(i) &= P(x_0 = e_i) \text{ satisfies } \pi_0(1) = 0. \end{aligned} \quad (4)$$

The $X \times X$ transition probability matrix P with elements $P_{ij} = P(x_{k+1} = e_j | x_k = e_i)$ is

$$P = \begin{bmatrix} 1 & 0 \\ \underline{P}_{(X-1) \times 1} & \bar{P}_{(X-1) \times (X-1)} \end{bmatrix}. \quad (5)$$

Let the “change time” τ^0 denote the time at which x_k enters the absorbing state 1, i.e.,

$$\tau^0 = \inf\{k : x_k = 1\}. \tag{6}$$

The distribution of τ^0 is determined by choosing the transition probabilities \underline{P}, \bar{P} in (5). To ensure that τ^0 is finite, assume states $2, 3, \dots, X$ are transient. This is equivalent to \bar{P} satisfying $\sum_{n=1}^{\infty} \bar{P}_{ii}^n < \infty$ for $i = 1, \dots, X - 1$ (where \bar{P}_{ii}^n denotes the (i, i) element of the n th power of matrix \bar{P}). The distribution of τ^0 (which is equivalent to the distribution of the absorption time to state 1) is given by

$$\nu_0 = \pi_0(1), \quad \nu_k = \bar{\pi}'_0 \bar{P}^{k-1} \underline{P}, \quad k \geq 1 \tag{7}$$

where $\bar{\pi}_0 = [\pi_0(2), \dots, \pi_0(X)]'$. The key idea is that by appropriately choosing the pair (π_0, P) and the associated state space dimension X , one can approximate any given discrete distribution on $[0, \infty)$ by the distribution $\{\nu_k, k \geq 0\}$; see [33, pp.240–243]. The event $\{x_k = 1\}$ means the change point has occurred at time k according to PH-distribution (7). Of course, in the special case when x is a 2-state Markov chain, the change time τ^0 is geometrically distributed.

Observation: At time k , the noisy observation $y_k \in \mathbb{Y}$ given state x_k has conditional probability distribution

$$P(y_k \leq \bar{y} | x_k = e_i) = \sum_{y \leq \bar{y}} B_{iy}, \quad i \in \mathbb{X}. \tag{8}$$

Here \sum_y denotes integration with respect to the Lebesgue measure (in which case $\mathbb{Y} \subset \mathbb{R}$ and B_{iy} is the conditional probability density function) or counting measure (in which case \mathbb{Y} is a subset of the integers and B_{iy} is the conditional probability mass function $B_{iy} = P(y_k = y | x_k = e_i)$).

Belief State: At time k , the belief state is the posterior probability mass function of x_k given the observation history y_1, \dots, y_k and past decisions u_1, \dots, u_{k-1} . That is

$$\begin{aligned} \pi_k &= (\pi_k(i), i \in \mathbb{X}) \\ \pi_k(i) &= P(x_k = e_i | y_1, \dots, y_k, u_1, \dots, u_{k-1}) \\ &\text{initialized by } \pi_0. \end{aligned} \tag{9}$$

Equivalently, denote the filtration

$$\mathcal{F}_k = \sigma\text{-algebra generated by } (y_1, \dots, y_k, u_1, \dots, u_{k-1}). \tag{10}$$

Then $\pi_k = \mathbb{E}\{x_k | \mathcal{F}_k\}$. (The notational advantage of choosing unit vectors (2) for the state space is that conditional probabilities and conditional expectations coincide).

The belief state is updated via the Bayesian (Hidden Markov Model) filter

$$\begin{aligned} \pi_k &= T(\pi_{k-1}, y_k), \text{ where} \\ T(\pi, y) &= \frac{B_y P' \pi}{\sigma(\pi, y)} \\ \sigma(\pi, y) &= \mathbf{1}'_X B_y P' \pi \\ B_y &= \text{diag}(P(y | x = e_i), i \in \mathbb{X}). \end{aligned} \tag{11}$$

Here $\mathbf{1}_X$ denotes the X -dimensional vector of ones. The belief state π in (1) is an X -dimensional probability vector. It belongs to the $X - 1$ -dimensional unit-simplex denoted as

$$\Pi(X) \triangleq \{\pi \in \mathbb{R}^X : \mathbf{1}'_X \pi = 1, 0 \leq \pi(i) \leq 1 \text{ for all } i \in \mathbb{X}\}. \tag{12}$$

For example, $\Pi(2)$ is a 1-dimensional simplex (unit line segment), $\Pi(3)$ is a 2-D simplex (equilateral triangle); $\Pi(4)$ is a tetrahedron, etc. The states e_1, e_2, \dots, e_X of the Markov chain x are the vertices of $\Pi(X)$.

Sequential Decision and Costs: At each time k , a decision u_k is taken where

$$\begin{aligned} u_k &= \mu(\pi_k) \in \mathcal{U} \\ &= \{1 \text{ (announce change and stop)}, 2 \text{ (continue)}\}. \end{aligned} \tag{13}$$

In (13), the policy μ belongs to the class of stationary decision policies denoted $\boldsymbol{\mu}$.

(i) *Cost of announcing change and stopping:* If decision $u_k = 1$ is chosen, then the problem terminates. If $u_k = 1$ is chosen before the change point τ^0 , then a false alarm and variance penalty is paid. If $u_k = 1$ is chosen at or after the change point τ^0 , then only a variance penalty is paid. Below we formulate these costs.

Let $g = (g_1, \dots, g_X)'$ specify the physical state levels associated with states $1, 2, \dots, X$ of the Markov chain x . The *variance penalty* is

$$\begin{aligned} \mathbb{E}\{ |(x_k - \pi_k)' g|^2 | \mathcal{F}_k \} &= G' \pi_k(i) - (g' \pi_k)^2 \\ \text{where } G_i &= g_i^2 \text{ and } G = (G_1, G_2, \dots, G_X). \end{aligned} \tag{14}$$

This conditional variance penalizes choosing the stop action if the uncertainty in the state estimate is large, see also [3].

Next, the false alarm event $\cup_{i \geq 2} \{x_k = e_i\} \cap \{u_k = 1\} = \{x_k \neq e_1\} \cap \{u_k = 1\}$ represents the event that a change is announced before the change happens at time τ^0 . To evaluate the *false alarm penalty*, let $f_i I(x_k = e_i, u_k = 1)$ denote the cost of a false alarm in state $e_i, i \in \mathbb{X}$, where $f_i \geq 0$. Of course, $f_1 = 0$ since a false alarm is only incurred if the stop action is picked in states $2, \dots, X$. The expected false alarm penalty is

$$\begin{aligned} \sum_{i \in \mathbb{X}} f_i \mathbb{E}\{ I(x_k = e_i, u_k = 1) | \mathcal{F}_k \} &= \mathbf{f}' \pi_k I(u_k = 1) \\ \text{where } \mathbf{f} &= (f_1, \dots, f_X)', f_1 = 0. \end{aligned} \tag{15}$$

The false alarm vector \mathbf{f} is chosen with increasing elements so that states further from state 1 incur larger penalties.

Then with α, β denoting non-negative constants that weight the relative importance of these costs, the expected stopping cost at time k is

$$\bar{C}(\pi_k, u_k = 1) = \alpha(G' \pi_k - (g' \pi_k)^2) + \beta \mathbf{f}' \pi_k. \tag{16}$$

One can also view α informally as a Lagrange multiplier in a stopping time problem that seeks to minimize a cumu-

lative cost (as in (20) below) subject to a variance stopping constraint.

(ii) *Delay cost of continuing*: We allow two possible choices for the delay costs:

- (a) If decision $u_k = 2$ is taken then $\{x_{k+1} = e_1, u_k = 2\}$ is the event that no change is declared at time k even though the state has changed at time $k + 1$. So with d denoting a non-negative constant, $dI(x_{k+1} = e_1, u_k = 2)$ depicts a *delay cost*. The expected delay cost for decision $u_k = 2$ is

$$\begin{aligned} \bar{C}(\pi_k, u_k = 2) &= d \mathbb{E}\{I(x_{k+1} = e_1, u_k = 2) | \mathcal{F}_k\} \\ &= de'_1 P' \pi_k. \end{aligned} \quad (17)$$

The above cost is motivated by applications (e.g., sensor networks) where if the decision maker chooses $u_k = 2$, then it needs to gather observation y_{k+1} thereby incurring an additional operational cost denoted as c . Strictly speaking, $\bar{C}(\pi, 2) = de'_1 P' \pi + c$. Without loss of generality set the constant c to zero, as it does not affect our structural results. The penalty $dI(x_{k+1} = e_1, u_k = 2)$ gives incentive for the decision maker to predict the state x_{k+1} .

- (b) Instead of the above, the more “classical” formulation is that a delay cost is incurred when the event $\{x_k = e_1, u_k = 2\}$ occurs. Then the expected delay cost is

$$\begin{aligned} \bar{C}(\pi_k, u_k = 2) &= d \mathbb{E}\{I(x_k = e_1, u_k = 2) | \mathcal{F}_k\} \\ &= de'_1 \pi_k. \end{aligned} \quad (18)$$

Remark: Due to the variance penalty, the cost $\bar{C}(\pi, 1)$ in (16) is quadratic in the belief state π . Therefore, the formulation cannot be reduced to a standard stopping problem with linear costs in the belief state.

B. Quickest Time Detection Objective

Let (Ω, \mathcal{F}) be the underlying measurable space where $\Omega = (\mathbb{X} \times \mathcal{U} \times \mathbb{Y})^\infty$ is the product space, which is endowed with the product topology and \mathcal{F} is the corresponding product sigma-algebra. For any $\pi_0 \in \Pi(X)$, and policy $\mu \in \boldsymbol{\mu}$, there exists a (unique) probability measure $\mathbb{P}_{\pi_0}^\mu$ on (Ω, \mathcal{F}) , see [16] for details. Let $\mathbb{E}_{\pi_0}^\mu$ denote the expectation with respect to the measure $\mathbb{P}_{\pi_0}^\mu$.²

Let τ denote a stopping time adapted to the sequence of σ -algebras $\mathcal{F}_k, k \geq 1$, defined in (10). That is, with u_k determined by decision policy (13),

$$\tau = \{\inf k : u_k = 1\}. \quad (19)$$

²The formulation on $\Omega = (\mathbb{X} \times \mathcal{U} \times \mathbb{Y})^\infty$ is as follows, see [16]. Augment $\Pi(X)$ to include the fictitious stopping state e_{X+1} which is cost free, i.e., $C(e_{X+1}, u) = 0$ for all $u \in \mathcal{U}$. When decision $u_k = 1$ is chosen, the belief state π_{k+1} transitions to e_{X+1} and remains there indefinitely. Then (20) is equivalent to $J_\mu(\pi_0) = \mathbb{E}_{\pi_0}^\mu \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \bar{C}(\pi_k, u_k = 2) + \rho^{\tau-1} \bar{C}(\pi_\tau, u_\tau = 1) + \sum_{k=\tau+1}^{\infty} \rho^{k-1} \bar{C}(e_{X+1}, u_k) \right\}$, where the last summation is zero.

For each initial distribution $\pi_0 \in \Pi(X)$, and policy μ , the following cost is associated:

$$J_\mu(\pi_0) = \mathbb{E}_{\pi_0}^\mu \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \bar{C}(\pi_k, u_k = 2) + \rho^{\tau-1} \bar{C}(\pi_\tau, u_\tau = 1) \right\}. \quad (20)$$

Here $\rho \in [0, 1]$ denotes an economic discount factor. Since $\bar{C}(\pi, 1), \bar{C}(\pi, 2)$ are non-negative and bounded for all $\pi \in \Pi(X)$, stopping is guaranteed in finite time, i.e., τ is finite with probability 1 for any $\rho \in [0, 1]$ (including $\rho = 1$).

Remark: For the special case $\alpha = 0, \mathbb{X} = \{1, 2\}$ (i.e., geometric distributed change time), $\mathbf{f} = e_2$, and delay cost (17), it is easily shown that

$$J_\mu(\pi_0) = d \mathbb{E}_{\pi_0}^\mu \{(\tau - \tau^0)^+\} + \beta \mathbb{P}_{\pi_0}^\mu(\tau < \tau^0) + dP_{21} \mathbb{E}_{\pi_0}^\mu \{(\tau^0 - 1)I(\tau^0 < \tau)\} \quad (21)$$

(where τ^0 is defined in (6) and τ is defined in (19)). For the delay cost (18), the cost function assumes the classical Kolmogorov-Shiryayev criterion for detection of disorder [43], namely

$$J_\mu(\pi_0) = d \mathbb{E}_{\pi_0}^\mu \{(\tau - \tau^0)^+\} + \beta \mathbb{P}_{\pi_0}^\mu(\tau < \tau^0). \quad (22)$$

■

The goal is to determine the change time τ^0 defined in (6) with minimal cost, that is, compute the optimal policy $\mu^* \in \boldsymbol{\mu}$ to minimize (20), i.e., $J_{\mu^*}(\pi_0) = \inf_{\mu \in \boldsymbol{\mu}} J_\mu(\pi_0)$. The existence of an optimal stationary policy μ^* follows from [7, Prop.1.3, Chapter 3]. Considering the above cost (20), the optimal stationary policy $\mu^* : \Pi(X) \rightarrow \{1, 2\}$ and associated value function $\bar{V}(\pi)$ are the solution of the following “Bellman’s dynamic programming equation”

$$\begin{aligned} \mu^*(\pi) &= \arg \min \left\{ \bar{C}(\pi, 1), \bar{C}(\pi, 2) + \rho \sum_{y \in \mathbb{Y}} \bar{V}(T(\pi, y)) \sigma(\pi, y) \right\} \\ J_{\mu^*}(\pi) &= \bar{V}(\pi) \\ &= \min \left\{ \bar{C}(\pi, 1), \bar{C}(\pi, 2) + \rho \sum_{y \in \mathbb{Y}} \bar{V}(T(\pi, y)) \sigma(\pi, y) \right\}. \end{aligned} \quad (23)$$

Before proceeding, we rewrite the above in a form that is more amenable³ for analysis. Define

$$\begin{aligned} V(\pi) &= \bar{V}(\pi) - (\alpha + \beta) \mathbf{f}' \pi, \quad C(\pi, 1) \\ &= \alpha(G' \pi - (g' \pi)^2) - \alpha \mathbf{f}' \pi \\ C(\pi, 2) &= \bar{C}(\pi, 2) - (\alpha + \beta) \mathbf{f}' \pi + \rho(\alpha + \beta) \mathbf{f}' P' \pi. \end{aligned} \quad (24)$$

³The reason for changing coordinates from $\bar{C}(\pi, 1), \bar{C}(\pi, 2)$ to $C(\pi, 1), C(\pi, 2)$ is to make our analysis compatible with existing results in quickest time detection. To ensure this compatibility, we need $C(\pi, 1)$ to be decreasing with respect to the MLR order (see Appendix) when $\alpha = 0$. As shown in the proof of Theorem 1 in the Appendix, $\bar{C}(\pi, 1)$ is MLR decreasing. In comparison $\bar{C}(\pi, 1)$ is not MLR decreasing. Of course, the stopping set \mathcal{R}_1 (see (6)) and optimal policy μ^* are invariant to the choice of coordinates. This idea of changing coordinates is described in [13], albeit for the simpler fully observed Markov decision process case.

Then clearly $V(\pi)$ satisfies Bellman's dynamic programming equation

$$\begin{aligned} \mu^*(\pi) &= \arg \min_{u \in \mathcal{U}} Q(\pi, u) \\ J_{\mu^*}(\pi) &= V(\pi) = \min_{u \in \{1,2\}} Q(\pi, u) \\ \text{where } Q(\pi, 2) &= C(\pi, 2) + \rho \sum_{y \in \mathcal{Y}} V(T(\pi, y)) \sigma(\pi, y) \\ Q(\pi, 1) &= C(\pi, 1). \end{aligned} \quad (25)$$

Thus, the goal is to determine the optimal stopping set

$$\begin{aligned} \mathcal{R}_1 &= \{\pi \in \Pi(X) : \mu^*(\pi) = 1\} \\ &= \left\{ \pi \in \Pi(X) : C(\pi, 1) < C(\pi, 2) + \rho \sum_{y \in \mathcal{Y}} V(T(\pi, y)) \sigma(\pi, y) \right\} \\ &= \left\{ \pi \in \Pi(X) : \bar{C}(\pi, 1) < \bar{C}(\pi, 2) + \rho \sum_{y \in \mathcal{Y}} \bar{V}(T(\pi, y)) \sigma(\pi, y) \right\}. \end{aligned} \quad (26)$$

In Section III-B, sufficient conditions are given to ensure that \mathcal{R}_1 is nonempty.

Value Iteration Algorithm and Methodology: We comment briefly here on our analysis methodology which is detailed in Section III. Let $k = 1, 2, \dots$, denote iteration number (the fact that we used k previously to denote time should not result in confusion). The value iteration algorithm is a fixed point iteration of Bellman's equation and proceeds as follows:

$$\begin{aligned} V_0(\pi) &= -(\alpha + \beta)\mathbf{f}'\pi, \quad V_{k+1}(\pi) \\ &= \min_{u \in \{1,2\}} Q_{k+1}(\pi, u), \quad \mu_{k+1}^*(\pi) \\ &= \arg \min_{u \in \{1,2\}} Q_{k+1}(\pi, u) \\ \text{where } Q_{k+1}(\pi, 2) &= C(\pi, 2) + \rho \sum_{y \in \mathcal{Y}} V_k(T(\pi, y)) \\ &\quad \sigma(\pi, y), \quad Q_{k+1}(\pi, 1) = C(\pi, 1). \end{aligned} \quad (27)$$

Let $\mathcal{B}(X)$ denote the set of bounded real-valued functions on $\Pi(X)$. Then for any V and $\tilde{V} \in \mathcal{B}(X)$, define the sup-norm metric $\sup \|V(\pi) - \tilde{V}(\pi)\|$, $\pi \in \Pi(X)$. Then $\mathcal{B}(X)$ is a Banach space. The value iteration algorithm (27) will generate a sequence of value functions $\{V_k\} \subset \mathcal{B}(X)$ that will converge uniformly (sup-norm metric) as $k \rightarrow \infty$ to $V(\pi) \in \mathcal{B}(X)$, the optimal value function of Bellman's equation. However, since the belief state space $\Pi(X)$ is an uncountable set, the value iteration algorithm (27) do not translate into practical solution methodologies as $V_k(\pi)$ needs to be evaluated at each $\pi \in \Pi(X)$, an uncountable set. Indeed, due to the nonlinearity in the belief states, the formulation is more complex than a partially observed Markov decision process which is known to be PSPACE hard [34]. Although value iteration is not useful from a computational point of view, in Section III, we exploit the structure of the value iteration recursion (25), (27) to prove that \mathcal{R}_1 is characterized by a threshold switching curve. We then exploit this structure to devise polynomial complexity algorithms for

approximating the optimal policy μ^* and thus determining the stopping set \mathcal{R}_1 .

Remark: Computational algorithms based on value iteration such as Sondik's algorithm, Monahan's algorithm, Cheng's algorithm, Witness algorithm (see [9], [10] for a tutorial description) and Lovejoy's suboptimal algorithms [29] solve POMDPs with linear costs (i.e., $\alpha = 0$) over finite horizons. These algorithms require finite observation spaces and are computationally intractable except for small \mathcal{X} and \mathcal{Y} . They are not applicable directly to stopping problems considered in this paper since we consider nonlinear penalty costs, possibly continuous observation space \mathcal{Y} (Examples 1, 2, and 3), and problems where the observation probabilities depend on the belief state (social learning in Examples 4 and 5).

III. EXAMPLE 1: QUICKEST TIME DETECTION WITH PH-DISTRIBUTED CHANGE TIME AND VARIANCE PENALTY

This section considers quickest time detection with PH-distributed change time and variance penalty. Section III-A below gives the main results of this paper, namely the optimal decision policy is characterized by a threshold curve. Sections III-B and III-C discuss the implications and main assumptions. Section III-D then parametrizes the optimal linear approximation to this threshold curve. Finally, Section III-E gives a stochastic optimization algorithm (Algorithm 1) to compute this optimal linear approximation.

The quickest time detection problem is a special case of the stochastic control problem formulated in Section II. The states $2, 3, \dots, X$ are fictitious and are defined to generate the change time τ^0 with PH-distribution (7). So states $2, 3, \dots, X$ are indistinguishable in terms of the observation y . That is, the observation probabilities B in (8) and Markov chain state levels g in (14) satisfy

$$\begin{aligned} B_{2y} &= B_{3y} = \dots = B_{Xy} \text{ for all } y \in \mathcal{Y} \\ g_1 &= 0, \quad g_2 = g_3 = \dots = g_X = 1. \end{aligned} \quad (28)$$

The above choice of $g = \mathbf{1}_X - e_1$ is without loss of generality since the variance penalty (14) is translation invariant with respect to $c\mathbf{1}_X$ for any c .

Notation: Notation and definitions regarding stochastic orders, lattice programming, the poset $[\Pi(X), \geq_r]$ and submodularity are given in Appendix A. Below \geq_r denotes the monotone likelihood ratio order, \geq_{L_X} denotes the likelihood ratio order on lines $\mathcal{L}(e_X, \bar{\pi})$, \geq_{L_1} denotes the likelihood ratio order on lines $\mathcal{L}(e_1, \bar{\pi})$, and \geq_s denotes first order stochastic dominance.

A. Main Result: Existence of Decision Curve Policy for Quickest Time Detection

This section gives three main results: The optimal policy for quickest detection with PH-distributed change time and variance penalty is characterized by a threshold curve (Theorems 1 and 2). Also for $\alpha = 0$, it is shown that the stopping set \mathcal{R}_1 is convex (Theorem 3).

Quickest Detection With Delay Penalty (17): For the stopping cost $\bar{C}(\pi, 1)$ in (16), choose $\mathbf{f} = [0, 1, \dots, 1]' = \mathbf{1}_X - e_1$. This weighs the states $2, \dots, X$ equally in the false alarm

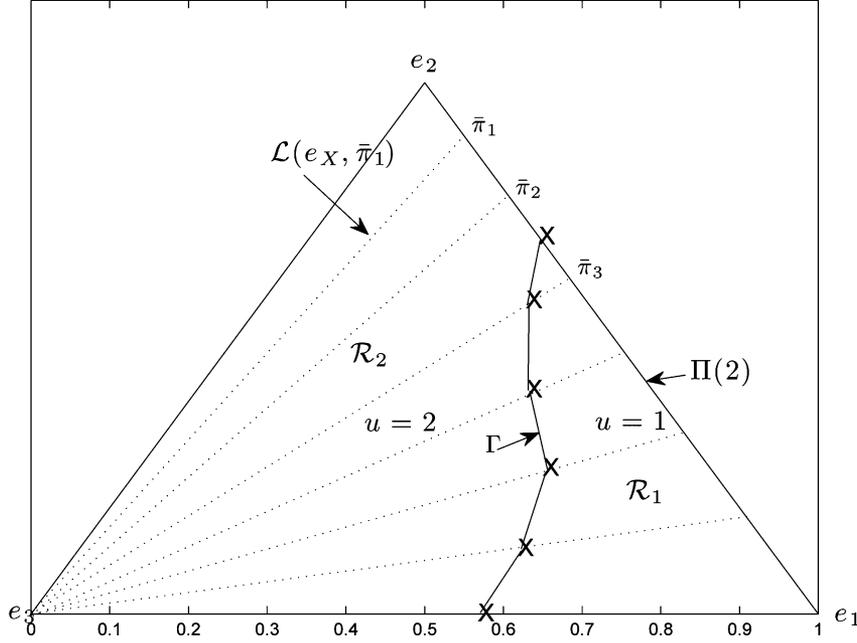


Fig. 1. Illustration of threshold switching decision curve Γ . Here $X = 3$ and hence $\Pi(X)$ is an equilateral triangle. Theorem 1 shows that the stopping region \mathcal{R}_1 is a connected and $e_1 \in \mathcal{R}_1$. Also \mathcal{R}_2 is connected. The lines segments $\mathcal{L}(e_X, \bar{\pi}_1)$ connecting the sub-simplex $\Pi(2)$ to e_3 are defined in (80). Theorem 1 says that the threshold curve Γ can intersect each line $\mathcal{L}(e_X, \bar{\pi})$ only once (and similarly intersect each line $\mathcal{L}(e_1, \bar{\pi})$ only once). In the special case $\alpha = 0$, Theorem 3 says that \mathcal{R}_1 is a convex set (not shown in this figure).

penalty. With assumption (28), the variance penalty (14) becomes $\alpha(e'_1\pi - (e'_1\pi)^2)$. The delay cost $\bar{C}(\pi, 2)$ is chosen as (17). To summarize (24)–(26) hold with

$$\begin{aligned}\bar{C}(\pi, 1) &= \alpha(e'_1\pi - (e'_1\pi)^2) + \beta(1 - e'_1\pi) \\ \bar{C}(\pi, 2) &= de'_1P'\pi.\end{aligned}\quad (29)$$

Theorem 1 below is our main result on the structure of the optimal decision policy $\mu^*(\pi)$. It is based on the following assumptions (discussed in Section III-C).

