

Markov decision processes with observation costs: framework and computation with a penalty scheme

Christoph Reisinger* Jonathan Tam*

December 5, 2023

Abstract

We consider Markov decision processes where the state of the chain is only given at chosen observation times and of a cost. Optimal strategies involve the optimisation of observation times as well as the subsequent action values. We consider the finite horizon and discounted infinite horizon problems, as well as an extension with parameter uncertainty. By including the time elapsed from observations as part of the augmented Markov system, the value function satisfies a system of quasi-variational inequalities (QVIs). Such a class of QVIs can be seen as an extension to the interconnected obstacle problem. We prove a comparison principle for this class of QVIs, which implies uniqueness of solutions to our proposed problem. Penalty methods are then utilised to obtain arbitrarily accurate solutions. Finally, we perform numerical experiments on three applications which illustrate our framework.

1 Introduction

In this article, we examine the observation cost model (OCM) for Markov Decision Processes (MDPs). A cost must be paid in order to observe the state of the underlying MDP, and only then can adjustments be made to the action which influences the dynamics of the MDP. We propose a penalty scheme for efficient numerical computation for the resulting system of equations.

MDPs are mathematical tools that model the optimisation of a random process, in order to maximise the expected profit over time. Applications are common in maintenance, portfolio optimisation, sensor detection, reinforcement learning and more. Most setups implicitly assume a fixed source of information upon which the user relies to select an optimal action. However, such a steady stream of information might not be available in situations where resources are constrained, either by the expensive cost of measurements, or by the impracticality of frequent sampling. This calls for an extra layer of optimisation, where the user has to decide on the optimal observation times of the information source, as well as the optimal sequence of actions to maximise the expected returns.

The literature involving observation control appear across several different fields, and appear under terms such as ‘optimal inspections’, ‘costly observations’ or ‘controllable observations’. To the best of our knowledge, the earliest works appear in [20, 22], which concerns the linear quadratic Gaussian (LQG) problem over a finite horizon with fixed number of measurements, as well as the papers [1, 2], which examines a costly optimal stopping problem in continuous time. Numerous applications have emerged in the literature over the years, which we list (non-exhaustively) below, broadly categorising into the following areas:

*Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK
(christoph.reisinger@maths.ox.ac.uk, jonathan.tam@maths.ox.ac.uk)

- environmental management control models [38–41],
- optimal sampling rates in communications [13, 14],
- optimal sensing problems [23, 32, 37],
- medical treatment cycles [33, 34],
- detection of drift in Brownian motion [4, 5, 10, 11],
- empirical works in reinforcement learning [6, 7, 19].

The standard approach is to formulate the problem in terms of a partially observable Markov decision process (POMDP). Solving the POMDP is then equivalent to solving a fully observable MDP on the belief state [15]. Dynamic programming for the value function leads to a search for the optimal observation time after the currently observed state, as well as the optimal action sequence between the observation times. In this article we will restrict ourselves to consider only *constant* actions between observations. Such an assumption allows us to parametrise the belief state, which takes values on the space of probability measures, with a finite number of variables, so that the augmented state space becomes countable. This assumption in our framework can be relaxed to include a finitely parametrised set of time-inhomogeneous open-loop controls, as we demonstrate in Section 2.4. The general case, however, suffers from the curse of dimensionality: as the time between observations is unbounded, the number of actions to be optimised also grows unbounded. Indeed, non-constant controls between observations are mostly only treated under the LQG framework [9, 32, 37]. For applications, the constant action assumption applies to models where actions cannot be feasibly changed without an accompanying observation, such as the medical treatment applications in [33, 34] or the environmental management control models [38–41].

For the characterisation of the belief MDP for the OCM, our construction includes the time elapsed (since the last observation) as variable within dynamic programming to obtain a low-dimensional Markovian structure. To our knowledge, only the works of [1, 2] and [14] model the OCM in this specific formulation, but the problems considered were restricted to fixed dynamics for the underlying Markov chain. The inclusion of time elapsed as a variable for the augmented Markovian system leads to a system of discrete quasi-variational inequalities (QVI). We assume that the Markov chain takes values in a finite state space \mathcal{X} and that its dynamics are known and are given by the transition matrices $\{P_a\}_{a \in A}$, where A is a finite action set. We also assume that the actions can only be adjusted at the observation times. The one-step reward function is given by $r(x, a) = r_{a,x}$ and the observation cost is given by a constant $c_{\text{obs}} > 0$. The inclusion of time elapsed as a variable in the Markov system leads to a system of discrete quasi-variational inequalities (QVI), which for the discounted infinite horizon problem, reads:

$$\min \left\{ v_{a,x}^n - \gamma v_{a,x}^{n+1} - \left(P_a^n r_a \right)_x, \right. \\ \left. v_{a,x}^n - \left(P_a^n \overline{\gamma v^1 + r} \right)_x + c_{\text{obs}} \right\} = 0, \quad (1.1)$$

where v is the value function, indexed by: $x \in \mathcal{X}$, the state of the chain at the previous observation; $n \in \mathbb{N}_{\geq 1}$, the time elapsed since the previous observation; and $a \in A$, the action applied at the previous observation. The vector $\left(\overline{\gamma v^1 + r} \right)_x = \max_{a \in A} (\gamma v_{a,x}^1 + r_{a,x})$ represents the ‘inner loop’ optimisation over the space of actions after an observation is made.