(A1-Ex1) $d \geq \rho(\alpha + \beta)$

(A2) The observation distribution B_{xy} in (8) is TP2 in (x, y) (see Defn.4(iii) in Appendix A. Equivalently, from (28), $B_{2y} \geq_r B_{1y}$).

(A3) The transition matrix P in (5) is TP2, i.e., all its second order minors are non-negative.

(S-Ex1) $(d - \rho(\alpha + \beta))(1 - P_{21}) \geq \alpha - \beta$

(A1-Ex1) and (S-Ex1) are constraints on the delay and stopping cost functions (that the decision maker can design), while (A2) and (A3) are assumptions on the underlying observation (8) and PH-distribution (6).

Theorem 1 (Switching Curve Optimal Policy): Consider the quickest time detection problem (20) with costs defined in (29) and PH-distributed change time τ^0 defined in (6). Then for $\pi \in \Pi(X)$, under (A1-Ex1), (A2), (A3), (S-Ex1), there exists an optimal policy $\mu^*(\pi)$ that is \geq_{L_X} increasing on lines $\mathcal{L}(e_X, \bar{\pi})$ and \geq_{L_1} increasing on lines $\mathcal{L}(e_1, \bar{\pi})$. As a consequence:

- (i) The stopping set \mathcal{R}_1 defined in (6) has the following structure: There exists a threshold switching curve Γ that partitions belief state space $\Pi(X)$ into two individually connected regions $\mathcal{R}_1, \mathcal{R}_2$, such that the optimal policy is

$$\mu^*(\pi) = \begin{cases} \text{continue} = 2 & \text{if } \pi \in \mathcal{R}_2 \\ \text{stop} = 1 & \text{if } \pi \in \mathcal{R}_1. \end{cases}\quad (30)$$

(A set is connected if it cannot be expressed as the union of two disjoint nonempty closed sets [42]). The threshold curve Γ intersects each line $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$ at most once.

- (ii) There exists an $i^* \in \{0, \dots, X\}$, such that $e_1, e_2, \dots, e_{i^*} \in \mathcal{R}_1$ and $e_{i^*+1}, \dots, e_X \in \mathcal{R}_2$.
- (iii) For geometric distributed change time τ^0 , there exists a unique threshold point $\pi^*(2)$ such that (1) holds. (Note (A3) holds trivially in this case). ■

Theorem 1 is proved in Appendix C and uses meta Theorem 13 in Appendix B as a key step. The intuition behind Theorem 1 is discussed in Sections III-B and III-C below. Fig. 1 gives a pictorial illustration. Note that if $\alpha = 0$, then (S-Ex1) holds trivially if (A1-Ex1) holds.

Quickest Detection With Delay Penalty (18): Next consider the “classical” delay cost $\bar{C}(\pi, 2)$ in (18) and stopping cost $\bar{C}(\pi, 1)$ in (16) with g in (28). Then (24)–(26) hold with

$$\begin{aligned}\bar{C}(\pi, 1) &= \alpha(e'_1\pi - (e'_1\pi)^2) + \beta\mathbf{f}'\pi_k \\ \bar{C}(\pi, 2) &= de'_1\pi.\end{aligned}\quad (31)$$

Below we show that Theorem 1 continues to hold, if the decision maker designs the false alarm vector \mathbf{f} to satisfy the following linear constraints:

- (i) (AS-Ex1) (i) $f_i \geq \max\{1, \rho\frac{\alpha+\beta}{\beta}\mathbf{f}'P'e_i + \frac{\alpha-d}{\beta}\}$, $i \geq 2$.
- (ii) $f_j - f_i \geq \rho\mathbf{f}'P'(e_j - e_i)$, $j \geq i$, $i \in \{2, \dots, X-2\}$
- (iii) $f_X - f_i \geq \frac{\rho(\alpha+\beta)}{\beta}\mathbf{f}'P'(e_X - e_i)$, $i \in \{2, \dots, X-1\}$.

Feasible choices of \mathbf{f} are easily obtained by a linear programming solver.

Theorem 2: Consider the quickest detection problem with delay and stopping costs in (31). Then under (AS-Ex1), (A2), (A3), Theorem 1 holds. ■

Convexity of Stopping Region When $\alpha = 0$: Finally, we present the following result for the case $\alpha = 0$, i.e., no variance penalty.

Theorem 3: For arbitrary PH-distributed change time τ^0 , and no variance penalty ($\alpha = 0$), the stopping region \mathcal{R}_1 is a convex subset of $\Pi(X)$. ■

The proof of Theorem 3 is in Appendix D; it was proved in [28] in a POMDP setting. Theorem 3 says that as long as costs $\bar{C}(\pi, 1)$, $\bar{C}(\pi, 2)$ are linear in π (i.e., no variance penalty), then the stopping set is convex for any size X (i.e., arbitrary PH-distribution); no assumptions are required on the transition matrix P or observation likelihood matrix B . However, even though \mathcal{R}_1 is convex (and therefore connected), Theorem 3 does not guarantee that \mathcal{R}_2 is connected. As described in Section III-C, Theorem 1 and Theorem 2 go much further than Theorem 3 in characterizing \mathcal{R}_1 and \mathcal{R}_2 , even for the case $\alpha = 0$.

B. Discussion of Theorems 1, 2

Theorems 1 and 2 imply that since the optimal decision policy is characterized by a threshold curve, quickest time detection for PH-distributed change times and variance penalty can be implemented efficiently, Section III-D. Without this result, the stopping set \mathcal{R}_1 is not necessarily a connected region as will be shown in Section VIII.

Geometric Distributed Change Time: When the change time τ^0 is geometrically distributed, since the state space $\mathbb{X} = \{1, 2\}$, $\Pi(X)$ is a 1-D simplex. Then the stochastic orders \geq_r , \geq_{L_i} and \geq_s defined in Appendix A coincide, and become total orders. Also (A3) holds automatically for this case. Below we discuss the cases of $\alpha \neq 0$ and $\alpha = 0$.

- (i) $\alpha \neq 0$: For geometric distributed change time, Theorems 1 and 2 say that the classical threshold policy depicted in (1) continues to hold when a nonlinear variance penalty is considered. For example, consider Theorem 2 with nonzero α , delay in (31) and false alarm vector $\mathbf{f}' = e_2$. So the false alarm cost is $\mathbf{f}'\pi = 1 - e_1'\pi$ (which is identical to (29)). One can view this as the Kolmogorov-Shiryayev criterion (22) with an additional variance penalty. Theorem 2 holds under the conditions (AS-Ex1) and (A2). Here (AS-Ex1) equivalent to the constraint that $\alpha \leq \frac{d+\beta(1-\rho P_{22})}{1+\rho P_{22}}$. [Choose $f_1 = 0$, $f_2 = 1$ in (AS-Ex1)(i)]. So for $\alpha \leq d/2$, (AS-Ex1) always holds.
- (ii) $\alpha = 0$: For quickest time detection with geometric distributed change time and no variance penalty ($\alpha = 0$), the well known existence of a threshold point (e.g., for the Kolmogorov-Shiryayev criterion (22)) follows trivially from Theorem 3. Since $\Pi(X)$ is a 1-D simplex, convexity of stopping set \mathcal{R}_1 (Theorem 3) implies that there is a threshold point π^* that satisfies (1).

Avoiding Trivial Cases: To ensure the stopping set \mathcal{R}_1 contains state e_1 , assume $\bar{C}(e_1, 1) < \bar{C}(e_1, 2)$. From (29) or (31) this is equivalent to $d > 0$. The strict inequality also implies that $\mathcal{R}_1 \cap \Pi^o(X)$ is nonempty, where $\Pi^o(X)$ denote the interior of the simplex $\Pi(X)$.

For the detection problem to be nontrivial, we want $\bar{C}(e_i, 1) > \bar{C}(e_i, 2)$ for $i \geq 2$, otherwise it is always optimal to stop at time 1. For the case of Theorem 1, from (29), a sufficient

condition is that $\beta > dP_{i1}$, $i = 2, \dots, X$. Since the transition matrix P is TP2, it follows⁴ that $P_{21} \geq P_{31} \geq \dots \geq P_{X1}$. Therefore, it is sufficient for the decision maker to choose constants β and d such that $\beta > dP_{21}$. For Theorem 2, from (31), $\bar{C}(e_i, 1) > \bar{C}(e_i, 2)$ always holds for $\beta > 0$ and $f_i > 0$, $i \geq 2$.

Nondegenerate Threshold Curve: Let $\Pi^o(X)$ and $\Pi^b(X)$, respectively, denote the interior and boundary of the simplex $\Pi(X)$. Determining the threshold switching curve Γ in Theorem 1 requires determining $\Gamma \cap \Pi^o(X)$ (portion of curve that lies in the interior of the simplex) and $\Gamma \cap \Pi^b(X)$ (portion of curve that lies on the boundary of the simplex). Since $\Pi^b(X)$ comprises of sub-simplices, to determine $\Gamma \cap \Pi^b(X)$ one would need to search for the threshold curve Γ within these sub-simplices. While conceptually straightforward, we can eliminate this search by ensuring that the belief state π always lives in the interior $\Pi^o(X)$ of the simplex. The following lemma gives sufficient conditions for the sequence of belief states π_k over time to lie in $\Pi^o(X)$.

Lemma 1: Suppose each column of transition matrix P has at least one nonzero element, and the observation likelihoods satisfy $B_{1y} \neq 0$ for $y \in \mathbb{Y}$. Then for initial belief state π_0 satisfying (4) with $\pi_0(i) \neq 0$, $i > 1$, subsequent belief states π_k lie in $\Pi^o(X) \cup \{e_1\}$ for all time $k \geq 1$.

The proof follows straightforwardly from the belief state update (11). Since the sequence of belief states π_k , $k \geq 1$, lives in $\Pi^o(X)$, one only needs to compute the threshold curve inside the simplex, i.e., $\Gamma \cap \Pi^o(X)$. Recall from the previous remark that $\mathcal{R}_1 \cap \Pi^o(X)$ is nonempty and consequently $\Gamma \cap \Pi^o(X)$ is nonempty.

C. Assumptions and Proofs of Theorem 1 and Theorem 2

Below we discuss the main assumptions of Theorem 1 and Theorem 2 in Section III-A, then outline the structure of the proof, and finally give intuitive examples that illustrate the structure of stopping set \mathcal{R}_1 .

Discussion of Assumptions: Recall (A1-Ex1) and (S-Ex1) are design constraints the decision maker uses to choose the stopping and delay costs. In contrast, (A2) and (A3) are assumptions on the underlying stochastic model.

As described in the Appendix B, (A1-Ex1) is sufficient for $C(\pi, 2)$ to be \geq_r decreasing. We also require $C(\pi, 1)$ in (29) to be \geq_r decreasing, but this holds trivially in our setup.

(S-Ex1) is a submodularity condition, see Defn.5 in Appendix. We refer to [1], [48] for extensive treatments of lattice programming and submodularity. The key idea is that if $Q(\pi, u)$ is submodular on the partially ordered set $[\Pi(X), \geq_r]$, then the optimal policy $\mu^*(\pi) = \arg \min_u Q(\pi, u)$ is monotone increasing with respect to \geq_r .

In our setting, submodularity of $Q(\pi, u)$ in (25) is equivalent to showing that $Q(\pi, 2) - Q(\pi, 1)$ is decreasing with respect to π (in terms of the MLR order \geq_r , see discussion of structure of proof given below). Since by (A1-Ex1), $C(\pi, 1)$ and $C(\pi, 2)$

⁴Proof: We prove the contrapositive, that is, $P_{i1} < P_{i+1,1}$ implies P is not TP2. Recall from (A3-Ex1), TP2 means that $P_{i1}P_{i+1,j} \geq P_{i+1,1}P_{ij}$ for all j . So assuming $P_{i1} < P_{i+1,1}$, to show that P is not TP2, we need to show that there is at least one j such that $P_{i+1,j} < P_{ij}$. But $P_{i1} < P_{i+1,1}$ implies $\sum_{k \neq 1} P_{i+1,k} < \sum_{k \neq 1} P_{ik}$, which in turn implies that at least for one j , $P_{i+1,j} < P_{ij}$.

are MLR decreasing in π , it will be proved that $V(\pi)$ is MLR decreasing in π providing (A2) and (A3) hold. Since the sum of submodular functions is submodular, establishing submodularity of $Q(\pi, u)$ in (25) is equivalent to establishing submodularity of $C(\pi, u)$. Clearly if $\alpha = 0$, then from (29), $C(\pi, 1) = 0$ is independent of π . So if $\alpha = 0$ then the submodular condition (S-Ex1) holds trivially since $C(\pi, 2)$ is decreasing in π via (A1-Ex1). So for quickest time detection with PH-distributed change time and no variance penalty, submodularity holds by construction. Note that when the variance penalty is included, (S-Ex1) always holds if $\alpha < \beta$.

AS-Ex1 in Theorem 2 ensures that $C(\pi, 1)$ and $C(\pi, 2)$ in (24) for the modified delay cost in (31) are monotone decreasing and that $C(\pi, u)$ is submodular. It is analogous to (A-Ex1) and (S-Ex1).

(A2) is required for preserving the MLR ordering with respect to observation y of the Bayesian filter update $T(\pi, y)$ —this is a key step in showing $V(\pi)$ is MLR decreasing in π . Theorem 13(ii) in the Appendix states that $T(\pi, y)$ is MLR increasing in y , iff (A2) holds.

(A2) is satisfied by numerous continuous and discrete distributions, see a classical detection theory book such as [37]. Examples include Gaussians, Exponential, Binomial, Poisson, etc.

(A3) is essential for the Bayesian update $T(\pi, y)$ preserving monotonicity with respect to π . Theorem 13 (1) in the appendix shows that $T(\pi, y)$ is MLR increasing in π iff $P'\pi$ is MLR increasing in π , and (A3) is a sufficient condition for the latter. TP2 stochastic orders and kernels have been studied in great detail in [18].

(A3) is satisfied by several classes of transition matrices; see [19], [20]. Consider, for example, a tridiagonal transition probability matrix P with $P_{ij} = 0$ for $j \geq i + 2$ and $j \leq i - 2$. As shown in [15, pp.99–100], a necessary and sufficient condition for tridiagonal P to be TP2 is that $P_{ii}P_{i+1,i+1} \geq P_{i,i+1}P_{i+1,i}$.

Structure of Proof of Theorem 1: The proof in the appendix comprises of three steps. Steps 1 and 2 below are proved in meta Theorem 13 in Appendix B under general conditions (A1), (A2), (A3) and (S).

Step 1: Step 1 We first show that the value function $V(\pi)$ is MLR decreasing (see Appendix A for definition). As shown in the proof, the general conditions (A1), (A2), (A3) are sufficient for $V(\pi)$ to be \geq_r decreasing on $\Pi(X)$. This involves showing that $C(\pi, 1)$, $C(\pi, 2)$ are MLR decreasing. (A1)(i) and (A1)(ii) in the appendix are sufficient conditions for this. The proof that (A1)(i) is sufficient for $C(\pi, 1)$ to be MLR decreasing is similar in spirit to the Schur-convexity proof (Theorem A.3 of [30]) with the difference that in Schur convexity the vectors π have elements in ascending order while in our case the elements of π can be in any order.

Conditions (A2) and (A3) are required for MLR monotone updates $T(\pi, y)$ of the belief state, and also first order stochastic dominance monotonicity of $\sigma(\pi, y)$, see Appendix for definition.

Step 2: Step 2 We then prove that $Q(\pi, u)$ is submodular on $[\mathcal{L}(e_X, \bar{\pi}), \geq_{L_X}]$ and $[\mathcal{L}(e_1, \bar{\pi}), \geq_{L_1}]$. (S) is sufficient for $C(\pi, u)$ to be submodular on lines

$\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$. Since we only require submodularity on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$, and these are chains (i.e., totally ordered subsets of a partially ordered set), the condition (S) is less restrictive than requiring submodularity on the entire simplex $\Pi(X)$. Finally (A1),(A2),(A3),(S) are sufficient for $Q(\pi, u)$ to be submodular on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$. So Theorem 13 in Appendix B implies a monotone policy on each chain $[\mathcal{L}(e_X, \bar{\pi}), \geq_{L_X}]$. So there exists a threshold belief state on each line where the optimal policy switches from 1 to 2. (A similar argument holds for lines $[\mathcal{L}(e_1, \bar{\pi}), \geq_{L_1}]$).

Step 3: Step 3 Step 3 is proved in Appendix C. The entire simplex $\Pi(X)$ can be covered by the union of lines $\mathcal{L}(e_X, \bar{\pi})$. The union of the resulting threshold belief states yields the threshold curve $\Gamma(\pi)$. This is illustrated in Fig. 1.

Some Intuition: Recall for $\mathbb{X} = \{1, 2, 3\}$, the belief state space $\Pi(3)$ is an equilateral triangle. So on $\Pi(3)$, more insight can be given to visualize what the above theorem says.⁵ In Fig. 2, six examples are given of decision regions that violate the theorem. To make these examples nontrivial, we have included $e_1 \in \mathcal{R}_1$ in all cases.

The decision regions in Fig. 2(a) violate the condition that $\mu^*(\pi)$ is increasing on lines towards e_3 . Even though \mathcal{R}_1 and \mathcal{R}_2 are individually connected regions, the depicted line $\mathcal{L}(e_3, \bar{\pi})$ intersects the boundary of \mathcal{R}_2 more than once (and so violates Theorem 1).

The decision regions in Fig. 2(b) satisfy Theorem 3 since the stopping set \mathcal{R}_1 is convex. As mentioned in Section III-A, Theorem 1 gives more structure to \mathcal{R}_1 and \mathcal{R}_2 . Indeed, the decision regions in Fig. 2(b) violate Theorem 1. They violate the statement that the policy is increasing on lines towards e_3 since the boundary of \mathcal{R}_1 (i.e., threshold curve Γ) cannot intersect a line from e_3 more than once. Therefore, Theorem 1 says a lot more about the structure of the boundary than convexity does. In particular, for the PH-distributed change time without variance penalty, Theorems 1 and 3 together say that the threshold curve Γ is convex and cannot intersect a line $\mathcal{L}(e_3, \bar{\pi})$ or a line $\mathcal{L}(e_1, \bar{\pi})$ more than once.

Fig. 2(c) also satisfies Theorem 3 since \mathcal{R}_1 is convex. But the decision regions in Fig. 2(c) violate Statement (ii) of Theorem 1. In particular, if e_1 and e_3 lie in \mathcal{R}_1 , then e_2 should also lie in \mathcal{R}_1 . Again this reveals that Theorem 1 says a lot more about the structure of the stopping region even for the case of zero variance penalty ($\alpha = 0$).

Fig. 2(d) also satisfies Theorem 3 since \mathcal{R}_1 is convex; but does not satisfy Theorem 1 since \mathcal{R}_2 is not a connected set. Indeed when \mathcal{R}_2 is not connected as shown in the proof of Theorem 1, the policy $\mu^*(\pi)$ is not monotone on the line $\mathcal{L}(e_3, \bar{\pi})$ since it goes from 2 to 1 to 2.

⁵It can be shown that the threshold curve is a Borel measurable function. Also, Γ intersects each line segment from vertex e_3 and each line segment from vertex e_1 at most once. This implies Γ can be parametrized by a pair of monotonically decreasing angles with respect to vertices e_1 and e_3 . By Lebesgue theorem [42], a monotone function is differentiable almost everywhere. So for $\mathbb{X} = \{1, 2, 3\}$, Γ is differentiable almost everywhere.

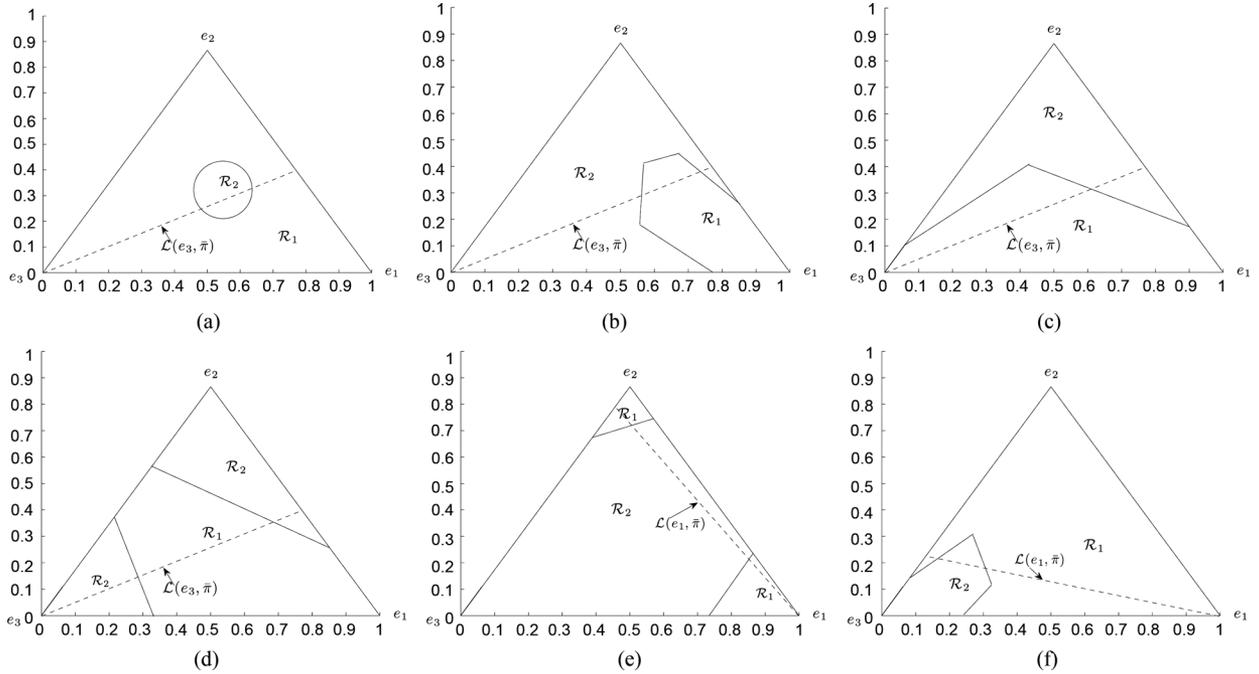


Fig. 2. Examples of decision regions that violate Theorem 1 on belief space $\Pi(3)$. (a) Example 1. (b) Example 2. (c) Example 3. (d) Example 4. (e) Example 5. (f) Example 6.

Figs. 2(e) and 2(f) violate Theorem 1 since the optimal policy $\mu^*(\pi)$ is not monotone on line $\mathcal{L}(e_1, \bar{\pi})$; it goes from 1 to 2 to 1. For the case $\alpha = 0$, Figs. 2(e) and 2(f) violate Theorem 3 since the stopping region \mathcal{R}_1 is nonconvex.

Since the conditions of Theorem 1 are sufficient conditions, what happens when they do not hold? In Section VIII, we will give a numerical example where (S-Ex1) is violated and \mathcal{R}_1 is no longer a connected set Fig. 6(d). It is straightforward to construct other examples where both \mathcal{R}_1 and \mathcal{R}_2 are disconnected regions when the assumptions of Theorem 1 are violated.

D. Characterization of Optimal Linear Decision Threshold

This subsection assumes that (A1-Ex1), (A2), (A3), (S-Ex1) of Section III-A hold. So Theorem 1 applies and computing the optimal policy μ^* reduces to estimating the threshold curve Γ . In general, any user-defined basis function approximation can be used to parametrize this curve. However, any such approximation needs to capture the essential feature of Theorem 1: the parametrized optimal policy needs to be MLR increasing on lines. [An identical discussion applies to Theorem 2 with assumptions (AS-Ex1), (A2), (A3)].

Below, we derive the optimal linear approximation to the threshold curve Γ on simplex $\Pi(X)$. Such a linear decision threshold has two attractive properties: (i) Estimating it is computationally efficient. (ii) We give conditions on the coefficients of the linear threshold that are necessary and sufficient for the resulting policy to be MLR increasing on lines. Due to the necessity and sufficiency of the condition, optimizing over the space of linear thresholds on $\Pi(X)$ yields the optimal linear approximation to threshold curve Γ .

On $\Pi(X)$, define the linear threshold policy $\mu_\theta(\pi)$ as

$$\mu_\theta(\pi) = \begin{cases} \text{stop} = 1, & \text{if } [0 \quad 1 \quad \theta']' \begin{bmatrix} \pi \\ -1 \end{bmatrix} < 0 \\ \text{continue} = 2, & \text{otherwise} \end{cases} \quad \pi \in \Pi(X). \quad (32)$$

Here $\theta = (\theta(1), \dots, \theta(X-1))' \in \mathbb{R}^{X-1}$ denotes the parameter vector of the linear threshold policy. (Since $\Pi(X) \subset \mathbb{R}^{X-1}$, a linear hyperplane on $\Pi(X)$ is parametrized by $X-1$ coefficients).

Theorem 4 below characterizes the optimal linear decision threshold approximation to the threshold curve on $\Pi(X)$. Assume conditions (A-Ex1), (A2), (A3), (S-Ex1) hold for the quickest detection problem (20) so that from Theorem 1, the optimal policy $\mu^*(\pi)$ is MLR increasing on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$. Assume the conditions of Lemma 1 hold, so that one only needs to search for the optimal linear threshold in the interior of $\Pi(X)$. Finally, the requirement that e_1 lies in the stopping set, means $\mu_\theta(e_1) < 0$ which implies $\theta(X-1) > 0$.

Theorem 4 (Optimal Linear Threshold Policy): For belief states $\pi \in \Pi(X)$, the linear threshold policy $\mu_\theta(\pi)$ defined in (32) is:

- (i) MLR increasing on lines $\mathcal{L}(e_X, \bar{\pi})$ iff $\theta(X-2) \geq 1$ and $\theta(i) \leq \theta(X-2)$ for $i < X-2$;
- (ii) MLR increasing on lines $\mathcal{L}(e_1, \bar{\pi})$ iff $\theta(i) \geq 0$, for $i < X-2$.

■

The proof of Theorem 4 is in Appendix E. As a consequence of Theorem 4, the optimal linear threshold approximation to

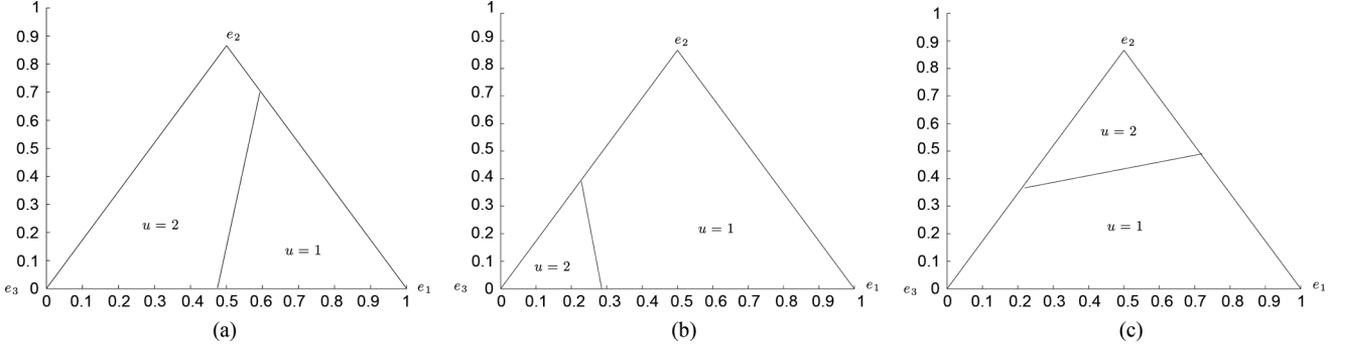


Fig. 3. Examples of Valid Linear Threshold Policies on belief space $\Pi(X)$ for $X = 3$ (Case 1 and Case 2). Case 3 is invalid.

threshold curve Γ of Theorem 1 is the solution of the following constrained optimization problem:

$$\begin{aligned} \theta^* &= \arg \min_{\theta \in \mathbb{R}^{X-1}} J_{\mu_\theta}(\pi_0) \\ \text{subject to } &0 \leq \theta(i) \leq \theta(X-2) \\ &\theta(X-2) \geq 1 \text{ and } \theta(X-1) > 0 \end{aligned} \quad (33)$$

where the cost $J_{\mu_\theta}(\pi_0)$ is obtained as in (20) by applying threshold policy μ_θ in (32).

Remark: The constraints in (33) are necessary and sufficient for the linear threshold policy (32) to be MLR increasing on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$. That is, (33) defines the set of all MLR increasing linear threshold policies—it does not leave out any MLR increasing policies; nor does it include any non MLR increasing policies. Therefore, optimizing over the space of MLR increasing linear threshold policies yields the optimal linear approximation to threshold curve Γ .

Intuition: Consider $X = 3$, $\mathbb{X} = \{1, 2, 3\}$ so that the belief space $\Pi(X)$ is an equilateral triangle. Then with $(\omega(1), \omega(2))$ denoting Cartesian coordinates in the equilateral triangle, clearly $\pi(2) = 2\omega(2)/\sqrt{3}$, $\pi(1) = \omega(1) - \omega(2)/\sqrt{3}$ and the linear threshold satisfies

$$\omega(2) = \frac{\sqrt{3}\theta(1)}{2 - \theta(1)}\omega(1) + (\theta(2) - \theta(1))\frac{\sqrt{3}}{2 - \theta(1)}. \quad (34)$$

So the conclusion of Theorem 4 that $\theta(1) \geq 1$ implies that the linear MLR increasing threshold has slope of 60° or larger. For $\theta(1) > 2$, it follows from (34) that the slope of the linear threshold becomes negative, i.e., more than 90° . For a nondegenerate threshold, the $\omega(1)$ intercept of the line should lie in $[0, 1]$ implying $\theta(1) > \theta(2)$ and $\theta(2) > 0$. Fig. 3 illustrates these results. Figs. 3(a) and 3(b) illustrate valid linear thresholds. In Figs. 3(a) and 3(b), the conditions of Theorem 4 hold (the slope is larger than 60° and $\omega(1)$ intercept is in $[0, 1]$). Fig. 3(c) shows an invalid threshold (since the slope is smaller than 60° and $\omega(1)$ intercept lies outside $[0, 1]$). In other words, Fig. 3(c) shows an invalid threshold since it violates the requirement that $\mu_\theta(\pi)$ is decreasing on lines towards e_3 on $\Pi(3)$. (A line segment $\mathcal{L}(e_3, \bar{\pi})$ starting from some point $\bar{\pi}$ on facet

(e_2, e_1) and connected to e_3 would start in the region $u = 2$ and then go to region $u = 1$. This violates the requirement that $\mu_\theta(\pi)$ is increasing on lines towards e_3).

E. Algorithm to Compute the Optimal Linear Decision Curve

In this section a stochastic approximation algorithm is presented to estimate the optimal threshold vector θ^* in (33). Because the cost $J_{\mu_\theta}(\pi_0)$ in (33) cannot be computed in closed form, we resort to simulation based stochastic optimization. Let $n = 1, 2, \dots$, denote iterations of the algorithm. The aim is to solve the following linearly constrained stochastic optimization problem:

$$\begin{aligned} \text{Compute } \theta^* &= \arg \min_{\theta \in \Theta} \mathbb{E}\{J_n(\mu_\theta)\} \\ \text{subject to } &0 \leq \theta(i) \\ &\leq \theta(X-2), \theta(X-2) \geq 1 \text{ and } \theta(X-1) > 0. \end{aligned} \quad (35)$$

Here, for each initial condition π_0 , the sample path cost $J_n(\mu_\theta, \pi_0)$ is evaluated as

$$\begin{aligned} J_n(\mu_\theta, \pi_0) &= \sum_{k=1}^{\infty} \rho^{k-1} C(\pi_k, u_k) \quad \text{where } u_k = \mu_\theta(\pi_k) \\ &\text{is computed via (32)} \\ J_n(\mu_\theta) &= \frac{1}{L} \sum_{l=1}^L J_n(\mu_\theta, \pi_0^{(l)}) \quad \text{where prior } \pi_0^{(l)} \\ &\text{is sampled uniformly from simplex } \Pi(X). \end{aligned} \quad (36)$$

A convenient way of sampling uniformly from $\Pi(X)$ is to use the Dirichlet distribution (i.e., $\pi_0(i) = x_i / \sum_i x_i$, where $x_i \sim$ unit exponential distribution).

The above constrained stochastic optimization problem can be solved by a variety of methods. One method is to convert it into an equivalent unconstrained problem via the following parametrization: Let $\phi = (\phi(1), \dots, \phi(X-1))' \in \mathbb{R}^{X-1}$ and parametrize θ as

$$\begin{aligned} \theta^\phi &= [\theta^\phi(1), \dots, \theta^\phi(X-1)]' \\ &\quad \text{where} \\ \theta^\phi(i) &= \begin{cases} \phi^2(X-1), & i = X-1 \\ 1 + \phi^2(X-2), & i = X-2 \\ (1 + \phi^2(X-2)) \sin^2(\phi(i)), & i = 1, \dots, X-3. \end{cases} \end{aligned} \quad (37)$$

Then θ^ϕ trivially satisfies constraints in (35). So (35) is equivalent to the following unconstrained stochastic optimization problem:

$$\begin{aligned} &\text{Compute } \mu_{\phi^*}(\pi) \text{ where } \phi^* = \arg \min_{\phi \in \mathbb{R}^{X-1}} \mathbb{E}\{J_n(\phi)\} \text{ and} \\ &J_n(\phi) \text{ is computed using (36) with policy } \mu_{\theta^\phi}(\pi) \\ &\text{evaluated according to (37).} \end{aligned} \quad (38)$$

Algorithm 1 below, uses the Simultaneous Perturbation Stochastic Approximation (SPSA) algorithm [47] to generate a sequence of estimates $\hat{\phi}_n$, $n = 1, 2, \dots$, that converges to a local minimum of the optimal linear threshold ϕ^* with policy $\mu_{\phi^*}(\pi)$.

Algorithm 1 Policy Gradient Algorithm for computing optimal linear threshold policy

Assume (A1-Ex1), (A2), (A3), (S-Ex1) hold so that the optimal social policy is characterized by a threshold switching curve in Theorem 1.

Step 1: Choose initial threshold coefficients $\hat{\phi}_0$ and linear threshold policy $\mu_{\hat{\phi}_0}$.

Step 2: For iterations $n = 0, 1, 2, \dots$

- Evaluate sample cost $J_n(\hat{\phi}_n)$ using (38). Compute gradient estimate

$$\begin{aligned} \hat{\nabla}_\phi J_n(\hat{\phi}_n) &= \frac{J_n(\hat{\phi}_n + \Delta_n \omega_n) - J_n(\hat{\phi}_n - \Delta_n \omega_n)}{2\Delta_n} \omega_n \\ \omega_n(i) &= \begin{cases} -1 & \text{with probability 0.5} \\ +1 & \text{with probability 0.5.} \end{cases} \end{aligned}$$

Here $\Delta_n = \Delta/(n+1)^\gamma$ denotes the gradient step size with $0.5 \leq \gamma \leq 1$ and $\Delta > 0$.

- Update threshold coefficients $\hat{\phi}_n$ via stochastic approximation algorithm

$$\begin{aligned} \hat{\phi}_{n+1} &= \hat{\phi}_n - \epsilon_{n+1} \hat{\nabla}_\phi J_n(\hat{\phi}_n), \quad \epsilon_n \\ &= \epsilon/(n+1+s)^\zeta, \quad 0.5 < \zeta \leq 1, \text{ and } \epsilon, s > 0. \end{aligned} \quad (39)$$

The above SPSA algorithm [47] picks a single random direction ω_n along which direction the derivative is evaluated at each batch n . Unlike the Kiefer-Wolfowitz finite difference algorithm to evaluate the gradient estimate $\hat{\nabla}_\phi J_n$ in (39), SPSA requires only 2 batch simulations, i.e., the number of evaluations is independent of dimension of parameter ϕ . Because the stochastic gradient algorithm (39) converges to local optima, it is necessary to try several initial conditions $\hat{\phi}_0$. The computational cost at each iteration is linear in the dimension of θ and is independent of the observation alphabet size.

For fixed θ , the samples $J_n(\mu_\theta)$ in (36) are simulated independently and have identical distribution. Thus, the proof that $\theta_n = \theta^{\hat{\phi}_n}$ generated by Algorithm 1 converges to a local optimum of $\mathbb{E}\{J_n(\mu_\theta)\}$ (defined in (35)) with probability one, is a straightforward application of techniques in [26] (which gives general convergence methods for Markovian dependencies).

Remark: More sophisticated gradient estimation methods can be used instead of the SPSA finite difference algorithm

given here. For example, [23], [35] present score function and weak derivative approaches for estimating the gradient of a Markov process with respect to a policy. In [4] the score function method is used to perform gradient-based reinforcement learning. These algorithms are applicable to solve the constrained stochastic optimization problem (35) thereby yielding the optimal linear threshold policy. If the change time distribution (specified by P) and the observation likelihoods (specified by B) are not completely specified, but (A2) and (A3) hold, Theorem 1 applies and reinforcement learning algorithms [4] can be used to solve (35). Moreover [25], [54] analyze the tracking properties of stochastic approximation algorithms when the transition and observation matrices and time varying.

IV. EXAMPLE 2: QUICKEST TRANSIENT DETECTION WITH VARIANCE PENALTY

Our second example deals with Bayesian quickest transient detection.⁶ We show below that under similar assumptions to quickest time detection, the threshold switching curve of Theorem 1 holds. Therefore, the linear threshold results and Algorithm 1 hold.

The set up is identical to Section II with state space $\mathbb{X} = \{1, 2, 3\}$. The transition probability matrix and initial distribution are

$$P = \begin{bmatrix} 1 & 0 & 0 \\ p_{21} & p_{22} & 0 \\ 0 & p_{32} & p_{33} \end{bmatrix}, \quad \pi_0 = e_3. \quad (40)$$

So the Markov chain starts in state 3. After some geometrically distributed time it jumps to the transient state 2. Finally after residing in state 2 for some geometrically distributed time, it then jumps to the absorbing state 1.

In quickest transient detection, we are interested in detecting transition to state 2 with minimum cost. The action space is $\mathcal{U} = \{1(\text{stop}), 2(\text{continue})\}$. The stop action $u = 1$ declares that transient state 2 was visited.

We choose the following costs (see [39] for other choices). Similar to (18), let $d_i I(x_k = e_i, u_k = 2)$ denote the delay cost in state e_i , $i \in \mathbb{X}$. Of course $d_3 = 0$ since $x_k = e_3$ implies that the transient state has not yet been visited. So the expected delay cost is

$$\begin{aligned} \sum_{i \in \mathbb{X}} d_i \mathbb{E}\{x_k = e_i, u_k = 2 | \mathcal{F}_k\} &= \mathbf{d}' \pi_k \text{ where} \\ \mathbf{d} &= (d_1, d_2, d_3)', \quad d_3 = 0. \end{aligned} \quad (41)$$

Typically the elements of the delay vector \mathbf{d} are chosen as $d_1 \geq d_2 > 0$ so that state 1 (final state) accrues a larger delay than the transient state. This gives incentive to declare that transient state 2 was visited when the current state is 2, rather than wait until the process reaches state 1.

The false alarm cost for declaring $u = 1$ (transient state 2 was visited) when $x = 1$ is zero since the final state 1 could only have been reached after visiting transient state 2. So the false alarm penalty is $\mathbb{E}\{I(x_k = e_3, u_k = 1) | \mathcal{F}_k\} = 1 - (e_1 + e_2)' \pi_k$

⁶The author gratefully acknowledges Dr. Venugopal Veeravalli at U. Illinois for describing quickest transient detection and giving access to the preprint [39].

for action $u_k = 1$. For convenience, in the variance penalty (1), we choose $g = [0, 0, 1]'$. So from (16), (18), the expected stopping cost and continuing costs are

$$\begin{aligned}\bar{C}(\pi, 1) &= \alpha(g'\pi - (g'\pi)^2) + \beta(1 - (e_1 + e_2)'\pi) \\ \bar{C}(\pi, 2) &= \mathbf{d}'\pi.\end{aligned}\quad (42)$$

The optimal decision policy $\mu^*(\pi)$ and stopping set \mathcal{R}_1 are as in (24)–(26).

Main Result: The following assumptions are similar to those in quickest time detection. Note that due to its structure, P in (40) is always TP2 (i.e., (A3) in Section III-A holds).

(S-Ex2) The scaling factor for the variance penalty satisfies $\alpha \leq \frac{d_2 + \beta - \rho\beta P_{33}}{1 + \rho P_{33}}$.

For $\alpha = 0$ (zero variance penalty) (S-Ex2) holds trivially.

Theorem 5: Consider the quickest transient detection problem with delay and stopping costs in (42). Then under (A2), (40), (S-Ex2), the conclusions of Theorem 1 hold. Also if the observation likelihoods are nonzero, then for $k \geq 2$, $\pi_k \in \Pi^o(X)$ and the threshold is nondegenerate, see Section III-B-3 and Lemma 1. Thus, Algorithm 1 estimates the optimal linear threshold. (The proof follows from meta Theorem 13 in Appendix B and Theorem 1). ■

PH-Distributed Change Times: More generally, suppose the process x jumps after a PH-distributed time to transient state. Then after another PH-distributed time period, it jumps to the absorbing state. We show that Theorem 5 continues to hold.

To model the two PH-distributed change times, let $\mathbb{1}$ denote the absorbing state, $\mathbb{T} = \{2, \dots, X_1 + 1\}$ denote the set of transient states and $\mathbb{S} = \{X_1 + 2, \dots, X_1 + X_2 + 1\}$ denote set of starting states. Define the $(X_1 + X_2 + 1) \times (X_1 + X_2 + 1)$ transition matrix

$$P = \begin{bmatrix} e'_{X_1} & \mathbf{0}_{X_2} \\ \underline{P}_{X_1 \times X_1} & \mathbf{0}_{X_1 \times X_2} \\ \mathbf{0}_{X_2} & \bar{P}_{(X_1 + X_2 - 1) \times X_2} \end{bmatrix}. \quad (43)$$

Suppose the Markov chain starts in \mathbb{S} . Then after a PH-distributed time it jumps to \mathbb{T} and finally after another PH-distributed time, jumps to state 1. Just as in Section II, \bar{P} and \underline{P} determine the PH-distribution in the start and transient states, respectively. Let \mathbf{d} and \mathbf{f} denote the delay and false alarm vectors. \mathbf{d} is a vector with decreasing elements with $d_i = 0$, $i \in \mathbb{S}$. \mathbf{f} is a vector with increasing elements with $f_i = 0$, $i \neq \mathbb{S}$. The delay and stopping costs are $\bar{C}(\pi, 1) = \beta\mathbf{f}'\pi$, $\bar{C}(\pi, 2) = \mathbf{d}'\pi$. Then the conclusions of Theorem 5 hold under (A2), (A3) if the decision maker designs \mathbf{f} and \mathbf{d} to satisfy the following linear constraints:

$$\begin{aligned}f_{X_1+2} &\geq 1, (\mathbf{d} + \beta(\rho P - I)\mathbf{f})'(e_i - e_{i+1}) \geq 0 \\ i &= 1, \dots, X_1 + X_2.\end{aligned}\quad (44)$$

The decision maker can design suitable \mathbf{f} and \mathbf{d} satisfying (44) using a linear programming solver.

Remark: In the formulation of [39], it is assumed that states 1 and 3 are indistinguishable in terms of observations, i.e., $B_{1y} = B_{3y}$ for all $y \in \mathbb{Y}$. In this case, obviously (A2) does not hold.

At this stage, we are unable to prove the structural result of Theorem 5 when (A2) does not hold. (Relabelling state 1 as 2 and state 2 as 1 does not work. Then B satisfies (A2) but the transition matrix P is longer TP2). Nevertheless, we have the following result which follows from Theorem 3.

Corollary: For $\alpha = 0$, (no variance penalty), then the stopping region \mathcal{R}_1 in quickest transient detection is a convex subset of $\Pi(X)$. ■

V. EXAMPLE 3: QUICKEST DETECTION WITH EXPONENTIAL PENALTY FOR DELAY

In this example, we generalize the results of Poor [36], which deals with exponential delay penalty and geometric change times. We consider exponential delay penalty with PH-distributed change time. Our formulation involves risk sensitive partially observed stochastic control, see Section I for motivation. We first show that the exponential penalty cost function in [36] is a special case of risk-sensitive stochastic control cost function when the state space dimension $X = 2$. We then use the risk-sensitive stochastic control formulation to derive structural results for PH-distributed change time. In particular, the main result below (Theorem 6) shows that the threshold switching curve still characterizes the optimal stopping region \mathcal{R}_1 . The assumptions and main results are conceptually similar to Theorem 1.

Since our aim is to interpret and extend the results of [36] using risk sensitive control, we consider the same costs as in [36], so $\alpha = 0$ (no variance penalty). Below, we will use $c(e_i, u = 1)$ to denote false alarm costs and $c(e_i, u = 2)$ to denote delay costs, where $i \in \mathbb{X}$.

Risk sensitive control [6] considers the exponential cost function

$$J_\mu(\pi_0) = \mathbb{E}_{\pi_0}^\mu \left\{ \exp \left(\epsilon \sum_{k=1}^{\tau-1} c(x_k, u_k = 2) + \epsilon c(x_\tau, u_\tau = 1) \right) \right\} \quad (45)$$

where $\epsilon > 0$ is the risk sensitive parameter.

Let us first show that the exponential penalty cost in [36] is a special case of (45) for consider the case $X = 2$ (geometric distributed change time). For the state $x \in \{e_1, e_2\}$, choose $c(x, u = 1) = \beta I(x \neq e_1, u = 1) = \beta(1 - e'_1 x)$ (false alarm cost), $c(x, u = 2) = dI(x = e_1, u = 2) = de'_1 x$ (delay cost). Then it is easily seen that $\sum_{k=1}^{\tau-1} c(x_k, u_k = 2) + c(x_\tau, u_\tau = 1) = d|\tau - \tau^0|^+ + \beta I(\tau < \tau^0)$. Therefore, [recall τ^0 is defined in (6) and τ is defined in (19)]; see (46) at the bottom of the next page, which is identical to Poor's exponential delay cost function [36, Eq.40]. Thus, the Bayesian quickest time detection with exponential delay penalty in [36] is a special case of a risk sensitive stochastic control problem.

We consider the delay cost as in (17); so for state $x \in \{e_1, \dots, e_X\}$, $c(x, u_k = 2) = de'_1 P^k x$. To get an intuitive feel for this modified delay cost function, for the case $X = 2$

$$\begin{aligned} &\sum_{k=1}^{\tau-1} c(x_k, u_k = 2) + c(x_\tau, u_\tau = 1) \\ &= d|\tau - \tau^0|^+ + \beta I(\tau < \tau^0) + dP_{21}(\tau^0 - 1)I(\tau^0 < \tau).\end{aligned}$$

Therefore, by using (21), for $X = 2$, the exponential delay cost function is

$$J_\mu(\pi_0) = (e^{\epsilon\beta} - 1)\mathbb{P}_{\pi_0}^\mu(\tau < \tau^0) + \mathbb{E}_{\pi_0}^\mu \{e^{\epsilon d} [|\tau - \tau^0|^+ + P_{21}(\tau_0 - 1)I(\tau_0 < \tau)]\}. \quad (47)$$

This is similar to (46) except for the additional term $P_{21}(\tau_0 - 1)I(\tau_0 < \tau)$ in the exponential.

With the above motivation, in the rest of this section we consider risk sensitive quickest time detection for PH-distributed change time, i.e., $X \geq 2$. Let π denote the risk sensitive belief state, see [14], [17] for extensive descriptions of the risk sensitive belief state and verification theorems for dynamic programming in risk sensitive control. It can be shown [14] that the value function $\bar{V}(\pi)$ satisfies

$$\bar{V}(\pi) = \min\{\bar{C}(\pi, 1), \sum_{y \in \mathcal{Y}} \bar{V}(T(\pi, Y))\sigma(\pi, y)\} \quad (48)$$

where with $R_1 = (1, e^{\epsilon\beta}, \dots, e^{\epsilon\beta})'$, $R_2 = (e^{\epsilon d}, e^{\epsilon d P_{21}}, \dots, e^{\epsilon d P_{X1}})'$, B_y defined in (11)

$$\begin{aligned} \bar{C}(\pi, 1) &= R_1' \pi, \quad T(\pi, y) = \frac{B_y P' \text{diag}(R_2) \pi}{\sigma(\pi, y)}, \sigma(\pi, y) \\ &= \mathbf{1}' B_y P' \text{diag}(R_2) \pi. \end{aligned} \quad (49)$$

As in Section II-B, define $V(\pi) = \bar{V}(\pi) - \bar{C}(\pi, 1)$. Then $V(\pi)$ satisfies Bellman's (25) with

$$C(\pi, 1) = 0, \quad C(\pi, 2) = R_1'(P' \text{diag}(R_2) - I)\pi. \quad (50)$$

Assume the following condition holds

(A1-Ex3) The elements of $R_1'(P' \text{diag}(R_2) - I)$ are decreasing wrt $i = 1, 2, \dots, X$.

Evaluating $C(\pi, 2) = R_1'(P' \text{diag}(R_2) - I)\pi$, then (A1-Ex3) is equivalent to

$$e^{\epsilon d} - 1 \geq e^{\epsilon d P_{21}}(P_{21} + e^{\epsilon\beta}(1 - P_{21})) - e^{\epsilon\beta} \text{ and } e^{\epsilon d P_{i1}}(P_{i1} + e^{\epsilon\beta}(1 - P_{i1})) \text{ decreasing in } i \in \{2, \dots, X\}.$$

For example, if $d = \epsilon = 1$, then for $\beta \geq 1$, the following are verified by elementary calculus:

- (i) (A1-Ex3) always holds for $\beta \geq 1$ when $X = 2$ (geometric distributed change time).
- (ii) For PH-distributed change time, if (A3) holds, then (A1-Ex3) always holds providing $P_{21} < 1/(e^\beta - 1)$.

Theorem 6: The stopping region \mathcal{R}_1 is a convex subset of $\Pi(X)$. Under (A1-Ex3), (A2), (A3), Theorem 1 holds. Thus, Algorithm 1 estimates the optimal linear threshold. ■

The proof is in Appendix F.

Remarks:

- (i) Delay Formulation in [36]: Consider the formulation in Poor [36] which is equivalent to (46). Then for the geometric distributed case $X = 2$, the convexity of \mathcal{R}_1 holds using a similar proof to above. Since $\Pi(X)$ is a 1-D simplex and $e_1 \in \mathcal{R}_1$, convexity implies there exists (a possible degenerate) threshold point π^* that characterizes \mathcal{R}_1 such that the optimal policy is of the form (1). As a sanity check, the analogous condition to (A1-Ex3) reads $e^{\epsilon d} - 1 > P_{21}(1 - e^{\epsilon\beta})$. This always holds for $\epsilon \geq 0$. Therefore, assuming (A2) holds, the above theorem holds for Poor's [36] exponential delay penalty case under (A2). (Recall (A3) holds trivially when $X = 2$). Finally, for $X > 2$, using a similar proof (see Theorem 2), one can again show that the conclusions of Theorem 6 hold.
- (ii) *Other Examples:* With the above risk sensitive formulation, the dynamic programming equation (48) for the exponential delay case is very similar to the other examples in this paper. Therefore, it is straightforward to generalize the above exponential penalty result to quickest transient detection (of Section IV, and social learning stopping time problems considered below.

VI. EXAMPLE 4 & 5: STOPPING TIME PROBLEMS IN MULTI-AGENT SOCIAL LEARNING

Here we consider stopping time problems in multi-agent social learning. We present two results:

- (i) Section VI-B (Example 4) considers a multi-agent system seeking to solve a Bayesian stopping time problem.
- (ii) Section VI-C (Example 5) deals with *constrained optimal* social learning which is formulated in Chamley [11] as a sequential stopping time problem. We show that the optimal policy has a threshold switching curve similar to Theorem 1.

A. Motivation: Social Learning Amongst Myopic Agents

Since social learning only serves as a motivation for subsequent subsections, our description is brief; see [11]. Consider a countable infinite number of agents performing social learning to estimate an underlying random state x . Each agent acts once in a predetermined sequential order indexed by $k = 1, 2, \dots$. One can also view k as the discrete time instant when agent k acts. A key difference between social learning compared to the formulation in previous sections is that agent k does not have access to the belief state or private observations of previous agents. Instead each agent k only has access to the actions taken by previous agents together with its own current private observation y_k .

$$\begin{aligned} J_\mu(\pi_0) &= \mathbb{E}_{\pi_0}^\mu \{ \exp(\epsilon d |\tau - \tau^0|^+ + \epsilon\beta I(\tau < \tau^0)) \} [I(\tau < \tau^0) + I(\tau = \tau^0) + I(\tau > \tau^0)] \\ &= \mathbb{E}_{\pi_0}^\mu \{ \exp(\epsilon\beta) I(\tau < \tau^0) + \exp(\epsilon d |\tau - \tau^0|^+) I(\tau > \tau^0) + 1 \} \\ &= \mathbb{E}_{\pi_0}^\mu \{ (e^{\epsilon\beta} - 1) I(\tau < \tau^0) + e^{\epsilon d |\tau - \tau^0|^+} \} \\ &= (e^{\epsilon\beta} - 1) \mathbb{P}_{\pi_0}^\mu(\tau < \tau^0) + \mathbb{E}_{\pi_0}^\mu \{ e^{\epsilon d |\tau - \tau^0|^+} \} \end{aligned} \quad (46)$$

Throughout this section, we assume that the observation space \mathbb{Y} is finite. Let $y_k \in \mathbb{Y} = \{1, 2, \dots, Y\}$ denote the private observation of agent k and $a_k \in \mathbb{A} = \{1, 2, \dots, A\}$ denote the action agent k takes. Define the sigma algebras

$$\begin{aligned} \mathcal{H}_k & \sigma\text{-algebra generated by } (a_1, \dots, a_{k-1}, y_k) \\ \mathcal{G}_k & \sigma\text{-algebra generated by } (a_1, \dots, a_{k-1}, a_k). \end{aligned} \quad (51)$$

The social learning model [8], [11] comprises of the following ingredients:

- (A) The state of nature x as in Section II-A except that the transition matrix is $P = I$. That is, the state of nature is a random variable with distribution π_0 (see (4)) instead of a random process.
- (B) At time k , agent k records a private observation $y_k \in \mathbb{Y}$ from the observation distribution $B_{iy} = P(y|x = e_i)$, $i \in \mathbb{X}$.
- (C) *Private belief*: Using the public belief π_{k-1} available at time $k-1$ (defined in Step (v) below), agent k then updates its Bayesian private belief η_k as in (11) with $P = I$. Here

$$\begin{aligned} \eta_k &= \mathbb{E}\{x|\mathcal{H}_k\} = (\eta_k(i), i \in \mathbb{X}) \\ \eta_k(i) &= P(x = e_i|a_1, \dots, a_{k-1}, y_k) \\ & \text{initialized with } \pi_0. \end{aligned} \quad (52)$$

- (D) *Myopic Action*: Agent k then takes action $a_k \in \mathbb{A} = \{1, 2, \dots, A\}$ to minimize myopically its expected cost $a_k = \arg \min_{a \in \mathbb{A}} \{c'_a \eta_k\}$. Here $c_a = (c(e_i, a), i \in \mathbb{X})$ denotes an X -dimensional cost vector, and $c(e_i, a)$ denotes the cost incurred when the underlying state is e_i and the agent picks action a . Thus, agent k chooses action

$$\begin{aligned} a_k &= a(\pi_{k-1}, y_k) = \arg \min_{a \in \mathbb{A}} \mathbb{E}\{c(x, a)|\mathcal{H}_k\} \\ &= \arg \min_{a \in \mathbb{A}} \{c'_a \eta_k\} \end{aligned} \quad (53)$$

- (E) *Social learning and Public belief*: Finally agent k broadcasts this action a_k to subsequent agents. Define the public belief π_k as the posterior distribution of the state x given all actions taken up to time k

$$\begin{aligned} \pi_k &= \mathbb{E}\{x|\mathcal{G}_k\} = (\pi_k(i), i \in \mathbb{X}) \\ \pi_k(i) &= P(x = e_i|a_1, \dots, a_k), \quad \text{initialized with } \pi_0. \end{aligned} \quad (54)$$

Based on the action a_k every agent (apart from k) perform social learning to update their public belief according to the following ‘‘social learning Bayesian filter’’:

$$\begin{aligned} \pi_k &= T(\pi_{k-1}, a_k), \text{ where} \\ T(\pi, a) &= \frac{R'_a \pi}{\sigma(\pi, a)}, \quad \sigma(\pi, a) = \mathbf{1}'_X R'_a \pi. \end{aligned} \quad (55)$$

In (55), $R'_a = \text{diag}(P(a|x = e_i, \pi), i \in \mathbb{X})$ with elements

$$P(a_k = a|x = e_i, \pi_{k-1} = \pi)$$

$$\begin{aligned} &= \sum_{y \in \mathbb{Y}} P(a_k = a|y, \pi) P(y|x = e_i) \\ &= \sum_{y \in \mathbb{Y}} \prod_{\tilde{a} \in \mathbb{A} - \{a\}} I(c'_a B_y \pi < c'_{\tilde{a}} B_y \pi) P(y|x = e_i) \end{aligned} \quad (56)$$

where $I(\cdot)$ denotes the indicator function and B_y is defined in (11).

The following well known result [8], [11] states that eventually after some finite time \bar{k} , all agents pick the same action and the private belief freezes. This is termed an information cascade. The proof follows via an elementary application of the martingale convergence theorem.

Theorem 7 [8]: The above social learning model leads to an information cascade (i.e., all agents herd) in finite time with probability 1. That is there exists a finite time \bar{k} after which social learning ceases, i.e, public belief $\pi_{k+1} = \pi_k$, $k \geq \bar{k}$, and all agents pick the same action, i.e., $a_{k+1} = a_k$, $k \geq \bar{k}$. ■

B. Example 4: Sequential Detection With Social Learning

Suppose a multi-agent system makes local decisions and performs social learning as above.

Given such a protocol, how can the multi-agent system make a global decision when to stop? As mentioned in Section I, such problems are motivated in decision systems where a global decision needs to be made based on local decisions of agents.

We consider a Bayesian sequential detection problem for state $x = e_1$. The main result below (Theorem 8) is that the global decision of when to stop is a multithreshold function of the belief state. This unusual behavior is because in social learning, the action likelihood probabilities $R'_a \pi$ in (56) depend on the belief state π .

Consider $\mathbb{X} = \mathbb{Y} = \{1, 2\}$ and $\mathbb{A} = \{1, 2\}$ and the social learning model of Section VI-A, where the costs $c(e_i, a)$ satisfy

$$c(e_1, 1) < c(e_1, 2), \quad c(e_2, 2) < c(e_2, 1). \quad (57)$$

Otherwise one action will always dominate the other action and the problem is un-interesting.

Redefine the sigma algebras in (51) to include the action history

$$\begin{aligned} \mathcal{H}_k & \sigma\text{-algebra generated by } (a_1, \dots, a_{k-1}, y_k, u_1, \dots, u_k) \\ \mathcal{G}_k & \sigma\text{-algebra generated by } (a_1, \dots, a_{k-1}, a_k, u_1, \dots, u_k). \end{aligned} \quad (58)$$

Let τ denote a stopping time adapted to the sequence of sigma-algebras \mathcal{G}_k , $k \geq 1$ (see (58)). In words, each agent has only the public belief obtained via social learning to make the global decision of whether to continue or stop. The goal is to solve the following sequential detection problem to detect state e_1 : Pick the stopping time τ to minimize (59) at the bottom of the page. As in previous sections, the first term is the delay cost and

$$J_\mu(\pi_0) = \mathbb{E}_{\pi_0}^\mu \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \mathbb{E} \{ dI(x = e_1) | \mathcal{G}_{k-1} \} + \rho^{\tau-1} \beta \mathbb{E} \{ I(x \neq e_1) | \mathcal{G}_{\tau-1} \} \right\} \quad (59)$$

penalizes the decision of choosing $u_k = 2$ (continue) when the state is e_1 by the non-negative constant d . The second term is the stopping cost incurred by choosing $u_\tau = 1$ (stop and declare state 1) at time $k = \tau$. It is the error probability of declaring state e_1 when the actual state is e_2 . In terms of the public belief,

$$\begin{aligned} & J_\mu(\pi_0) \\ &= \mathbb{E}_{\pi_0}^\mu \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \bar{C}(\pi_{k-1}, u_k=2) + \rho^{\tau-1} \bar{C}(\pi_{\tau-1}, u_\tau=1) \right\} \\ & \bar{C}(\pi, 2) = d e_1' \pi, \quad \bar{C}(\pi, 1) = \beta e_2' \pi. \end{aligned} \quad (60)$$

The global decision $u_k = \mu(\pi_{k-1}) \in \{1(\text{stop}), 2(\text{continue})\}$ is a function of the public belief π_{k-1} updated according to the social learning protocol (52), (55). The optimal policy $\mu^*(\pi)$ and value function $V(\pi)$ satisfy Bellman's (25) with

$$\begin{aligned} Q(\pi, 2) &= C(\pi, 2) + \rho \sum_{a \in \mathbb{A}} V(T(\pi, a)) \sigma(\pi, a) \text{ where} \\ C(\pi, 2) &= \bar{C}(\pi, 2) - (1 - \rho) \bar{C}(\pi, 1), \quad Q(\pi, 1) \\ &= C(\pi, 1) = 0. \end{aligned} \quad (61)$$

Here $T(\pi, a)$ and $\sigma(\pi, a)$ are obtained from the social learning Bayesian filter (55).

Since $\mathbb{X} = \{1, 2\}$, the public belief state $\pi = [1 - \pi(2), \pi(2)]'$ is parametrized by the scalar $\pi(2) \in [0, 1]$, i.e., $\Pi(X)$ is the interval $[0, 1]$. In order to state the main result, define the following four intervals which form a partition of the interval $[0, 1]$:

$$\begin{aligned} \mathcal{P}_l &= \{\pi(2) : \eta_l < \pi(2) \leq \eta_{l-1}\}, \quad l = 1, \dots, 4 \text{ where} \\ \eta_0 &= 1 \\ \eta_1 &= \frac{(c(e_1, 2) - c(e_1, 1))B_{11}}{(c(e_1, 2) - c(e_1, 1))B_{11} + (c(e_2, 1) - c(e_2, 2))B_{21}} \\ \eta_2 &= \frac{(c(e_1, 2) - c(e_1, 1))}{(c(e_1, 2) - c(e_1, 1)) + (c(e_2, 1) - c(e_2, 2))} \\ \eta_3 &= \frac{(c(e_1, 2) - c(e_1, 1))B_{12}}{(c(e_1, 2) - c(e_1, 1))B_{12} + (c(e_2, 1) - c(e_2, 2))B_{22}} \\ \eta_4 &= 0. \end{aligned}$$

Note that η_0 corresponds to belief state e_2 , and η_4 corresponds to belief state e_1 . (See discussion at the end of this section for more intuition about the intervals \mathcal{P}_i).

It is readily verified that if the observation matrix B is TP2, then $\eta_3 \leq \eta_2 \leq \eta_1$. The following is the main result.

Theorem 8: Consider the stopping time problem (60) where agents perform social learning using the social learning Bayesian filter (55). Assume (57), $d \geq \rho\beta$, and B is TP2 symmetric. Then the optimal stopping policy $\mu^*(\pi)$ has the following structure: The stopping set \mathcal{R}_1 is the union of at most

three intervals. That is $\mathcal{R}_1 = \mathcal{R}_1^a \cup \mathcal{R}_1^b \cup \mathcal{R}_1^c$ where $\mathcal{R}_1^a, \mathcal{R}_1^b, \mathcal{R}_1^c$ are possibly empty intervals. Here:

- (i) The stopping interval $\mathcal{R}_1^a \subseteq \mathcal{P}_1 \cup \mathcal{P}_4$ and is characterized by a threshold point. That is, if \mathcal{P}_1 has a threshold point π^* , then $\mu^*(\pi) = 1$ for all $\pi(2) \in \mathcal{P}_4$ and

$$\mu^*(\pi) = \begin{cases} 2, & \text{if } \pi(2) \geq \pi^*, \pi(2) \in \mathcal{P}_1. \\ 1, & \text{otherwise} \end{cases}, \quad (62)$$

Similarly, if \mathcal{P}_4 has a threshold point π_4^* , then $\mu^*(\pi) = 2$ for all $\pi(2) \in \mathcal{P}_1$.

- (ii) The stopping intervals $\mathcal{R}_1^b \subseteq \mathcal{P}_2$ and $\mathcal{R}_1^c \subseteq \mathcal{P}_3$
- (iii) The intervals \mathcal{P}_1 and \mathcal{P}_4 are regions of information cascades. That is, if $\pi_k \in \mathcal{P}_1 \cup \mathcal{P}_4$, then social learning ceases and $\pi_{k+1} = \pi_k$ (see Theorem 7 for definition of information cascade). ■

The proof of Theorem 8 is in Appendix G. The proof depends on properties of the social learning filter and these are summarized in Lemma 7 in Appendix G.

Examples:

- (i) To illustrate the multiple threshold structure of the above theorem, consider the stopping time problem (60) with the following parameters:

$$\begin{aligned} \rho &= 0.9, \quad d = 1.8, \quad B = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix} \\ c(e_i, a) &= \begin{bmatrix} 4.57 & 5.57 \\ 2.57 & 0 \end{bmatrix}, \quad \beta = 2. \end{aligned} \quad (63)$$

Figs. 4(a) and 4(b) show the optimal policy and value function. These were computed by constructing a grid of 500 values for $\Pi(X) = [0, 1]$. The double threshold behavior of the stopping time problem when agents perform social learning is due to the discontinuous dynamics of the Bayesian social learning filter (55).

- (ii) Consider the following generalization of the sequential detection problem (60). In addition to the delay and error probability costs, we consider the total social learning cost incurred by all the agents. So now instead of (59) we have (64) at the bottom of the page.

Here $\mathcal{G}_k, \mathcal{H}_k$ are defined in (58). The first and last terms above constitute the total social learning cost (53) from time 1 to τ . In terms of the public belief

$$\begin{aligned} & \mathbb{E}\{\min_a \mathbb{E}\{c(x, a) | \mathcal{H}_k\} | \mathcal{G}_{k-1}\} \\ &= \sum_{y \in \mathbb{Y}} \min_a c'_a T(\pi_{k-1}, y) \sigma(\pi_{k-1}, y) \\ &= \sum_{y \in \mathbb{Y}} \min_a c'_a B_y \pi_{k-1} \end{aligned}$$

$$\begin{aligned} J_\mu(\pi_0) &= \mathbb{E}_{\pi_0}^\mu \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \mathbb{E} \left\{ \min_a \mathbb{E}\{c(x, a) | \mathcal{H}_k\} + dI(x = e_1) \middle| \mathcal{G}_{k-1} \right\} \right. \\ & \quad \left. + \rho^{\tau-1} \beta \mathbb{E}\{I(x \neq e_1) | \mathcal{G}_{\tau-1}\} + \rho^{\tau-1} \mathbb{E}\{\min_a \mathbb{E}\{c(x, a) | \mathcal{H}_\tau\} | \mathcal{G}_{\tau-1}\} \right\} \end{aligned} \quad (64)$$

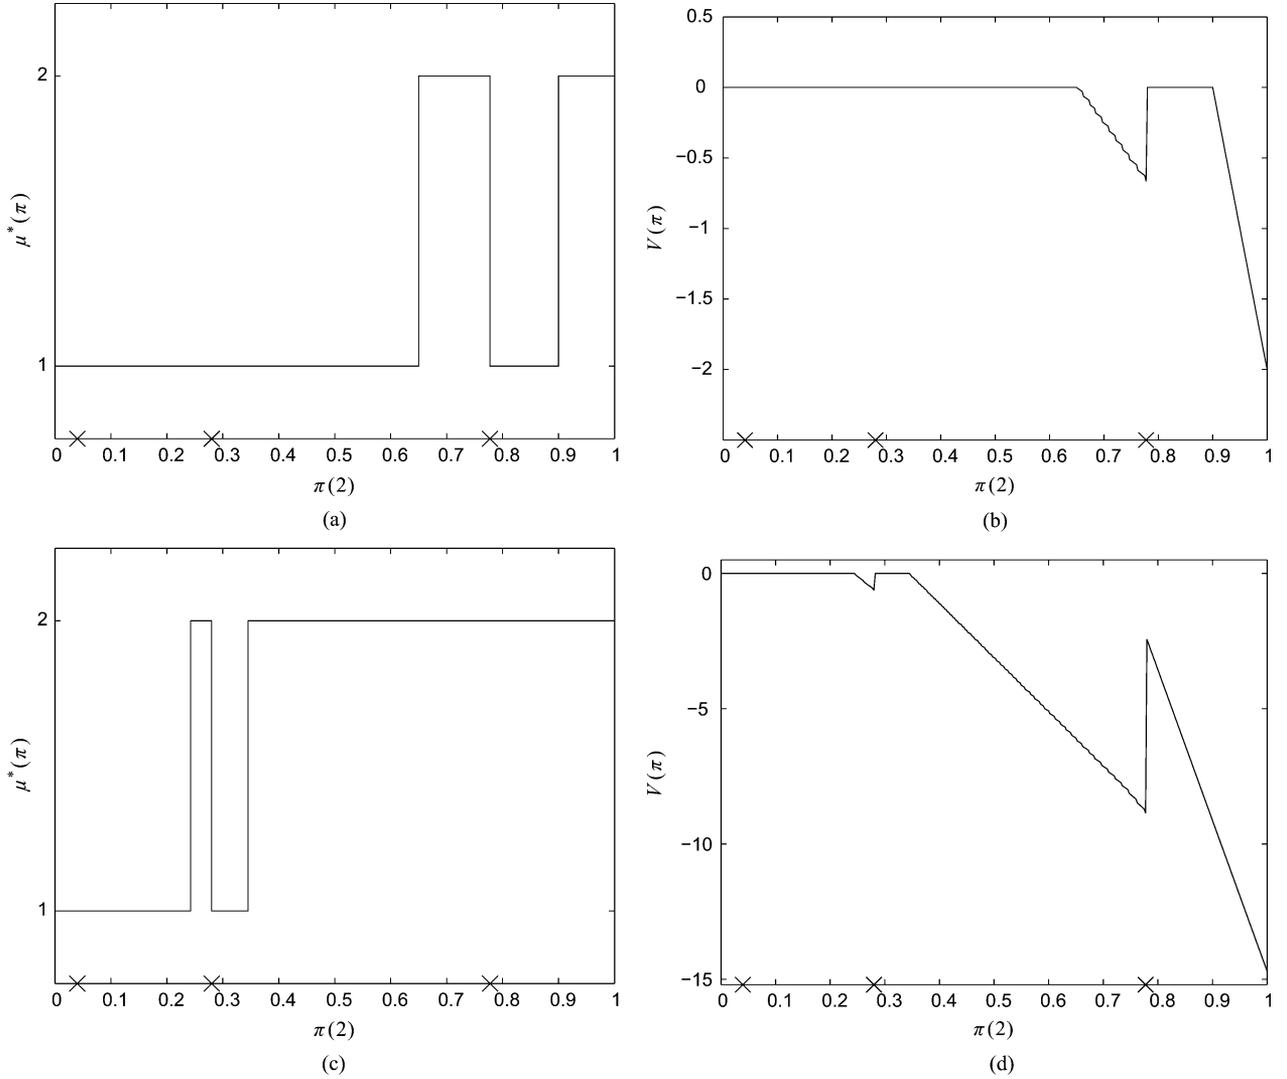


Fig. 4. Double Threshold Policy in stopping time problem involving social learning. The parameters are specified in (63) and (65). (a) and (c) show the optimal policy $\mu^*(\pi)$ and (b) and (d) show the value function $V(\pi)$ computed using (61).

Then it can be shown that Theorem 7 holds providing $c(e_i, a)$ is decreasing in i . We chose the following parameters: Last term in (64) set to zero

$$\rho = 0.9, d = 1, B = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}$$

$$c(e_i, a) = \begin{bmatrix} 2.1 & 3.1 \\ 3.1 & 0.53 \end{bmatrix}, \beta = 20. \quad (65)$$

Figs. 4(c) and 4(d) show the optimal policy and value function. The optimal stopping policy is again a double threshold, and the value function is monotone on individual intervals.

Discussion: The multiple threshold behavior (nonconvex stopping set \mathcal{R}_1) of Theorem 8 is unusual. One would have thought that if it was optimal to “continue” for a particular belief $\pi^*(2)$, then it should be optimal to continue for all beliefs $\pi(2)$ larger than $\pi^*(2)$. The multiple threshold optimal policy shows that this is not true. Fig. 4(a) shows that as the public belief $\pi(2)$ of state 2 decreases, the optimal decision switches from “continue” to “stop” to “continue” and finally “stop”.

Thus, the global decision (stop or continue) is a nonmonotone function of public beliefs obtained from local decisions.

The main reason for this unusual behavior is the dependence of the action likelihood R_a^π on the belief state π . This causes the social learning Bayesian filter to have a discontinuous update. The value function is no longer concave on $\Pi(X)$ and the optimal policy is not necessarily monotone. As shown in the proof of Theorem 7, the value function $V(\pi)$ is concave on each of the intervals $\mathcal{P}_l, l = 1, \dots, 4$.

To explain the final claim of the theorem, let us define the intervals \mathcal{P}_1 and \mathcal{P}_4 more explicitly

$$\mathcal{P}_1 = \{\pi : \arg\min_a c_a' B_y \pi = 2, \forall y \in \mathbb{Y}\}$$

$$\mathcal{P}_4 = \{\pi : \arg\min_a c_a' B_y \pi = 1, \forall y \in \mathbb{Y}\}. \quad (66)$$

For public belief $\pi \in \mathcal{P}_1$, the optimal local action is $a = 2$ irrespective of the observation y ; similarly for $\pi \in \mathcal{P}_4$, the optimal local action is $a = 1$ irrespective of the observation y . Therefore, on intervals \mathcal{P}_1 and \mathcal{P}_4 , there is no social learning since the

local action a reveals nothing about the observation y to subsequent agents. Social learning only takes place when the public belief is in \mathcal{P}_2 and \mathcal{P}_3 .

Finally, we comment on the intervals \mathcal{P}_l , $l = 1, \dots, 4$. They form a partition of $\Pi(X)$ such that if $\pi \in \mathcal{P}_l$, then $T(\pi, 1) \in \mathcal{P}_{l+1}$ and $T(\pi, 2) \in \mathcal{P}_{l-1}$ (with obvious modifications for $l = 1$ and $l = 4$), see Lemma 7 in the appendix for details. In fact, η_1 and η_3 are fixed points of the composition Bayesian maps: $\eta_1 = T(T(\eta_1, 1), 2)$ and $\eta_3 = T(T(\eta_3, 2), 1)$. Given that the updates of the social Bayesian filter can be localized to specific intervals, we can then inductively prove that the value function is concave on each such interval. This is the main idea behind Theorem 8.

C. Example 5: Constrained Social Optimum and Sequential Detection

In this subsection, we consider the constrained social optimum formulation in Chamley [11, Chapter 4.5]. We show in Theorem 9 that the resulting stopping time problem has a threshold switching curve. Thus, Chamley's optimal social learning can be implemented efficiently in a multi-agent system. This is in contrast to the multithreshold behavior of the stopping time problem in Sec. Section VI-B when agents were selfish in choosing their local actions.

The constrained social optimum formulation in [11] is motivated by the following question:⁷ How can agents aid social learning by acting benevolently and choosing their action to sacrifice their local cost but optimize a social welfare cost? In Section VI-B, agents ignore the information benefit their action provides to others resulting in information cascades where social learning stops. By constraining the choice by which agents pick their local action to two specific decision rules, the optimal choice between the two rules becomes a sequential decision problem.

Assume $\mathbb{A} = \mathbb{Y}$. As in the social learning model above, let c_a denote the cost vector for picking local action a and η_k [see (52)] denote the private belief of agent k .

Let τ denote a stopping time adapted to the sequence of sigma-algebras \mathcal{G}_k , $k \geq 1$ defined in (58). As in Section VI-B, the goal is to solve the following sequential detection problem to detect state e_1 : Pick the stopping time τ to minimize (67) at the bottom of the page. Similar to (64), the second and third terms are the delay cost and error probability in stopping and

⁷The author gratefully acknowledges discussions with Dr C. Chamley who authored the book [11]. The constrained social optimum formulation presented in this section is based on [11]. Our formulation considers a multistate version together with sequential detection of state e_1 .

announcing state e_1 . The first and last terms model the total social welfare cost involving all agents based on their local action. Let us explain these two terms. The key difference compared to (64) is that agents now pick their local action according to the decision rule $a(\pi, y, \mu(\pi))$ (instead of myopically) as follows. As in [11], we constrain decision rule $a(\pi, y, \mu(\pi))$ to two possible modes:

$$a_k = a(\pi_{k-1}, y_k, u_k) = \begin{cases} y_k, & \text{if } u_k = \mu(\pi_{k-1}) \\ & = 2 \text{ (reveal observation)} \\ \arg \min_a c'_a \pi_{k-1}, & \text{if } u_k = \mu(\pi_{k-1}) \\ & = 1 \text{ (stop)}. \end{cases} \quad (68)$$

Here the stationary policy $\mu : \pi_{k-1} \rightarrow u_k$ specifies which one of the two modes the benevolent agent k chooses. In mode $u_k = 2$, the agent k sacrifices immediate cost $c(x, a_k)$ and picks action $a_k = y_k$ to reveal full information to subsequent agents, thereby enhancing social learning.

In mode $u_k = 1$ the agent "stops and announces state 1". Equivalently, using the terminology of [11], the agent "herds" in mode $u_k = 1$. It ignores its private observation y_k , and chooses its action selfishly to minimize its cost given the public belief π_{k-1} . So agent k chooses $a_k = \arg \min_a c'_a \pi_{k-1} = a_{k-1}$. Then clearly from (56), $P(a|e_x, \pi)$ is functionally independent of x since $P(y|x = e_i)$ is independent of i . Therefore, from (55), if agent k herds, then $\pi_k = T(\pi_{k-1}, a_k) = \pi_{k-1}$, i.e., the public belief remains frozen. The total cost incurred in herding is then equivalent to final term in (67)

$$\begin{aligned} \sum_{k=\tau}^{\infty} \rho^{k-1} \min_{a \in \mathbb{Y}} \mathbb{E}\{c(x, a) | \mathcal{G}_{k-1}\} &= \sum_{k=\tau}^{\infty} \rho^{k-1} \min_a c'_a \pi_{\tau-1} \\ &= \frac{\rho^{\tau-1}}{1-\rho} \min_a c'_a \pi_{\tau-1}. \end{aligned}$$

Define the *constrained social optimal policy* μ^* such that $J_{\mu^*}(\pi_0) = \inf_{\mu \in \mathbb{M}} J_{\mu}(\pi_0)$. The sequential stopping problem (67) seeks to determine the optimal policy μ^* to achieve the optimal tradeoff between stopping and announcing state 1 and the cost incurred by agents that are acting benevolently. In analogy to Theorem 1, we show that $\mu^*(\pi)$ is characterized by a threshold curve.

Similar to (24) define the costs in terms of the belief state as

$$\begin{aligned} C(\pi, 1) &= \frac{1}{1-\rho} \min_{a \in \mathbb{Y}} c'_a \pi, \quad C(\pi, 2) \\ &= \sum_{y \in \mathbb{Y}} c'_y B_y \pi + (d + (1-\rho)\beta) e'_1 \pi - (1-\rho)\beta. \end{aligned} \quad (69)$$

$$\begin{aligned} J_{\mu}(\pi_0) &= \mathbb{E}_{\pi_0}^{\mu} \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} \mathbb{E}\{c(x, a_k) | \mathcal{G}_{k-1}\} + \mathbb{E}_{\pi_0}^{\mu} \left\{ \sum_{k=1}^{\tau-1} \rho^{k-1} dI(x = e_1) | \mathcal{G}_{k-1} \right\} \right. \\ &\quad \left. + \rho^{\tau-1} \mathbb{E}_{\pi_0}^{\mu} \{ \beta I(x \neq e_1) | \mathcal{G}_{\tau-1} \} + \frac{\rho^{\tau-1}}{1-\rho} \min_{a \in \mathbb{Y}} \mathbb{E}\{c(x, a) | \mathcal{G}_{\tau-1}\} \right\} \end{aligned} \quad (67)$$

Below we list the assumptions and main structural result which is similar to Theorem 1. These assumptions involve the social learning cost $c(e_i, a)$ and are not required if these costs are zero.

- (i) (A1-Ex5) $c(e_i, a) - c(e_{i+1}, a) \geq 0$
- (ii) (S-Ex5) (i) $c(e_X, a) - c(e_i, a) \geq (1 - \rho) \sum_y (c(e_X, y)B_{Xy} - c(e_i, y)B_{iy})$.
- (iii) $(1 - \rho) \sum_y (c(e_1, y)B_{1y} - c(e_i, y)B_{iy}) \geq c(e_1, a) - c(e_i, a)$.

Similar to the discussion in Section III-C, (A1-Ex5) is sufficient for $C(\pi, 1)$ and $C(\pi, 2)$ to be \geq_r decreasing in $\pi \in \Pi(X)$. This implies that the costs $C(e_i, u)$ are decreasing in i , i.e., state 1 is the most costly state.

(S-Ex5) is sufficient for $C(\pi, u)$ to be submodular. It implies that $C(e_i, 2) - C(e_i, 1)$ is decreasing in i . This gives economic incentive for agents to herd when approaching the state e_1 , since the differential cost between continuing and stopping is largest for e_1 . Intuitively, the decision to stop (herd) should be made when the state estimate is sufficiently accurate so that revealing private observations is no longer required.

Theorem 9: Consider the sequential detection problem for state e_1 with social welfare cost in (67) and constrained decision rule (68). Then:

- (i) Under (A1-Ex5), (A2), (S-Ex5), constrained social optimal policy $\mu^*(\pi)$ satisfies the structural properties of Theorem 1. (Thus, a threshold switching curve exists).
- (ii) The stopping set \mathcal{R}_1 is the union of $|\mathbb{Y}|$ convex sets (where $|\mathbb{Y}|$ denotes cardinality of \mathbb{Y}). Note also that \mathcal{R}_1 is a connected set by Statement (i). (Recall $\mathbb{Y} = \mathbb{A}$ in our formulation).

■

The proof of Theorem 9 is in Appendix H. The main implication of Theorem 9 is that the constrained optimal social learning scheme formulated in Chamley [11] has a monotone structure. This is in contrast to the multithreshold behavior of the stopping time problem in Section VI-B when agents were selfish in picking their local actions. In [11, Chapter 4.5], the above formulation is used for pricing information externalities in social learning. From an implementation point of view, the existence of a threshold switching curve implies that the protocol only needs individual agents to store the optimal linear MLR policy (computed, for example, using Algorithm 1). Finally, \mathcal{R}_1 is the union of $|\mathbb{Y}|$ convex sets and is nonconvex in general. This is different to standard stopping problems where the stopping set is convex.

VII. EXAMPLE 6: MULTI-AGENT SCHEDULING IN A CHANGING WORLD

So far we have considered models where the underlying state is a constant Section VI, or a Markov chain that jumps once into an absorbing state (change detection of Section III or jumps twice (transient detection of Section IV). In this section, we consider a more general model where the target state x evolves on the same time scale as the observation process. The target state x jumps with time according to a finite state Markov chain over the state space \mathbb{X} with transition probability matrix P . Also, unlike previous sections, decision $u = 1$ does not “stop” the evolu-

tion of the belief state. So instead of a stopping time problem, we have a more general partially-observed stochastic control problem.

As mentioned in Section I, this section is motivated by two questions: (i) How can the optimal policy be bounded? (ii) How does the optimal achievable cost vary with transition probability? The main results of this section are two-fold. First, using Blackwell dominance, we show that the optimal policy is lower bounded by a myopic policy (Theorem 10). Next, Theorem 11 shows that for the underlying Markovian state, the larger the transition matrix (in an order defined in (75)), the cheaper the expected optimal cost. This is useful in comparing the optimal achievable cost of quickest time detection with different PH-distributions.

A. Myopic Policy Bound to Optimal Decision Policy

Consider a countable number of agents where each agent acts once in a predetermined sequential order indexed by $k = 1, 2, \dots$ as follows: Based on the current belief state π_{k-1} , agent k chooses mode

$$u_k \in \{1 \text{ (low resolution)}, 2 \text{ (high resolution)}\}.$$

Depending on its mode u_k , agent k views the world according to this mode—that is, it obtains observation from a distribution that depends on u_k . Assume that for mode $u \in \{1, 2\}$, the observation $y^{(u)} \in \mathbb{Y}^{(u)} = \{1, \dots, Y^{(u)}\}$ is obtained from the matrix of conditional probabilities

$$B^{(u)} = \left(B_{iy^{(u)}}^{(u)}, i \in \mathbb{X}, y^{(u)} \in \mathbb{Y}^{(u)} \right) \text{ where} \\ B_{iy^{(u)}}^{(u)} = P \left(y^{(u)} | x = e_i, u \right). \quad (70)$$

The notation $\mathbb{Y}^{(u)}$ allows for mode dependent observation spaces. In sensor scheduling [21], the tradeoff is as follows: Mode $u = 2$ yields more accurate observations of the state than mode $u = 1$, but the cost of choosing mode $u = 2$ is higher than mode $u = 1$. Thus, there is a tradeoff between the cost of acquiring information and the value of the information. The assumption that mode $u = 2$ yields more accurate observations than mode $u = 1$ is modelled by

$$B^{(1)} = B^{(2)}Q, \quad (71)$$

Here Q is a $Y^{(2)} \times Y^{(1)}$ stochastic matrix. Q can be viewed as a *confusion matrix* that maps $\mathbb{Y}^{(2)}$ probabilistically to $\mathbb{Y}^{(1)}$. (In a communications context, one can view Q as a noisy discrete memoryless channel with input $y^{(2)}$ and output $y^{(1)}$).

When agent k chooses mode $u_k \in \{1, 2\}$, it incurs the expected cost

$$C(\pi_{k-1}, u_k) = \alpha_{u_k} \mathbb{E} \{ \| (x_k - \mathbb{E}\{x_k | \mathcal{F}_k\})' g \|^2 | \mathcal{F}_{k-1} \} \\ + \mathbb{E} \{ c(x_k, u_k) | \mathcal{F}_{k-1} \} \\ = \alpha_{u_k} (G' P' \pi_{k-1} - (g' P' \pi_{k-1})^2) + c'_{u_k} P' \pi_{k-1} \quad (72)$$

where $c_u = (c(x = e_i, u), i \in \mathbb{X})$, \mathcal{F}_k is defined in (10), g and G are defined in (14). In (72), the tradeoff between information obtained from a mode and the cost of operating in the mode is

modelled as follows: Choose $\alpha_1 > \alpha_2$ to penalize choosing the less accurate mode 1 in terms of the variance, while $c(e_i, 1) < c(e_i, 2)$ since mode 1 incurs a cheaper operating cost.

The goal is to compute the optimal policy $\mu^*(\pi) \in \{1, 2\}$ to minimize the overall cost incurred by all the agents

$$J_\mu(\pi_0) = \mathbb{E}_{\pi_0}^\mu \left\{ \sum_{k=1}^{\infty} \rho^{k-1} C(\pi_{k-1}, u_k) \right\}. \quad (73)$$

The above problem is not a stopping time problem, since if mode $u = 1$ is chosen, the problem does not terminate. The mode u chosen by each agent will affect the modes chosen by subsequent agents, and hence affects the total cost. For such a partially observed stochastic control problem, determining the optimal policy $\mu^*(\pi)$ is computationally intractable. However, using Blackwell dominance, we show below that a myopic policy forms a lower bound for the optimal policy.

The value function $V(\pi)$ and optimal policy $\mu^*(\pi)$ satisfy the dynamic programming equation

$$\begin{aligned} V(\pi) &= \min_{u \in \mathcal{U}} Q(\pi, u) \\ \mu^*(\pi) &= \arg \min_{u \in \mathcal{U}} Q(\pi, u) \\ J_{\mu^*}(\pi) &= V(\pi) \\ Q(\pi, u) &= C(\pi, u) + \rho \sum_{y^{(u)} \in \mathbb{Y}^{(u)}} V\left(T(\pi, y^{(u)})\right) \sigma(\pi, y^{(u)}) \\ T(\pi, y^{(u)}) &= \frac{B_{y^{(u)}} P' \pi}{\sigma(\pi, a)} \\ \sigma(\pi, y^{(u)}) &= \mathbf{1}'_X B_{y^{(u)}} P' \pi. \end{aligned} \quad (74)$$

We now present the structural result. Let $\Pi^s \subset \Pi(X)$ denote the set of belief states for which $C(\pi, 2) < C(\pi, 1)$. Define the myopic policy

$$\bar{\mu}(\pi) = \begin{cases} 2 & \pi \in \Pi^s \\ 1 & \text{otherwise} \end{cases}.$$

Theorem 10: The myopic policy $\bar{\mu}(\pi)$ satisfies the following property: For all $\pi \in \Pi^s$, $\mu^*(\pi) = \bar{\mu}(\pi)$, i.e., it is optimal to pick action 2. Therefore, $\bar{\mu}(\pi)$ forms a lower bound for the optimal policy $\mu^*(\pi)$, i.e., $\mu^*(\pi) \geq \bar{\mu}(\pi)$ for all $\pi \in \Pi(X)$. ■

Theorem 10 is proved in Appendix A. The usefulness of Theorem 10 stems from the fact that $\bar{\mu}(\pi)$ is trivial to compute. It forms a rigorous lower bound to the computationally intractable optimal policy $\mu^*(\pi)$. What the theorem says is that $\bar{\mu}(\pi)$ lower bounds the optimal policy, and coincides with the optimal policy in region Π^s . Since $\bar{\mu}$ is sub-optimal, it incurs a higher cost. This cost can be evaluated via simulation and forms an upper bound to the optimal achievable cost.

Theorem 10 is nontrivial. Just because at some time k , the expected instantaneous costs satisfy $C(\pi_k, 2) < C(\pi_k, 1)$, does not necessarily imply that the myopic policy $\bar{\mu}(\pi)$ coincides with the optimal policy $\mu^*(\pi)$, since the optimal policy applies to the infinite trajectory of the dynamical system.

The proof uses Blackwell dominance of measures. The first instance of a similar proof using Blackwell dominance for POMDPs was given in [49], see also [41]. Our proof is similar and uses concavity of the value function $V(\pi)$ and Blackwell

dominance of observation probabilities. In particular, observation $y^{(2)}$ is more informative than (Blackwell dominates) observation $y^{(1)}$, if (71) holds, see [41]. The proof of Theorem 10 in the appendix comprises of first proving concavity of $V(\pi)$. The proof is a non trivial extension, since $C(\pi, 1)$ is nonlinear in π .

B. Effect of State Transition Matrix

The next structural result establishes how the optimal expected cost varies with different transition matrices. The model we consider applies to both stopping time problems (such as quickest time detection) and multi-agent scheduling considered above.

Suppose $P^{(1)}, P^{(2)}$ are two distinct transition probability matrices corresponding to two distinct models of Markov state evolution. Let $V(\pi; P^{(1)})$ and $V(\pi; P^{(2)})$ denote the corresponding optimal value function in (74). The question we pose is: How does $V(\pi; P)$ vary with transition matrix P ? For example, in the quickest detection problem, do certain phase-type distributions result in larger total optimal cost compared to other phase-type distributions? A similar question can be posed for the stochastic control problem considered above.

We consider costs that are linear in the belief state. To show the explicit dependence on the transition matrix, define the costs (recall $c_u = (c(e_i, u), i \in \mathbb{X})$)

$$C(\pi, u; P) = c'_u P' \pi, \quad u \in \{1, 2\}.$$

The result below also applies to the case where $C(\pi, u) = c'_u \pi$ (i.e., the cost at each stage is not an explicit function of transition matrix).

Define the following ordering of transition matrices $P^{(1)}$ and $P^{(2)}$:

$$P^{(1)} \succeq P^{(2)} \text{ if } P_{ij}^{(1)} P_{m,l}^{(2)} \leq P_{ij}^{(2)} P_{m,l}^{(1)}, \quad l > j, \quad i, j, l, m \in \mathbb{X}. \quad (75)$$

Theorem 11: Assume $P^{(1)} \succeq P^{(2)}$, $c(e_i, u)$ is decreasing in i , $C(\pi, u; P^{(1)}) \leq C(\pi, u; P^{(2)})$ and (A2), (A3) hold. Then:

- (i) $V(\pi; P^{(1)}) \leq V(\pi; P^{(2)})$. That is, the larger the transition matrix (with respect to the partial ordering (75)), the lower the optimal expected cost incurred when making optimal decisions.
- (ii) Consider the quickest time detection problem with $\alpha = 0$ and costs in (29). The optimal expected cost $\bar{V}(\pi; P)$ (see (23)) with change time distribution $P^{(1)}$ is less than that with $P^{(2)}$. ■

Theorem 11 is proved in Appendix J. We now present three examples.

Example (i): Consider $\mathbb{X} = \{1, 2\}$ and the dynamic decision making formulation of Sec. Section VII-A. Then using (75) it can be verified that the transition matrices

$$P^{(1)} = \begin{bmatrix} 0.2 & 0.8 \\ 0.1 & 0.9 \end{bmatrix} \succeq P^{(2)} = \begin{bmatrix} 0.8 & 0.2 \\ 0.7 & 0.3 \end{bmatrix}.$$

Note that $P^{(1)}$ and $P^{(2)}$ above are TP2 [as required by (A3)]. So Theorem 11 applies.

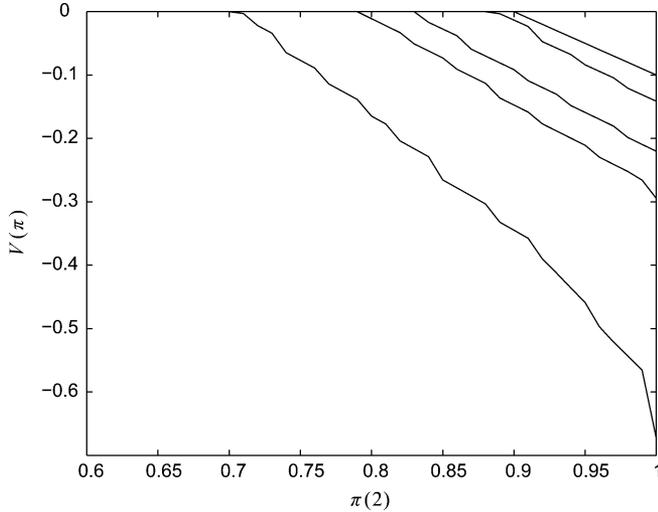


Fig. 5. Monotone behaviour of optimal expected cost (value function $V(\pi; P^{(p)})$) versus probability p for quickest detection with geometric change time, see (76) for notation. The values of p are 0.99 (lowest plot), 0.95, 0.9, 0.8, and 0.01 (highest plot).

Example (ii): The transition matrix $\begin{bmatrix} p & (1-p) \\ p & (1-p) \end{bmatrix}$ corresponding to a two state iid process is decreasing with respect to the order (29) as p increases from 0 to 1. So the theorem says that the smaller p is, the cheaper the optimal expected cost. So even though the underlying process has maximum uncertainty (entropy) when $p = 0.5$, Theorem 11 says that the largest total cost incurred is when $p = 1$.

Example (iii): Consider the quickest detection costs (29) with $\alpha = 0$. Consider first the geometric distributed change time case ($X = 2$) with transition matrix $P^{(p)}$ and parameters in (29)

$$P^{(p)} = \begin{bmatrix} 1 & 0 \\ 1-p & p \end{bmatrix}, \quad B = \begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}, \quad \rho = 0.9, \quad d = 0.9. \quad (76)$$

It can be verified that $P^{(p)}$ is increasing (in terms of (75)) in p . So Theorem 11 says that the larger p is, the smaller the average optimal cost (value function $V(\pi; P^{(p)})$) for quickest time detection. In Fig. 5 we plot the value function $V(\pi; P^{(p)})$ for several values of p , to illustrate this behavior.

Next, consider quickest detection with PH-distributed change times (7) modelled by the following transition matrices: in quickest detection:

$$P^{(1)} = \begin{bmatrix} 1 & 0 & 0 \\ 0.5 & 0.3 & 0.2 \\ 0.3 & 0.4 & 0.3 \end{bmatrix} \\ \succeq_{\text{TP2}} P^{(2)} = \begin{bmatrix} 1 & 0 & 0 \\ 0.9 & 0.1 & 0 \\ 0.8 & 0.15 & 0.05 \end{bmatrix}.$$

Since $P^{(1)}$ and $P^{(2)}$ are TP2 by (A3), Theorem 11 implies that the optimal expected cost incurred in quickest change detection with PH-distributed change time $P^{(1)}$ is less than that of $P^{(2)}$.

Discussion:

- (i) *TP2 Dominance Versus Dominance in (75)* It is shown in [28], [41] that if transition matrices are ordered in the TP2 sense, namely $P^{(1)} \succeq_{\text{TP2}} P^{(2)}$ (see Defn.4(i) in Appendix), then Theorem 11 holds under the same assumptions as above. It is easy to prove that for 2×2 case, only matrices with identical rows, i.e., transition matrices modelling independent and identically distributed (iid) finite state processes, satisfy $P^{(1)} \succeq_{\text{TP2}} P^{(2)}$. Our conjecture is that the only examples of transition matrices that satisfy $P^{(1)} \succeq_{\text{TP2}} P^{(2)}$ are transition matrices corresponding to iid processes. So TP2 dominance is less useful than the ordering (75).
- (ii) *Kolmogorov-Shiryayev criterion* If the Kolmogorov-Shiryayev criterion (22) is considered then a similar proof to Theorem 11 shows that $V(\pi; P)$ is increasing with P . The reason is that in this case $C(\pi, 2; P)$ in (24) (with $de'_1 P' \pi$ replaced by $de'_1 \pi$, see Theorem 2) has the term $-\rho(\alpha + \beta)e'_1 P' \pi$. This is increasing in P (wrt ordering (75)).

VIII. NUMERICAL EXAMPLES

For state-space $\mathbb{X} = \{1, 2, 3\}$, i.e., $X = 3$ states, the belief state space $\Pi(X)$ is an equilateral triangle, and the various results of this paper can be illustrated visually. There is much flexibility for choice of parameters that satisfy the general assumptions (A1), (A2), (A3), (S) in the appendix.

Example 1: We illustrate the structural result Theorem 1 for quickest time change detection with PH-distributed change time and variance penalty. The following parameters were chosen in (29):

$$\beta = 1, \quad \rho = 1, \quad P = \begin{bmatrix} 1 & 0 & 0 \\ 0.3 & 0.1 & 0.6 \\ 0 & 0.02 & 0.98 \end{bmatrix}.$$

Assume Gaussian observation noise with variance 0.01; so $\mathbb{Y} = \mathbb{R}$ and the observation likelihoods are $B_{1y} \sim N(0, 0.01)$, $B_{2y} \sim N(1, 0.01)$. The operational cost $c = 10^{-3}$ (see discussion below (17)).

The optimal policy was computed by forming a grid of 230 values in the 2-D unit simplex, and then solving the value iteration algorithm (27) over this grid and a horizon length of 200. The optimal policy is shown in Fig. 6(a) for four different choices of α and d . The first three examples, namely α, d specified in Figs. 6(a), 6(b) and 6(c), satisfy assumptions (A1-Ex1), (A2), (A3) and (S-Ex1). Therefore, the optimal policy $\mu^*(\pi)$ satisfies Theorem 1 and is characterized by a threshold curve $\Gamma(\pi)$. The figures clearly show the existence of a threshold curve that partitions $\Pi(X)$ into two individually connected regions. Recall that for the case $\alpha = 0$, Theorem 1 says that the optimal stopping set \mathcal{R}_1 is a convex set.

Is the stopping set \mathcal{R}_1 always a connected set even when the assumptions of Theorem 1 do not hold? Recall the assumptions in Theorem 1 are sufficient conditions. The parameters $\alpha = 10$, $d = 5$ in Fig. 6(d) does not satisfy condition (S-Ex1) in

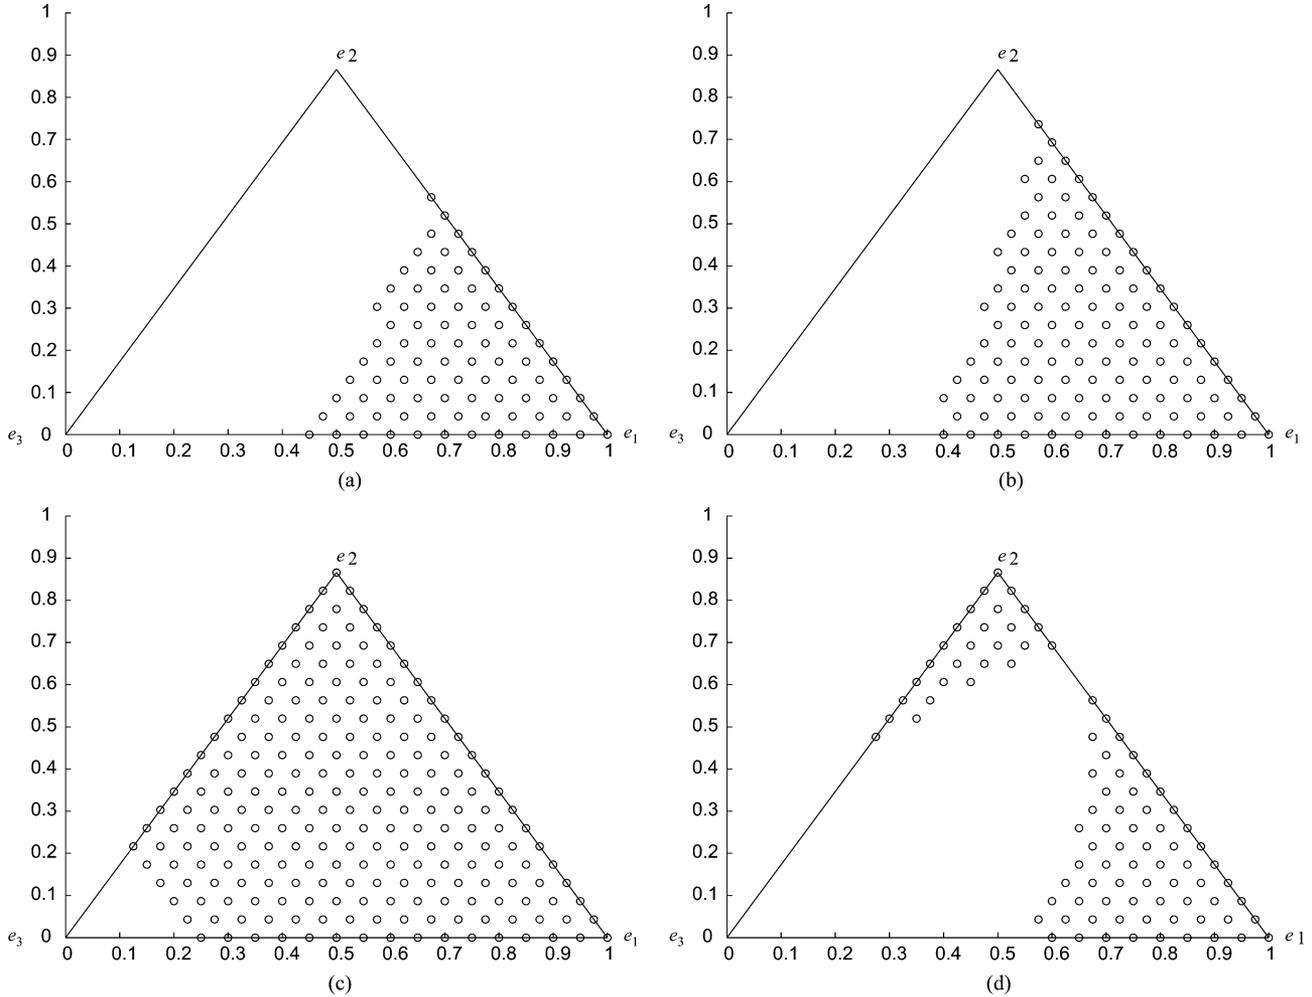


Fig. 6. Quickest Time Detection illustrating Theorem 1. The \circ region represents the stopping set \mathcal{R}_1 where decision $u = 1(\text{stop})$ is optimal. The empty region in the simplex represents \mathcal{R}_2 where $u = 2(\text{continue})$ is optimal. Cases (a), (b) and (c) satisfy the assumptions of Theorem 1. The optimal linear threshold for these 3 cases was estimated using Algorithm 1. Case (d) does not satisfy Assumptions (S-Ex1) and \mathcal{R}_1 is not a connected set.

Appendix B. As shown in Fig. 6(d), the optimal stopping set \mathcal{R}_1 is no longer connected. This highlights the importance of developing useful sufficient conditions that result in monotone policies $\mu^*(\pi)$, such as the assumptions presented in this paper.

Example 2: Here we consider the classical delay cost (18). We illustrate how the optimal stopping region \mathcal{R}_1 varies with transition probabilities of the PH-distribution for change time. Since all these constraints in \mathbf{f} are linear, determining feasible choices is straightforward using the Matlab command `linprog`.

We chose $B_{1y} \sim N(0, 4)$, $B_{2y} = N(1, 4)$ and the following parameters in (31):

$$\alpha = 0.5, \beta = 1, d = 1, \rho = 0.75, \mathbf{f} = \begin{bmatrix} 0 \\ 1 \\ 2 \end{bmatrix}$$

$$P = \begin{bmatrix} 1 & 0 & 0 \\ 0.3 & 0.6 & 0.1 \\ 0.1 & p & 0.9 - p \end{bmatrix}. \quad (77)$$

Then (A3) holds, i.e., P is TP2 for $p \in [0.2, 0.7714]$. Also it can be verified that all the other assumptions of Theorem 2 hold.

Figs. 7(a), 7(b) illustrate the optimal stopping region \mathcal{R}_1 for $p = 0.2$ and $p = 0.77$, respectively. Fig. 7(c) plots the PH-distribution probability mass function ν_k in (7) versus time k for the transition probabilities in Example 1 and Example 2 for $p = 0.2$ and $p = 0.77$. It can be seen that even with a 3-state Markov chain the behavior is quite different to a geometric distribution.

IX. CONCLUSIONS

This paper has presented structural results for Bayesian quickest time detection with PH-distributed change time and variance penalty. The main result is Theorem 1 which proves the existence of a threshold switching curve for optimal decision making under general assumptions (A1), (A2), (A3), (S) given in the appendix. Theorem 4 gave necessary and sufficient conditions for a linear threshold policy to approximate the threshold curve. Then several examples were considered, namely quickest transient detection, quickest time detection with exponential penalty, stopping time problems in social learning, constrained optimal social learning, and multi-agent scheduling in a changing world. In the case of exponential penalty we used a risk sensitive stochastic control formulation. In all these examples, under similar assumptions to

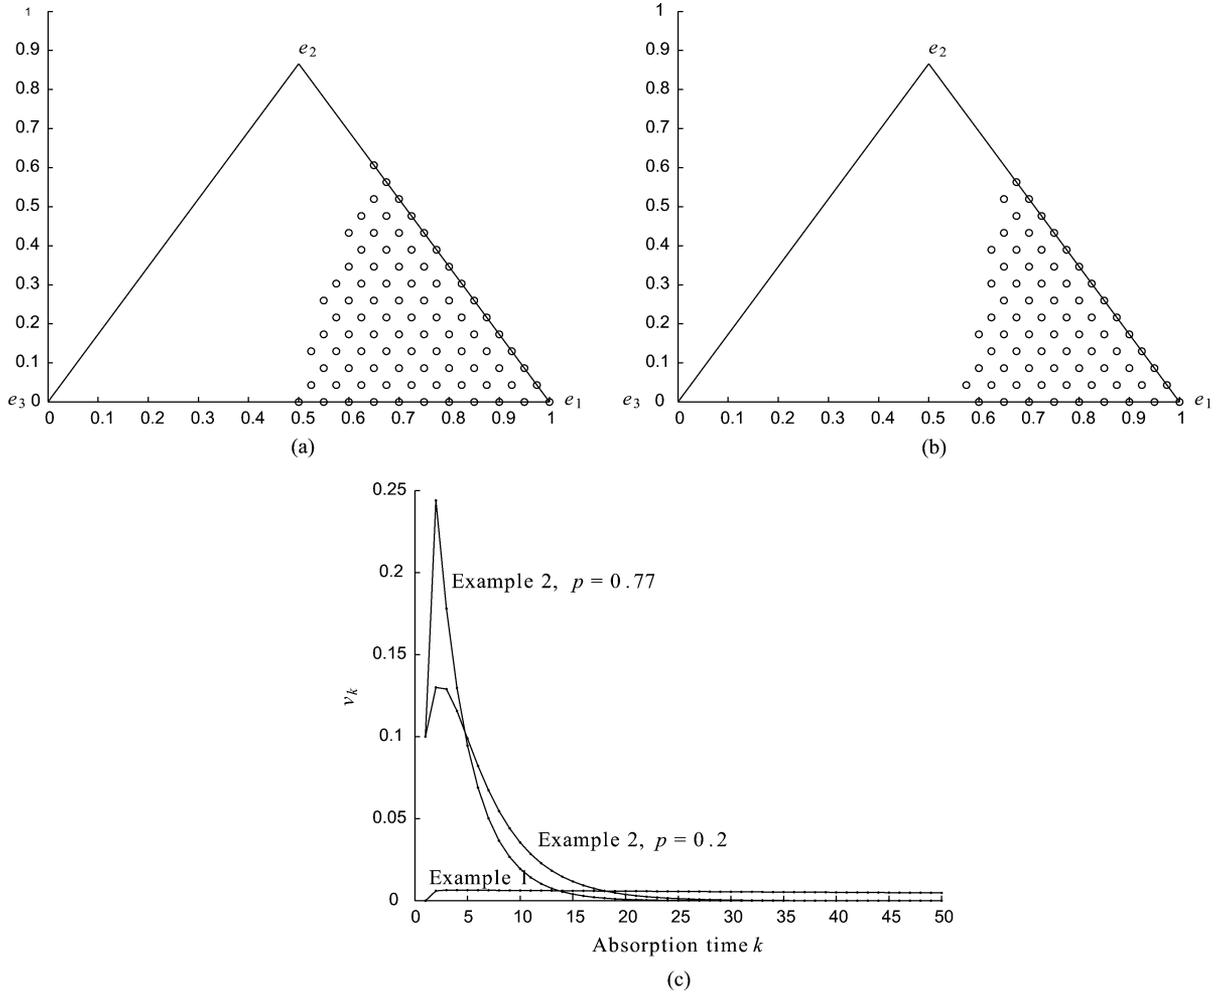


Fig. 7. Effect of PH-distribution probabilities (77) on optimal stopping region \mathcal{R}_1 . The region marked with \circ denotes \mathcal{R}_1 .

Theorem 1, the threshold switching curve holds. The proofs of the results use lattice programming and stochastic orders on the unit simplex. The structural results of this paper are class type results, that is, for parameters belonging to a set, the results hold. Hence, there is an inherent robustness in these results since even if the underlying parameters are not exactly specified but still belong to the appropriate sets, the results still hold. It would be useful to do a performance analysis of the various optimal detectors proposed in this paper—see [40] and references therein.

APPENDIX

A) *Stochastic Orders and Submodularity*: Theorem 1 below requires proving that the quickest time detection policy $\mu^*(\pi)$ is monotonically increasing in belief state π . That is, $\pi \leq \tilde{\pi}$ (in a sense to be made clear below), implies $\mu^*(\pi) \leq \mu^*(\tilde{\pi})$. In order to compare belief states π and $\tilde{\pi}$, we will use the monotone likelihood ratio (MLR) stochastic ordering and a specialized version of the MLR order restricted to lines in the simplex $\Pi(X)$. This stochastic order is useful since it is preserved under conditional expectations [18], [32], [41], [50]. Below we introduce several important definitions that will be used subsequently.

Definition 1 (MLR Ordering, [32]): Let $\pi_1, \pi_2 \in \Pi(X)$ be any two belief state vectors. Then π_1 is greater than π_2 with respect to the MLR ordering—denoted as $\pi_1 \geq_r \pi_2$, if

$$\pi_1(i)\pi_2(j) \leq \pi_2(i)\pi_1(j), \quad i < j, \quad i, j \in \{1, \dots, X\}. \quad (78)$$

Similarly $\pi_1 \leq_r \pi_2$ if \leq in (78) is replaced by a \geq .

Definition 2 (First Order Stochastic Dominance, [32]): Let $\pi_1, \pi_2 \in \Pi(X)$. Then π_1 first order stochastically dominates π_2 —denoted as $\pi_1 \geq_s \pi_2$ —if $\sum_{i=j}^X \pi_1(i) \geq \sum_{i=j}^X \pi_2(i)$ for $j = 1, \dots, X$.

Result 1 ([32]):

- (i) Let $\pi_1, \pi_2 \in \Pi(X)$. Then $\pi_1 \geq_r \pi_2$ implies $\pi_1 \geq_s \pi_2$.
- (ii) Let \mathcal{V} denote the set of all S -dimensional vectors v with nondecreasing components, i.e., $v_1 \leq v_2 \leq \dots \leq v_X$. Then $\pi_1 \geq_s \pi_2$ iff for all $v \in \mathcal{V}$, $v'\pi_1 \geq v'\pi_2$.

For state-space dimension $X = 2$, MLR is a complete order and coincides with first order stochastic dominance. For state-space dimension $X > 2$, MLR is a *partial order*, i.e., $[\Pi(X), \geq_r]$ is a partially ordered set (poset) since it is not always possible to order any two belief states $\pi \in \Pi(X)$. However, on line segments in the simplex defined below, MLR is a total ordering.

For $i \in \mathbb{X}$, define the sub simplex $\mathcal{H}_i \subset \Pi(X)$ as

$$\mathcal{H}_i = \{\pi \in \Pi(X) : \pi(i) = 0\}. \quad (79)$$

Denote belief states that lie in \mathcal{H}_i by $\bar{\pi}$. For each $\bar{\pi} \in \mathcal{H}_i$, construct the line segment $\mathcal{L}(e_i, \bar{\pi})$ that connects $\bar{\pi}$ to e_i . Thus, $\mathcal{L}(e_i, \bar{\pi})$ comprises of belief states π of the form

$$\begin{aligned} \mathcal{L}(e_i, \bar{\pi}) &= \{\pi \in \Pi(X) : \pi = (1 - \epsilon)\bar{\pi} + \epsilon e_i, 0 \leq \epsilon \leq 1\} \\ \bar{\pi} &\in \mathcal{H}_i. \end{aligned} \quad (80)$$

Definition 3 (MLR Ordering \geq_{L_i} on Lines): π_1 is greater than π_2 with respect to the MLR ordering on the line $\mathcal{L}(e_i, \bar{\pi})$ —denoted as $\pi_1 \geq_{L_i} \pi_2$, if $\pi_1, \pi_2 \in \mathcal{L}(e_i, \bar{\pi})$ for some $\bar{\pi} \in \mathcal{H}_i$, i.e., π_1, π_2 are on the same line connected to vertex e_i of simplex $\Pi(X)$, and $\pi_1 \geq_r \pi_2$.

Note that $[\Pi(X), \geq_{L_X}]$ and $[\Pi(X), \geq_{L_1}]$ are chains,⁸ i.e., all elements $\pi, \tilde{\pi} \in \mathcal{L}(e_X, \bar{\pi})$ are comparable, i.e., either $\pi \geq_{L_X} \tilde{\pi}$ or $\tilde{\pi} \geq_{L_X} \pi$ (and similarly for $\mathcal{L}(e_1, \bar{\pi})$) In Lemma 2, we summarize useful properties of $[\Pi(x), \geq_{L_i}]$ that will be used in our proofs.

Lemma 2: The following properties hold on $[\Pi(X), \geq_r]$, $[\mathcal{L}(e_X, \bar{\pi}), \geq_{L_X}]$.

- (i) On $[\Pi(X), \geq_r]$, e_1 is the least and e_X is the greatest element. On $[\mathcal{L}(e_X, \bar{\pi}), \geq_{L_i}]$, $\bar{\pi}$ is the least and e_X is the greatest element.
- (ii) Convex combinations of MLR comparable belief states form a chain. For any $\gamma \in [0, 1]$, $\pi \leq_r \tilde{\pi} \implies \pi \leq_r \gamma\pi + (1 - \gamma)\tilde{\pi} \leq_r \tilde{\pi}$. (iii) All points on a line $\mathcal{L}(e_X, \bar{\pi})$ are MLR comparable. Consider any two points $\pi^{\gamma_1}, \pi^{\gamma_2} \in \mathcal{L}(e_X, \bar{\pi})$ (80). Then $\gamma_1 \geq \gamma_2$, implies $\pi^{\gamma_1} \geq_{L_i} \pi^{\gamma_2}$.

Let $\mathbf{i} = (i_1, \dots, i_L)$ and $\mathbf{j} = (j_1, \dots, j_L)$ denote the indices of two L -variate probability mass functions Denote the lattice operators

$$\begin{aligned} \mathbf{i} \wedge \mathbf{j} &= [\min(i_1, j_1), \dots, \min(i_L, j_L)]' \\ \mathbf{i} \vee \mathbf{j} &= [\max(i_1, j_1), \dots, \max(i_L, j_L)]'. \end{aligned} \quad (81)$$

Definition 4 (TP2 Ordering and Reflexive TP2 Distributions): Let P and Q denote any two L -variate probability mass functions. Then:

- (i) $P \geq_{\text{TP2}} Q$ if $P(\mathbf{i})Q(\mathbf{j}) \leq P(\mathbf{i} \vee \mathbf{j})Q(\mathbf{i} \wedge \mathbf{j})$. If P and Q are univariate, then this definition is equivalent to the MLR ordering $P \geq_r Q$ defined above.
- (ii) A multivariate distribution P is said to be multivariate TP2 (MTP2) if $P \geq_{\text{TP2}} P$ holds, i.e., $P(\mathbf{i})P(\mathbf{j}) \leq P(\mathbf{i} \vee \mathbf{j})P(\mathbf{i} \wedge \mathbf{j})$.
- (iii) If $\mathbf{i}, \mathbf{j} \in \{1, \dots, X\}$ are scalar indices, Statement (ii) is equivalent to saying that an $X \times X$ matrix P is TP2 if all second order minors are non-negative. Equivalently, $P_{i+1, \cdot} \geq_r P_{i, \cdot}$, where $P_{i, \cdot}$ denotes the i th row of matrix P .

To prove the existence of a threshold switching curve, we will show that $Q(\pi, u)$ in (25) is a submodular function on chains $[\Pi(X), \geq_{L_X}]$ and $[\Pi(X), \geq_{L_1}]$.

⁸A chain is totally ordered subset of a partially ordered set.

Definition 5 (Submodular Function [48]): Suppose $i = 1$ or X . Then $f : \mathcal{L}(e_i, \bar{\pi}) \times \{1, 2\} \rightarrow \mathbb{R}$ is submodular (antitone differences) if $f(\pi, u) - f(\pi, \bar{u}) \leq f(\tilde{\pi}, u) - f(\tilde{\pi}, \bar{u})$, for $\bar{u} \leq u, \pi \geq_{L_i} \tilde{\pi}$.

The following result says that for a submodular function $Q(\pi, u)$, $\mu^*(\pi) = \operatorname{argmin}_u Q(\pi, u)$ is increasing in its argument π . This implies $\mu^*(\pi)$ is MLR increasing on the line segments $\mathcal{L}(e_1, \bar{\pi})$ and $\mathcal{L}(e_X, \bar{\pi})$, which in turn will be used to prove the existence of as threshold decision curve.

Theorem 12 ([48]): Suppose $i = 1$ or X . If $f : \mathcal{L}(e_i, \bar{\pi}) \times \{1, 2\} \rightarrow \mathbb{R}$ is submodular, then there exists a $\mu^*(\pi) = \operatorname{argmin}_{u \in \{1, 2\}} f(\pi, u)$, that is increasing on $[\mathcal{L}(e_i, \bar{\pi}), \geq_{L_i}]$, i.e., $\tilde{\pi} \geq_{L_i} \pi \implies \mu^*(\pi) \leq \mu^*(\tilde{\pi})$.

B) Meta-Theorem: Lattice Programming on Simplex: We start with the following general assumptions and meta-theorem. The meta-theorem is a major step in all our structural results.

- (A) (A1) $C(\pi, 1)$ and $C(\pi, 2)$ are first order stochastic decreasing in π .
- (B) (A2) B_{xy} is TP2.
- (C) (A3) P is TP2.
- (D) $C(\pi, 2) - C(\pi, 1)$ is $[\mathcal{L}(e_X, \bar{\pi}), \geq_{L_X}]$ decreasing and $[\mathcal{L}(e_1, \bar{\pi}), \geq_{L_1}]$ decreasing.

Theorem 13: Consider the generic dynamic programming (25) and the above assumptions. Then the following properties hold.

- 1) $\pi_1 \geq_r \pi_2$ implies $T(\pi_1, y) \geq_r T(\pi_2, y)$ if (A3) holds. Under (A2) and (A3), $\sigma(\pi_1, y) \geq_s \sigma(\pi_2, y)$.
- 2) For $y, \bar{y} \in \mathbb{Y}$, $y > \bar{y}$ implies $T(\pi_1, y) \geq_r T(\pi_1, \bar{y})$ iff (A2) holds.
- 3) Assumptions (A1-Ex1) in Section III-A, (AS-Ex1)(i) and (ii) in Section III-A, (A1-Ex2) in Section IV, (A1-Ex3) in Section V, and (A1-Ex5) in Section VI are sufficient conditions for (A1).
- 4) Under (A1), (A2), (A3), $Q(\pi, u)$ is MLR decreasing wrt \geq_{L_1} and \geq_{L_X} .
- 5) Assumptions (S-Ex1) in Section III-A, (AS-Ex1)(i) and (iii) in Section III-A, (S-Ex2) in Section IV, (S-Ex3) in Section V and (S-Ex5) in Section VI-C are sufficient conditions for (S).
- 6) Under (A1), (A2), (A3), (S), $Q(\pi, u)$ is submodular on $[\Pi(X), \geq_{L_X}]$ and $[\Pi(X), \geq_{L_1}]$. Thus, by Theorem 12, the optimal policy $\mu^*(\pi)$ is MLR increasing on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$.

Part 1 and Part 2 use elementary properties of positive matrices and are proved in [41, Lemma 4.1 and 4.2].

Proof of Part 3 for $C(\pi, 1)$: To give sufficient conditions for $C(\pi, 1)$ to \geq_r decrease wrt $\pi \in \Pi(X)$, we start with the following convenient parametrization of the family of belief states that first-order stochastic dominate another belief state, see [32] for proof.

Lemma 3: (i) For any $\pi, \tilde{\pi} \in \Pi(X)$, all belief states $\tilde{\pi} \leq_s \pi$ are of the form

$$\begin{aligned} \tilde{\pi} &= \pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}} \triangleq \pi + \epsilon_1 (e_1 - e_2) + \epsilon_2 (e_2 - e_3) + \dots \\ &\quad + \epsilon_{X-1} (e_{X-1} - e_X). \end{aligned} \quad (82)$$

where the variables ϵ_j satisfy $0 \leq \epsilon_j \leq \min\{1 - \pi(j), \pi(j + 1)\}$, $j = 1, \dots, X - 1$. Moreover, the ϵ -parametrized belief state $\pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}}$ is stochastically decreasing in the elements ϵ_j . That is, for any $j = 1, \dots, X - 1$, $\epsilon_j \leq \bar{\epsilon}_j$ implies that $\pi^{\epsilon_1, \dots, \epsilon_j, \dots, \epsilon_{S-1}} \geq_s \pi^{\bar{\epsilon}_1, \dots, \bar{\epsilon}_j, \dots, \epsilon_{S-1}}$. ■

Remark: The above constraints on ϵ_j ensure that $\pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}}$ is a valid belief state.

In light of the above lemma, it suffices to prove that $C(\pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}}, 1)$ is increasing in ϵ_i , $i = 1, 2, \dots, X$. We introduce the following lemma. The proof follows straightforwardly using $\partial F(\pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}})/\partial \epsilon_i \geq 0$ and is omitted.

Lemma 4: Suppose $F(\pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}}) = \phi' \pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}} - \alpha(h' \pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}})^2$, where $\phi, h \in \mathbb{R}^X$. Then if $h_i > 0$, a sufficient condition for $F(\pi^{\epsilon_1, \epsilon_2, \dots, \epsilon_{S-1}})$ to be \geq_s increasing wrt ϵ is

$$\phi_i - \phi_{i+1} \geq 2\alpha h' \pi (h_i - h_{i+1}) \quad \forall \pi \in \Pi(X). \quad (83)$$

If $h_i \geq 0$ is either monotone increasing or decreasing in i , then a sufficient condition for (83) is

$$\phi_i - \phi_{i+1} \geq 2\alpha h_1 (h_i - h_{i+1}). \quad (84)$$

■

- Theorem 1: Set $\phi = 2\alpha e_1$, $h = e_1$ in (84). This yields $2\alpha \geq 0$ and $2\alpha \geq 2\alpha$ which always hold. So $C(\pi, 1)$ is \geq_r decreasing in π for any non-negative α .
- Theorem 2: Set $\phi = \alpha e_1 - \alpha \mathbf{f}$ in (84). This yields $f_2 \geq 1$ and $f_{i+1} \geq f_i$. Clearly (AS-Ex1)(i) and (ii) are sufficient conditions for this. In particular, P TP2 and (AS-Ex1)(ii) implies $f_{i+1} \geq f_i$.
- Theorem 5: Note that the variance constraint $\alpha(e'_3 \pi - (e'_3 \pi)^2) = \alpha((\pi_1 + \pi_2) - (\pi_1 + \pi_2)^2)$. Accordingly, set $\phi = 2\alpha(e_1 + e_2)$, $h = e_1 + e_2$ in (84). This yields $2\alpha \geq 0$ and $2\alpha \geq 2\alpha$ which always hold for $\alpha \geq 0$.
- In Theorem 6, $C(\pi, 1) = 0$, see (50), so there is nothing to prove.
- In Theorem 9, to show that $C(\pi, 1)$ is \geq_r decreasing in π , it suffices to show that for each $a \in \mathbb{A}$, that $c'_a \pi$ is \geq_r decreasing in π . So in (84) choose $\phi = c_a$, $h = 0$. This yields (A1-Ex5).

Proof of Part 3 for $C(\pi, 2)$: Since $C(\pi, 2)$ is linear in π , to show that $C(\pi, 2) \geq_r$ decreases wrt π , from Result 1(ii) (Appendix A, it suffices that

$$C(\epsilon_i, 2) \geq C(\epsilon_{i+1}, 2), \quad i = 1, \dots, X - 1. \quad (85)$$

Theorem 1: This yields (A1-Ex1).

Theorem 2: (85) is equivalent to $f_2 \geq \rho \mathbf{f}' P' e_2 - \frac{d}{\alpha + \beta}$ and (AS-Ex1)(ii). Note (AS-Ex1)(i) is sufficient for $f_2 \geq \rho \mathbf{f}' P' e_2 - \frac{d}{\alpha + \beta}$.

Theorem 5: (85) is equivalent to $-\rho(\alpha + \beta) + d_1 + (\alpha + \beta) \geq -\rho(\alpha + \beta) + d_2 + (\alpha + \beta) \geq -\rho(\alpha + \beta) P_{32}$ which always holds for $d_1 \geq d_2$.

Proof of Part 4: The proof is by mathematical induction on the value iteration recursion (27). Clearly $V_0(\pi) = -\beta(1 - e'_1 \pi)$ in (27) is a MLR decreasing function of π . Consider (27) at any stage k . Assume that $V_k(\pi)$ is MLR decreasing in π .

From Part 1 and 2, it follows that under (A2), (A3), the term $\sum_{y \in \mathcal{Y}} V_k(T(\pi, y)) \sigma(\pi, y)$ is MLR decreasing in π . From Part 3, under (A1), $C(\pi, u)$ is MLR decreasing. Since the sum of decreasing functions is decreasing, the result follows.

Proof of Part 5 : Here We Show That General Assumption (s) At the Beginning of Appendix B Takes On the Forms (s-ex1) To (s-ex5) For the Various Examples. We Start With the Following Characterization of Belief States On Lines $\mathcal{L}(e_x, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$ and Submodularity On These Lines.

Lemma 5:

- (i) $\pi_1 \geq_{L_X} \pi_2$ is equivalent to $\pi_i = (1 - \epsilon_i)\bar{\pi} + \epsilon_i e_X$ and $\epsilon_1 \geq \epsilon_2$ for $\bar{\pi} \in \mathcal{H}_X$ where \mathcal{H}_X is defined in (79). So submodularity on $\mathcal{L}(e_X, \bar{\pi})$ is equivalent to showing

$$\begin{aligned} \pi^\epsilon &= (1 - \epsilon)\bar{\pi} + \epsilon e_X \\ &\Rightarrow C(\pi^\epsilon, 2) - C(\pi^\epsilon, 1) \text{ decreasing wrt } \epsilon. \quad (86) \end{aligned}$$

- (ii) $\pi_1 \geq_{L_1} \pi_2$ is equivalent to $\pi_i = (1 - \epsilon_i)\bar{\pi} + \epsilon_i e_X$ and $\epsilon_1 \leq \epsilon_2$ for $\bar{\pi} \in \mathcal{H}_1$ where \mathcal{H}_1 is defined in (79). So submodularity on $\mathcal{L}(e_1, \bar{\pi})$ is equivalent to showing

$$\begin{aligned} \pi^\epsilon &= (1 - \epsilon)\bar{\pi} + \epsilon e_1 \\ &\Rightarrow C(\pi^\epsilon, 2) - C(\pi^\epsilon, 1) \text{ increasing wrt } \epsilon. \quad (87) \end{aligned}$$

■

The proof of Lemma 5 follows from Lemma 2 and is omitted.

Suppose $C(\pi, 2) - C(\pi, 1)$ is of the form $\phi' \pi + \alpha(h' \pi)^2$. Then from (86), (87), sufficient conditions for submodularity on $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$ are for $\bar{\pi} \in \mathcal{H}_X$ and \mathcal{H}_1 , respectively

$$\begin{aligned} \phi_X - \phi' \bar{\pi} + 2\alpha h' \pi^\epsilon (h_X - h' \bar{\pi}) &\leq 0 \\ \phi_1 - \phi' \bar{\pi} + 2\alpha h' \pi^\epsilon (h_1 - h' \bar{\pi}) &\geq 0. \quad (88) \end{aligned}$$

In particular if $h_i \geq 0$ and monotone increasing or decreasing in i , then (88) is equivalent to

$$\begin{aligned} \phi_X - \phi' \bar{\pi} + 2\alpha h_X (h_X - h' \bar{\pi}) &\leq 0 \\ \phi_1 - \phi' \bar{\pi} + 2\alpha h_X (h_1 - h' \bar{\pi}) &\geq 0 \quad (89) \end{aligned}$$

where $\bar{\pi} \in \mathcal{H}_X$ and $\bar{\pi} \in \mathcal{H}_1$, respectively.

- Theorem 1: Set $\phi_i = (d - \rho(\alpha + \beta))P' e_i + (\beta - \alpha)e_1$, $h = e_1$ in (89). The first inequality is equivalent to: (i) $(d - \rho(\alpha + \beta))(P_{X1} - P_{i1}) \leq 0$ for $i \geq 2$ and (ii) $(d - \rho(\alpha + \beta))(1 - P_{X1}) \geq \alpha - \beta$. Note that (i) holds if $d \geq \rho(\alpha + \beta)$. The second inequality in (89) is equivalent to $(d - \rho(\alpha + \beta))(1 - P_{i1}) \geq \alpha - \beta$. Since P is TP2, from footnote 4 in Section III-B it follows that (S-Ex1) is sufficient for these inequalities to hold.
- Theorem 2: Set $\phi_1 = d - \alpha$, $\phi_i = -\beta f_i + \rho(\alpha + \beta) \mathbf{f}' P' e_i$, $i = 2, \dots, X$ in (89). The first inequality yields (i) $\rho(\alpha + \beta) \mathbf{f}' P' e_X \leq d - \alpha + \beta f_X$ and (ii) $\rho(\alpha + \beta) \mathbf{f}' P' (e_X - e_i) \leq \beta(f_X - f_i)$, $i = 2, \dots, X - 1$. The second inequality in (89) yields $q_i \geq \frac{\rho}{\beta}(\alpha + \beta) \mathbf{f}' P' e_i + \frac{\alpha - d}{\beta}$. These inequalities imply (AS-Ex1)(i) and (iii).
- Theorem 5: Recall that the variance constraint $\alpha(e'_3 \pi - (e'_3 \pi)^2) = \alpha((\pi_1 + \pi_2) - (\pi_1 + \pi_2)^2)$. Set $\phi_1 = d_1 + \beta - \alpha - \rho(\alpha + \beta)$, $\phi_2 = d_2 + \beta - \alpha - \rho(\alpha + \beta)$, $f_3 = -\rho(\alpha + \beta) P_{32}$ in (89). The first inequality is equivalent to $\rho(\alpha + \beta)(1 -$

$P_{32}) \leq \beta + d_i - \alpha$, $i = 1, 2$. The second inequality is equivalent to $d_1 \geq 0$ and $d_1 + \beta - \alpha \geq \rho(\alpha + \beta)(1 - P_{32})$. A sufficient condition for these is (S-Ex2).

- In Theorem 6, since $C(\pi, 1) = 0$, $C(\pi, 2)$ decreasing on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$ implies submodularity. This is implied by (A1-Ex3).
- In Theorem 9, $f = \sum_y B_y c_y - c_a / (1 - \rho)$, $h = 0$ in (89) yields (S-Ex5).

Proof of Part 6: From Definition 5, to show that $Q(\pi, u)$ is submodular, requires showing that $Q(\pi, 1) - Q(\pi, 2)$ is \geq_{L_i} on lines $\mathcal{L}(e_X, \bar{\pi})$ for $i = 1$ and X . Part 4 shows by induction that for each k , $V_k(\pi)$ is \geq_r decreasing on $\Pi(X)$ if (A1), (A2), (A3) hold. This implies that $V_k(\pi)$ is \geq_{L_i} decreasing on lines $\mathcal{L}(e_X, \bar{\pi})$ and $\mathcal{L}(e_1, \bar{\pi})$. So to prove $Q_k(\pi, u)$ in (5) is submodular, we only need to show that $C(\pi, 1) - C(\pi, 2)$ is \geq_{L_i} decreasing on $\mathcal{L}(e_i, \bar{\pi})$, $i = 1, X$. But this is implied by (S) as shown in Part 5 above. Since submodularity is closed under pointwise limits [48, Lemma 2.6.1 and Corollary 2.6.1], it follows that $Q(\pi, u)$ is submodular on \geq_{L_i} , $i = 1, X$. Having established $Q(\pi, u)$ is submodular on \geq_{L_i} , $i = 1, X$, Theorem 12 in Appendix A implies that the optimal policy $\mu^*(\pi)$ is \geq_{L_1} and \geq_{L_X} increasing on lines.

C) *Proof of Theorem 1:* With the above key theorem, we can now prove Theorem 1. The statement of Theorem 1 that $\mu^*(\pi)$ is \geq_{L_1} and \geq_{L_X} increasing is proved above.

Statement (i): (a) Characterization of Switching Curve Γ . For each $\bar{\pi} \in \mathcal{H}_X$ (80), construct the line segment $\mathcal{L}(e_X, \bar{\pi})$ connecting \mathcal{H}_X to e_X as in (80). By Lemma 2 in Appendix A, on the line segment connecting $(1 - \epsilon)\bar{\pi} + \epsilon e_X$, all belief states are MLR orderable. Since $\mu^*(\pi)$ is monotone increasing for $\pi \in \mathcal{L}(e_X, \bar{\pi})$, moving along this line segment towards e_X , pick the largest ϵ for which $\mu^*(\pi) = 1$. (Since $\mu^*(e_X) = 1$, such an ϵ always exists). The belief state corresponding to this ϵ is the threshold belief state. Denote it by $\Gamma(\bar{\pi}) = \pi^{\epsilon^*, \bar{\pi}} \in \mathcal{L}(e_X, \bar{\pi})$ where $\epsilon^* = \sup\{\epsilon \in [0, 1] : \mu^*(\pi^{\epsilon, \bar{\pi}}) = 1\}$.

The above construction implies that on $\mathcal{L}(e_X, \bar{\pi})$, there is a unique threshold point $\Gamma(\bar{\pi})$. Note that the entire simplex can be covered by considering all pairs of lines $\mathcal{L}(e_X, \bar{\pi})$, for $\bar{\pi} \in \mathcal{H}_X$, i.e., $\Pi(X) = \cup_{\bar{\pi} \in \mathcal{H}} \mathcal{L}(e_X, \bar{\pi})$. Combining all points $\Gamma(\bar{\pi})$ for all pairs of lines $\mathcal{L}(e_X, \bar{\pi})$, $\bar{\pi} \in \mathcal{H}_X$, yields a unique threshold curve in $\Pi(X)$ denoted $\Gamma = \cup_{\bar{\pi} \in \mathcal{H}} \Gamma(\bar{\pi})$.

Statement (i): (b) Connectedness of regions \mathcal{R}_1 and \mathcal{R}_2 .

Connectedness of \mathcal{R}_1 : Since $e_1 \in \mathcal{R}_1$, call \mathcal{R}_{1a} the subset of \mathcal{R}_1 that contains e_1 . Suppose \mathcal{R}_{1b} was a subset of \mathcal{R}_1 that was disconnected from \mathcal{R}_{1a} . Recall that every point in $\Pi(X)$ lies on a line segment $\mathcal{L}(e_1, \bar{\pi})$ for some $\bar{\pi}$. Then such a line segment starting from $e_1 \in \mathcal{R}_{1a}$ would leave the region \mathcal{R}_{1a} , pass through a region where action 2 was optimal, and then intersect the region \mathcal{R}_{1b} where action 1 is optimal. But this violates the requirement that $\mu(\pi)$ is increasing on $\mathcal{L}(e_1, \bar{\pi})$. Hence, \mathcal{R}_{1a} and \mathcal{R}_{1b} have to be connected. (Note in the special case $\alpha = 0$, then since \mathcal{R}_1 is convex (by Theorem 3), and so is obviously connected).

Connectedness of \mathcal{R}_2 : Assume $e_X \in \mathcal{R}_2$, otherwise $\mathcal{R}_2 = \emptyset$ and there is nothing to prove. Call the region \mathcal{R}_2 that contains e_X as \mathcal{R}_{2a} . Suppose $\mathcal{R}_{2b} \subset \mathcal{R}_2$ is disconnected from \mathcal{R}_{2a} . Since every point in $\Pi(X)$ can be joined by the line segment $\mathcal{L}(e_X, \bar{\pi})$ to e_X . Then such a line segment starting from $e_X \in \mathcal{R}_{2a}$ would

leave the region \mathcal{R}_{2a} , pass through a region where action 1 was optimal, and then intersect the region \mathcal{R}_{2b} (where action 2 is optimal). But this violates the requirement that $\mu(\pi)$ is increasing on $\mathcal{L}(e_X, \bar{\pi})$. Hence, \mathcal{R}_{2a} and \mathcal{R}_{2b} have to be connected.

Statement (ii): Suppose $e_i \in \mathcal{R}_1$. Then considering lines $\mathcal{L}(e_i, \bar{\pi})$ and ordering \geq_{L_i} , it follows that $e_{i-1} \in \mathcal{R}_1$. Similarly if $e_i \in \mathcal{R}_2$, then considering lines $\mathcal{L}(e_{i+1}, \bar{\pi})$ and ordering $\geq_{L_{i+1}}$, it follows that $e_{i+1} \in \mathcal{R}_2$.

Statement (iii) follows trivially since for $X = 2$, $\Pi(X)$ is a 1-D simplex.

D) *Proof of Theorem 3:* We first prove in the following lemma that $V(\pi)$ is concave in π .

Lemma 6: $V(\pi)$ in (25) is concave in $\pi \in \Pi(X)$.

Proof of Lemma 6: Our proof constructs an outer approximation $V_k^{(n)}(\pi)$ (defined below) to $V_k(\pi)$ and comprises of two steps: Step 1: $V_k^{(n)}(\pi)$ is concave; Step 2: $V_k^{(n)}(\pi) \rightarrow V_k(\pi)$ uniformly as $n \rightarrow \infty$. This establishes that $V_k(\pi)$ is concave, and therefore, $V(\pi)$ is concave.

Consider n arbitrary but distinct belief states $\pi^1, \dots, \pi^n \in \Pi(X)$. Let $\gamma^i(\pi)$ denote the gradient vector of $Q(\pi, 1)$ in (27) at $\pi = \pi^i$, $i = 1, \dots, n$. That is, $Q(\pi^i, 1) = \pi^i \gamma^i(\pi^i)$. Now construct a piecewise linear function out of these gradient vectors as $Q^{(n)}(\pi, 1) = \min_i \gamma^i \pi$. It is easily seen from (24) that $Q(\pi, 1)$ is concave. Therefore, $Q^{(n)}(\pi, 1)$ is piecewise linear and concave in π since a piecewise linear function composed of tangents to a concave function is concave.

Construct the following auxiliary value function $V^{(n)}(\pi)$ via value iteration similar to (27)

$$\begin{aligned} V_{k+1}^{(n)}(\pi) &= \min\{Q^{(n)}(\pi, 1), Q_{k+1}^{(n)}(\pi, 2)\}, \mu_{k+1}^*(\pi) \\ &= \operatorname{argmin}\{Q^{(n)}(\pi, 1), Q_{k+1}^{(n)}(\pi, 2)\} \end{aligned}$$

where

$$Q_{k+1}^{(n)}(\pi, 2) = C(\pi, 2) + \rho \sum_{y \in \mathcal{Y}} V_k^{(n)}(T(\pi, y)) \sigma(\pi, y)$$

$$Q^{(n)}(\pi, 1) = \min_i \gamma^i \pi \quad \pi \in \Pi(X). \quad (90)$$

Step 1: Proof of concavity of $V_k^{(n)}(\pi)$: We prove this by induction on the value iteration algorithm (90) for $k = 1, 2, \dots$ and fixed n . Start with arbitrary concave $V_0^{(n)}(\pi)$. As mentioned below (27), the VI algorithm converges for any choice of initialization.

Since both $C(\pi, 2)$ (see (24)) and $Q^{(n)}(\pi, 1)$ are piecewise linear in π , it is easily seen from (90) that at each iteration k , $V_k^{(n)}(\pi)$ is positively homogeneous, i.e., $V_k^{(n)}(\lambda\pi) = \lambda V_k^{(n)}(\pi)$ for any $\lambda \geq 0$. As a result (90) yields, $Q_{k+1}^{(n)}(\pi, 2) = C(\pi, 2) + \rho \sum_{y \in \mathcal{Y}} V_k^{(n)}(B_y \pi)$. Now use mathematical induction. Assume $V_k^{(n)}(\pi)$ is concave. Since $B_y \pi$ is concave, and the composition of concave functions is concave, $V_k^{(n)}(B_y \pi)$ is concave. Since $C(\pi, 2)$ is piecewise linear and concave, and the sum of concave functions is concave, it follows that $Q_{k+1}^{(n)}(\pi, 2)$ is concave. Finally since minimization preserves concavity, it follows that $V_{k+1}^{(n)}(\pi) = \min\{Q^{(n)}(\pi, 1), Q_{k+1}^{(n)}(\pi, 2)\}$ is concave. This completes the inductive proof.

Step 2: Concavity of $V_k(\pi)$: Next, we show that $V_k^{(n)}(\pi) \rightarrow V_k(\pi)$ uniformly in $\Pi(X)$ implying that $V_k(\pi)$ is concave. Since $Q(\pi, 1)$ is concave, it follows that $Q^{(n)}(\pi, 1) > Q(\pi, 1)$ and also $Q^{(n)}(\pi, 1)$ is a monotone sequence of decreasing functions in n . (Intuitively, the piecewise linear function composed of tangents always upper bound a concave function and they become tighter as more piecewise linear segments are considered). From (25) and (90) this implies that $V_k^{(n)}(\pi) \geq V_k(\pi)$ for all k and also $V_k^{(n+1)}(\pi) < V_k^{(n)}(\pi)$. Finally, since (i) $V_k^{(n)}(\pi)$ converges to $V_k(\pi)$ pointwise, (ii) $\Pi(X)$ is compact, (iii) $V_k^{(n)}(\pi)$ is continuous and $V_k^{(n)}(\pi)$ is monotone decreasing sequence in n , it follows from [42, Theorem 7.13] that $V_k^{(n)}(\pi) \rightarrow V_k(\pi)$ uniformly on $\Pi(X)$. Therefore, $V_k(\pi)$ is concave for $k = 1, 2, \dots$. Finally as discussed below (27), $V_k(\pi)$ converges uniformly to $V(\pi)$; so $V(\pi)$ is concave. Thus, Lemma 6 is proved. ■

The rest of the proof of Theorem 3 follows from arguments in [28]. We repeat this for completeness here. Our goal is to show that \mathcal{R}_1 is convex. Pick any two belief states $\pi_1, \pi_2 \in \mathcal{R}_1$. To demonstrate convexity of \mathcal{R}_1 , we need to show for any $\lambda \in [0, 1]$, $\lambda\pi_1 + (1 - \lambda)\pi_2 \in \mathcal{R}_1$. Since $V(\pi)$ is concave and $\alpha = 0$

$$\begin{aligned} & V(\lambda\pi_1 + (1 - \lambda)\pi_2) \\ & \geq \lambda V(\pi_1) + (1 - \lambda)V(\pi_2) \\ & = \lambda Q(\pi_1, 1) + (1 - \lambda)Q(\pi_2, 1) \\ & \quad (\text{since } \pi_1, \pi_2 \in \mathcal{R}_1) \\ & = Q(\lambda\pi_1 + (1 - \lambda)\pi_2, 1) \\ & \quad (\text{since } Q_1(\pi, 1) \text{ is linear in } \pi) \\ & \geq V(\lambda\pi_1 + (1 - \lambda)\pi_2) \\ & \quad (\text{since } V(\pi) \text{ is the optimal value function}). \end{aligned} \quad (91)$$

Thus, all the inequalities above are equalities, and $\lambda\pi_1 + (1 - \lambda)\pi_2 \in \mathcal{R}_1$.

E) Proof of Theorem 4: Given any $\pi_1, \pi_2 \in \mathcal{L}(e_X, \bar{\pi})$ with $\pi_2 \geq_{L_X} \pi_1$, we need to prove: $\mu_\theta(\pi_1) \leq \mu_\theta(\pi_2)$ iff $\theta(X - 2) \geq 1$, $\theta(i) \leq \theta(X - 2)$ for $i < X - 2$. But from the structure of (32), obviously $\mu_\theta(\pi_1) \leq \mu_\theta(\pi_2)$ is equivalent to $[0 \ 1 \ \theta']' \begin{bmatrix} \pi_1 \\ -1 \end{bmatrix} \leq [0 \ 1 \ \theta']' \begin{bmatrix} \pi_2 \\ -1 \end{bmatrix}$, or equivalently, $[0 \ 1 \ \theta(1) \ \dots \ \theta(X - 2)](\pi_1 - \pi_2) \leq 0$.

Now from Lemma 2(iii), $\pi_2 \geq_{L_X} \pi_1$ implies that $\pi_1 = \epsilon_1 e_X + (1 - \epsilon_1)\bar{\pi}$, $\pi_2 = \epsilon_2 e_X + (1 - \epsilon_2)\bar{\pi}$ and $\epsilon_1 \leq \epsilon_2$. Substituting these into the above expression, we need to prove

$$(\epsilon_1 - \epsilon_2)(\theta(X - 2) - [0 \ 1 \ \theta(1) \ \dots \ \theta(X - 2)]' \bar{\pi}) \leq 0 \quad \forall \bar{\pi} \in \mathcal{H}_X$$

iff $\theta(X - 2) \geq 1$, $\theta(i) \leq \theta(X - 2)$, $i < X - 2$. This is obviously true.

A similar proof shows that on lines $\mathcal{L}(e_1, \bar{\pi})$ the linear threshold policy satisfies $\mu_\theta(\pi_1) \leq \mu_\theta(\pi_2)$ iff $\theta(i) \geq 0$ for $i < X - 2$.

F) Proof of Theorem 6: The only difference compared to the meta-theorem is the update of the belief state (49) which now includes the term $\text{diag}(R_u)$. The elements of R_u are

non-negative and functionally independent of the observation y . Therefore, the three main requirements that $T(\pi, y)$ is MLR increasing in π , $T(\pi, y)$ is MLR increasing in Y , and $\sigma(\pi, \cdot) \geq_r$ increasing in π continue to hold. Then the rest of the proof is identical to Theorem 1.

G) Proof of Theorem 8: The proof is more complex than that of Theorem 3 since now $V(\pi)$ is not necessarily concave over $\Pi(X)$, since $T(\cdot)$ and $\sigma(\cdot)$ are functions of R_a^π (56) which itself is an explicit (and in general nonconcave) function of π .

Define the matrix $R^\pi = (R^\pi(i, a), i = \{1, 2\}, a \in \{1, 2\})$, where $R^\pi(i, a) = P(a|x = e_i, \pi)$. It can be verified from (56) that there are only 3 possible values for R^π , namely

$$\begin{aligned} R^\pi &= \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}, \pi \in \mathcal{P}_1 \\ R^\pi &= B, \pi \in \mathcal{P}_2 \cup \mathcal{P}_3 \\ R^\pi &= \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}, \pi \in \mathcal{P}_4. \end{aligned} \quad (92)$$

Thus, based on the dynamic programming equation (61), the value iteration algorithm reads

$$\begin{aligned} V_{n+1}(\pi) &= \min\{C(\pi, 2) + \rho V_n(\pi)I(\pi \in \mathcal{P}_1) \\ & \quad + \rho \sum_{a \in \mathcal{A}} V_n(T(\pi, a))\sigma(\pi, a)[I(\pi \in \mathcal{P}_2) \\ & \quad + I(\pi \in \mathcal{P}_3)] + \rho V_n(\pi)I(\pi \in \mathcal{P}_4), 0\}. \end{aligned} \quad (93)$$

Assuming $V_n(\pi)$ is MLR decreasing on $\mathcal{P}_1 \cup \mathcal{P}_4$ straightforwardly implies $V_{n+1}(\pi)$ is MLR decreasing on $\mathcal{P}_1 \cup \mathcal{P}_4$ since $C(\pi, 2)$ is MLR decreasing. This proves claim (i).

We now prove inductively that $V_n(\pi)$ is piecewise linear concave on each interval \mathcal{P}_l , $l = 1, \dots, 4$. The proof of concavity on \mathcal{P}_1 and \mathcal{P}_4 follows straightforwardly since $C(\pi, 2)$ is piecewise linear and concave. The proof for intervals \mathcal{P}_2 and \mathcal{P}_3 is more delicate.

We need the following property of the social learning Bayesian filter. We use the following slight abuse in notation. Define the 2-D vector $\eta_i = (1 - \eta_i, \eta_i)'$.

Lemma 7: Consider the social learning Bayesian filter (55). Then $T(\eta_1, 1) = \eta_2$, $T(\eta_3, 2) = \eta_2$. Furthermore if B is symmetric TP2, then $T(\eta_2, 2) = \eta_1$, $T(\eta_2, 1) = \eta_3$ and $\eta_3 \leq \eta_2 \leq \eta_1$. So:

- (i) $\pi \in \mathcal{P}_2$ implies $T(\pi, 2) \in \mathcal{P}_1$ and $T(\pi, 1) \in \mathcal{P}_3$.
- (ii) $\pi \in \mathcal{P}_3$ implies $T(\pi, 2) \in \mathcal{P}_2$ and $T(\pi, 1) \in \mathcal{P}_4$.

■

The proof of Lemma 7 is as follows. Recall from (92) that on intervals \mathcal{P}_2 and \mathcal{P}_3 , $R^\pi = B$. Then it is straightforwardly verified from (55) that $T(\eta_1, 1) = T(\eta_3, 2) = \eta_2$. Next, using (55) it follows that $B_{12}B_{11} = B_{22}B_{21}$ is a sufficient condition for $T(\eta_2, 2) = \eta_1$ and $T(\eta_2, 1) = \eta_3$. Also, applying Theorem 13 (2), B TP2 implies $\eta_3 \leq \eta_2 \leq \eta_1$. So B symmetric TP2 is sufficient for the claims of the lemma to hold. Statements (i) and (ii) then follow straightforwardly. In particular, from Theorem 13 (1), $\eta_1 \geq_r \pi \geq_r \eta_2$ implies $T(\eta_1, 1) = \eta_2 \geq_r T(\pi, 1) \geq_r T(\eta_2, 1) = \eta_3$, which implies Statement (i) of the Lemma. Statement (ii) follows similarly.

Returning to the proof of Theorem 8. Assume now that $V_n(\pi)$ is piecewise linear and concave on each interval \mathcal{P}_l , $l = 1, \dots, 4$. That is, for 2-D vectors γ_{m_l} in the set Γ_l

$$V_n(\pi) = \sum_l \min_{m_l \in \Gamma_l} \gamma'_{m_l} \pi I(\pi \in \mathcal{P}_l).$$

Consider $\pi \in \mathcal{P}_2$. From (92), since $R_a^\pi = B_a$, $a = 1, 2$, Lemma 7 (i) together with the value iteration algorithm (93) yields

$$V_{n+1}(\pi) = \min \left\{ C(\pi, 2) + \rho \left[\min_{m_3 \in \Gamma_3} \gamma'_{m_3} B_1 \pi + \min_{m_1 \in \Gamma_1} \gamma'_{m_1} B_2 \pi \right], 0 \right\}.$$

Since each of the terms in the above equation are piecewise linear and concave, it follows that $V_{n+1}(\pi)$ is piecewise linear and concave on \mathcal{P}_2 . A similar proof holds for \mathcal{P}_3 and this involves using Lemma 7(ii). As a result the stopping set on each interval \mathcal{P}_l is a convex region, i.e., an interval. This proves claim (ii).

H) *Proof of Theorem 9:* Part (i) follows directly from the proof of Theorems 13 and 1. For Part (ii), define the convex polytopes $\mathcal{P}_a = \{\pi : c'_a \pi < c'_a \bar{\pi}, \bar{a} \neq a\}$. Then on each convex polytope \mathcal{P}_a , since $C(\pi, 1) = c'_a \pi$ (recall $\alpha = 0$), we can apply the argument of (91) which yields that $\mathcal{R}_1 \cap \mathcal{P}_a$ is a convex region. Thus, \mathcal{R}_1 is the union of A convex regions and is in general nonconvex. However, it is still a connected set by part (i) of the theorem.

I) *Proof of Theorem 10:* A similar proof to Theorem 3 then establishes $V(\pi)$ is concave on $\Pi(X)$. We then use the Blackwell dominance condition (71). In particular

$$T(\pi, y^{(1)}) = \sum_{y^{(2)} \in \mathcal{Y}^{(2)}} T(\pi, y^{(2)}) \frac{\sigma(\pi, y^{(2)})}{\sigma(\pi, y^{(1)})} P(y^{(1)} | y^{(2)}) \quad \text{and}$$

$$\sigma(\pi, y^{(1)}) = \sum_{y^{(2)} \in \mathcal{Y}^{(2)}} \sigma(\pi, y^{(2)}) P(y^{(1)} | y^{(2)}).$$

Therefore, $\frac{\sigma(\pi, y^{(2)})}{\sigma(\pi, y^{(1)})} P(y^{(1)} | y^{(2)})$ is a probability measure wrt $y^{(2)}$. Since $V(\cdot)$ is concave, using Jensen's inequality it follows that

$$\begin{aligned} & V\left(T(\pi, y^{(1)})\right) \\ &= V\left(\sum_{y^{(2)} \in \mathcal{Y}^{(2)}} T(\pi, y^{(2)}) \frac{\sigma(\pi, y^{(2)})}{\sigma(\pi, y^{(1)})} P(y^{(1)} | y^{(2)})\right) \\ &\geq \sum_{y^{(2)} \in \mathcal{Y}^{(2)}} V(T(\pi, y^{(2)})) \frac{\sigma(\pi, y^{(2)})}{\sigma(\pi, y^{(1)})} P(y^{(1)} | y^{(2)}) \\ &\quad \text{implying} \\ &\sum_{y^{(1)}} V\left(T(\pi, y^{(1)})\right) \sigma(\pi, y^{(1)}) \\ &\geq \sum_{y^{(2)}} V(T(\pi, y^{(2)})) \sigma(\pi, y^{(2)}). \end{aligned}$$

Therefore, for $\pi \in \Pi^s$

$$\begin{aligned} C(\pi, 2) + \sum_{y^{(2)}} V\left(T(\pi, y^{(2)})\right) \sigma(\pi, y^{(2)}) \\ \leq C(\pi, 1) + \sum_{y^{(1)}} V\left(T(\pi, y^{(1)})\right) \sigma(\pi, y^{(1)}). \end{aligned}$$

So for $\pi \in \Pi^s$, the optimal policy $\mu^*(\pi) = \arg \min_{u \in \mathcal{U}} Q(\pi, u) = 2$. So $\bar{\mu}(\pi) = \mu^*(\pi)$ for $\pi \in \Pi^s$ and $\bar{\mu}(\pi) = 1$ otherwise, implying that $\bar{\mu}(\pi)$ is a lower bound for $\mu^*(\pi)$.

J) *Proof of Theorem 11:* Identical to the proof of meta Theorem 13 in Appendix B, under the assumptions $c(e_i, u)$ decreasing in i , (A2), (A3), it follows that $V(\pi; P^{(1)})$ and $V(\pi; P^{(2)})$ are MLR decreasing for $\pi \in \Pi(X)$. We next introduce the following lemma.

Lemma 8: [18, Theorem 2.4] $P^{(1)} \succeq_r P^{(2)}$ implies $P^{(1)'} \pi \succeq_r P^{(2)'} \pi$ where \succeq is defined in (75).

The proof of the lemma is as follows: By definition $P^{(1)'} \pi \succeq_r P^{(2)'} \pi$ is equivalent to

$$\sum_{i \in \mathcal{X}} \sum_{m \in \mathcal{X}} \left(P_{ij}^{(1)} P_{m,j+1}^{(2)} - P_{ij}^{(2)} P_{m,j+1}^{(1)} \right) \pi_i \pi_m \leq 0.$$

Thus, clearly (75) is a sufficient condition for $P^{(1)'} \pi \succeq_r P^{(2)'} \pi$.

Returning to the proof of the theorem, if (A2), (A3) hold, it follows from Lemma 8 and meta Theorem 13 (Statements (1) and (2)) that for actions $u \in \{1, 2\}$

$$\begin{aligned} T(\pi, y, u; P^{(1)}) &\succeq_r T(\pi, y, u; P^{(2)}) \\ \sigma(\pi, \cdot, u; P^{(1)}) &\succeq_s \sigma(\pi, \cdot, u; P^{(2)}). \end{aligned} \quad (94)$$

The rest of the proof is by induction on the value iteration algorithm (74). Assume $V_k(\pi; P^{(1)}) \leq V_k(\pi; P^{(2)})$ for $\pi \in \Pi(X)$. Then from (94)

$$\begin{aligned} V_k\left(T\left(\pi, y, u; P^{(1)}\right); P^{(1)}\right) &\leq V_k\left(T\left(\pi, y, u; P^{(1)}\right); P^{(2)}\right) \\ &\leq V_k\left(T\left(\pi, y, u; P^{(2)}\right); P^{(2)}\right). \end{aligned}$$

Therefore

$$\begin{aligned} \sum_y V_k\left(T\left(\pi, y, u; P^{(1)}\right); P^{(1)}\right) \sigma\left(\pi, y, u; P^{(1)}\right) \\ \leq \sum_y V_k\left(T\left(\pi, y, u; P^{(2)}\right); P^{(2)}\right) \sigma\left(\pi, y, u; P^{(2)}\right). \end{aligned}$$

Next since $C(\pi, u; P^{(1)}) \leq C(\pi, u; P^{(2)})$, it follows that

$$\begin{aligned} & C\left(\pi, u; P^{(1)}\right) \\ &+ \sum_y V_k\left(T\left(\pi, y, u; P^{(1)}\right); P^{(1)}\right) \sigma\left(\pi, y, u; P^{(1)}\right) \\ &\leq C\left(\pi, u; P^{(2)}\right) \\ &+ \sum_y V_k\left(T\left(\pi, y, u; P^{(2)}\right); P^{(2)}\right) \sigma\left(\pi, y, u; P^{(2)}\right). \end{aligned}$$

Taking the minimum with respect to u yields $V_{k+1}(\pi; P^{(1)}) \leq V_{k+1}(\pi; P^{(2)})$.

To prove the second claim of the theorem; first recall from (23) that $\bar{V}(\pi; P)$ is the actual optimal expected cost. Recall that the transformation from $\bar{V}(\pi; P)$ to $V(\pi; P)$ was made to prove that the optimal policy is monotone. It is readily verified that the quickest detection problem (29) satisfies all the assumptions of the theorem. So $V(\pi; P^{(1)}) \leq V(\pi; P^{(2)})$. The actual optimal cost is $\bar{V}(\pi; P) = V(\pi; P) + (\alpha + \beta)(1 - e_1' \pi)$ (see (29)).

Since $(\alpha + \beta)(1 - e_1' \pi)$ is functionally independent of P , it then follows that $\bar{V}(\pi; P^{(1)}) \leq \bar{V}(\pi; P^{(2)})$.

REFERENCES

- [1] R. Amir, "Supermodularity and complementarity in economics: An elementary survey," *Southern Econ. J.*, vol. 71, no. 3, pp. 636–660, 2005.
- [2] A. Banerjee, "A simple model of herd behavior," *Quart. J. Econ.*, vol. 107, pp. 797–817, 1992.
- [3] J. S. Baras and A. Bensoussan, "Optimal sensor scheduling in nonlinear filtering of diffusion processes," *SIAM J. Control Optim.*, vol. 27, no. 4, pp. 786–813, Jul. 1989.
- [4] P. Bartlett and J. Baxter, "Estimation and approximation bounds for gradient-based reinforcement learning," *J. Comput. Syst. Sci.*, vol. 64, no. 1, pp. 133–150, 2002.
- [5] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes—Theory and Applications Information and System Sciences Series*. Princeton, NJ: Prentice-Hall, 1993.
- [6] A. Bensoussan, *Stochastic Control of Partially Observable Systems*. Cambridge, U.K.: Cambridge Univ. Press, 1992.
- [7] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 2000, vol. 1 and 2.
- [8] S. Bikchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as information cascades," *J. Polit. Econ.*, pp. 100:992–1026, 1992.
- [9] A. R. Cassandra, "Exact and Approximate Algorithms for Partially Observed Markov Decision Process," Ph.D. dissertation, Brown Univ., Providence, RI, 1998.
- [10] A. R. Cassandra, Tony's POMDP Page [Online]. Available: <http://www.cs.brown.edu/research/ai/pomdp/index.html>
- [11] C. Chamley, *Rational Herds: Economic Models of Social Learning*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [12] Y. S. Chow, H. Robbins, and D. Siegmund, *Great Expectations: The Theory of Optimal Stopping*. Boston, MA: Houghton Mifflin, 1971.
- [13] D. P. Heyman and M. J. Sobel, *Stochastic Models in Operations Research*. New York: McGraw-Hill, 1984, vol. 2.
- [14] R. J. Elliott, L. Aggoun, and J. B. Moore, *Hidden Markov Models—Estimation and Control*. New York: Springer-Verlag, 1995.
- [15] F. R. Gantmacher, *Matrix Theory*. New York: Chelsea, 1960, vol. 2.
- [16] O. Hernández-Lerma and J. L. Bernard, *Discrete-Time Markov Control Processes: Basic Optimality Criteria*. New York: Springer-Verlag, 1996.
- [17] M. R. James, J. S. Baras, and R. J. Elliott, "Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems," *IEEE Trans. Autom. Control*, vol. 39, no. 4, pp. 780–792, Apr. 1994.
- [18] S. Karlin and Y. Rinott, "Classes of orderings of measures and related correlation inequalities. I. multivariate totally positive distributions," *J. Multivariate Anal.*, vol. 10, pp. 467–498, 1980.
- [19] J. Keilson and A. Kester, "Monotone matrices and monotone Markov processes* 1," *Stochastic Processes Appl.*, vol. 5, no. 3, pp. 231–241, 1977.
- [20] M. Kijima, *Markov Processes for Stochastic Modelling*. New York: Chapman and Hall, 1997.
- [21] V. Krishnamurthy, "Algorithms for optimal scheduling and management of hidden Markov model sensors," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1382–1397, Jun. 2002.
- [22] V. Krishnamurthy and D. V. Djonin, "Optimal threshold policies for multivariate POMDPs in radar resource management," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 3954–3969, Oct. 2009.
- [23] V. Krishnamurthy, F. V. Abad, and K. Martin, "Implementation of gradient estimation to a constrained Markov decision problem," presented at the IEEE Conf. Decision and Control, Maui, HI, 2003.
- [24] V. Krishnamurthy and B. Wahlberg, "POMDP multiarmed bandits—structural results," *Math. Oper. Res.*, vol. 34, no. 2, pp. 287–302, May 2009.
- [25] V. Krishnamurthy and G. Yin, "Recursive algorithms for estimation of hidden Markov models and autoregressive models with Markov regime," *IEEE Trans. Inf. Theory*, vol. 48, no. 2, pp. 458–476, Feb. 2002.
- [26] H. J. Kushner and G. Yin, *Stochastic Approximation Algorithms and Recursive Algorithms and Applications*, 2nd ed. New York: Springer-Verlag, 2003.
- [27] I. Lobel, D. Acemoglu, M. Dahleh, and A. E. Ozdaglar, "Preliminary results on social learning with partial observations," presented at the 2nd Int. Conf. Performance Evaluation Methodologies and Tools, Nantes, France, 2007.
- [28] W. S. Lovejoy, "On the convexity of policy regions in partially observed systems," *Oper. Res.*, vol. 35, no. 4, pp. 619–621, Jul.-Aug. 1987.
- [29] W. S. Lovejoy, "Computationally feasible bounds for partially observed Markov decision processes," *Oper. Res.*, vol. 39, no. 1, pp. 162–175, Jan.-Feb. 1991.
- [30] A. W. Marshall and I. Olkin, *Inequalities: Theory of Majorization and Its Applications*. New York: Academic, 1979.
- [31] G. B. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Statist.*, vol. 14, pp. 1379–1387, 1986.
- [32] A. Muller and D. Stoyan, *Comparison Methods for Stochastic Models and Risk*. Hoboken, NJ: Wiley, 2002.
- [33] M. F. Neuts, *Structured Stochastic Matrices of M/G/1 Type and Their Applications*. New York: Marcel Dekker, 1989.
- [34] C. H. Papadimitrou and J. N. Tsitsiklis, "The complexity of Markov decision processes," *Math. Oper. Res.*, vol. 12, no. 3, pp. 441–450, 1987.
- [35] G. Pflug, *Optimization of Stochastic Models: The Interface between Simulation and Optimization*. Boston, MA: Kluwer, 1996.
- [36] H. V. Poor, "Quickest detection with exponential penalty for delay," *Ann. Statist.*, vol. 26, no. 6, pp. 2179–2205, 1998.
- [37] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer-Verlag, 1993.
- [38] H. V. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [39] K. Premkumar, A. Kumar, and V. V. Veeravalli, "Bayesian quickest transient change detection," to be published.
- [40] V. Raghavan and V. Veeravalli, "Bayesian quickest change process detection," in *Proc. Inst. Symp. Information Theory*, Seoul, 2009, pp. 644–648.
- [41] U. Rieder, "Structural results for partially observed control models," *Meth. Models Oper. Res.*, vol. 35, pp. 473–490, 1991.
- [42] W. Rudin, *Principles of Mathematical Analysis*. New York: McGraw-Hill, 1976.
- [43] A. N. Shiryaev, "On optimum methods in quickest detection problems," in *Theory of Probability and its Applications*. Philadelphia, PA: SIAM, 1963.
- [44] A. N. Shiryaev, *Optimal Stopping Rules*. New York: Springer-Verlag, 1978.
- [45] L. Smith and P. Sorensen, *Informational Herding and Optimal Experimentation Economics*. Oxford, U.K.: Univ. Oxford, 1997, p. 139, Economics Group, Nuffield College.
- [46] L. Smith and P. Sorensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.
- [47] J. Spall, *Introduction to Stochastic Search and Optimization*. Hoboken, NJ: Wiley, 2003.
- [48] D. M. Topkis, *Supermodularity and Complementarity*. Princeton, NJ: Princeton Univ. Press, 1998.
- [49] C. C. White and D. P. Harrington, "Application of Jensen's inequality to adaptive suboptimal design," *J. Optim. Theory Appl.*, vol. 32, no. 1, pp. 89–99, 1980.
- [50] W. Whitt, "Multivariate monotone likelihood ratio and uniform conditional stochastic order," *J. Appl. Probab.*, vol. 19, pp. 695–701, 1982.
- [51] B. Yakir, "Optimal detection of a change in distribution when the observations form a Markov chain with a finite state space," *Lecture Notes-Monograph Ser.*, vol. 23, pp. 346–358, 1994.
- [52] B. Yakir, "A note on optimal detection of a change point in distribution," *Ann. Statist.*, vol. 25, pp. 2117–2126, 1997.
- [53] B. Yakir, A. M. Krieger, and M. Pollak, "Detecting a change in regression: First-order optimality," *Ann. Statist.*, vol. 27, no. 6, pp. 1896–1913, 1999.
- [54] G. Yin, V. Krishnamurthy, and C. Ion, "Regime switching stochastic approximation algorithms with application to adaptive discrete stochastic optimization," *SIAM J. Optim.*, vol. 14, no. 4, pp. 117–1215, 2004.

Vikram Krishnamurthy (S'90–M'91–SM'99–F'05) was born in 1966. He received his bachelor's degree from the University of Auckland, New Zealand in 1988, and the Ph.D. from the Australian National University in 1992. He currently is a professor and Canada Research Chair at the Department of Electrical Engineering, University of British Columbia, Vancouver, BC, Canada.

Dr. Krishnamurthy's current research interests include computational game theory, stochastic control in sensor networks, and stochastic dynamical systems for modeling of biological ion channels and biosensors. Dr. Krishnamurthy currently serves as Editor-in-Chief of the IEEE JOURNAL SELECTED TOPICS IN SIGNAL PROCESSING. He has served as an associate editor for several journals, including the IEEE TRANSACTIONS AUTOMATIC CONTROL and the IEEE TRANSACTIONS ON SIGNAL PROCESSING. In 2009–2010, he served as Distinguished lecturer for the IEEE signal processing society.