We note here that our formulation differs from [17], which uses the term ‘self-triggered MDPs’ to refer to the OCM with constant action between observations. There, the time elapsed variable is not considered as part of the augmented Markov system: in our framework we parametrise the belief state by the variables (n, x, a) as in (1.1); in [17] the parametrisation is only with (x, a) . We can interpret our formulation as a special case of unknown initial conditions in a filtering problem, where actions might not have been applied optimally in the past. Such situations

can occur in important settings with external factors and constrained resources — e.g., long NHS (National Health Service in the UK) waiting times especially after the pandemic — which clearly lead to suboptimal times of scans (observations) and treatment (actions). Suboptimal actions can also arise due to users' general inattentiveness and inertia to reacting quickly to information [25, 26].

More generally, we can write the QVI (1.1) obtained from dynamic programming in the following abstract form.

Problem 1.1. Find a function $u : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}^A$, with $u_a(\cdot, \cdot)$ as its components, satisfying

$$\min \{F_a(n, x, u(n, x), u_a), u_a(n, x) - \mathcal{M}u(n, x)\} = 0, \quad n \in \mathbb{N}, x \in \mathcal{X}, a \in A, \quad (1.2)$$

where $F_a : \mathbb{N} \times \mathcal{X} \times \mathbb{R}^A \times \mathbb{R}^{\mathbb{N} \times \mathcal{X}} \rightarrow \mathbb{R}$ satisfies Assumption 1.2, and $\mathcal{M}u : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$ is defined for all $(n, x) \in \mathbb{N} \times \mathcal{X}$ by

$$\mathcal{M}u(n, x) = \left(Q_n \overline{u^1 - c} \right)_x, \quad \left(\overline{u^1 - c} \right)_x = \max_{a \in A} (u_{a,x}^1 - c_{a,x}),$$

for a given vector $c \in \mathbb{R}^{\mathcal{X} \times A}$, and $\{Q_n\} \subset \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is a sequence of substochastic matrices, i.e., with non negative row sums bounded by 1.

Assumption 1.2. The functions $\{F_a\}_{a \in A}$ in Problem 1.1 satisfy:

1. There exists $\beta > 0$ such that for any $n \in \mathbb{N}$, $x \in \mathcal{X}$, bounded functions $u, v : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}^A$ such that $u \leq v$, and $r, s \in \mathbb{R}^A$ such that $\theta := r_a - s_a = \max_{j \in A} (r_j - s_j) \geq 0$, then

$$F_a(n, x, r, u_a + \theta) - F_a(n, x, s, v_a) \geq \beta \theta.$$

2. For any $n \in \mathbb{N}$, $a \in A$, $r \in \mathbb{R}^A$ and continuous bounded function $\phi : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$, the functions $x \mapsto F_a(n, x, r, \phi)$ are continuous and bounded. Moreover, the functions $r \mapsto F_a(n, x, r, \phi)$ are uniformly continuous for bounded r , uniformly with respect to $(n, x) \in \mathbb{N} \times \mathcal{X}$.

The QVI (1.2) is a generalisation of a monotone system with interconnected obstacles [28], which can arise from the discretisation of optimal switching problems in continuous time. In our case, we shall refer to the operator \mathcal{M} as the inspection operator. Much like the systems with interconnected obstacles, the QVI (1.2) is typically not amenable to the use of policy iteration, as the matrices arising from the inspection operator do not necessarily satisfy the M-matrix or weakly chain diagonally dominant conditions [3]. We propose instead a penalty scheme, which sees use on variational inequalities [12, 16, 18] and extensions to HJB VIs [30, 35, 36]/ QVIs [28], as an approximation. Penalty schemes have seen comparable computational performance to policy iteration in HJB QVIs, and is robust to the choice of initial estimates [35, 36]. An adaptation of the penalty scheme to the QVI (1.2) circumvents issues with numerical instabilities arising from computing iterates of the policy update, and the penalised equation can be solved with semismooth Newton methods. We demonstrate in Section 4 that the penalty method achieves quick convergence within a few iterations on a large system whilst also mapping out accurately the optimal policy.

The main contributions of our paper are as follows:

- We formulate the observation cost model (OCM) for Markov decision processes where the time elapsed after an observation is considered as part of the augmented Markov system. We present the optimality equations obtained from dynamic programming for the finite horizon problem, discounted infinite horizon problem, and the respective problems with parameter uncertainty. In all cases the optimality equations are in the form of a QVI, which are structurally different to the Bellman-type equations from existing approaches in the literature.

- We establish a comparison principle for the class of QVIs (1.2), of which the solution to the OCM belongs to. The class of QVIs are a generalisation of monotone systems with interconnected obstacles as seen in [28]. We propose a penalty scheme for this class of QVIs (1.2), and demonstrate the monotone convergence of the penalised solutions towards the solutions of said QVI, thereby establishing constructively the existence of solutions.
- We demonstrate the numerical performance of our model by applying it to the time-discretised version of the HIV-treatment model [34]. Our framework is compatible with the original results, and also shows qualitatively different optimal behaviour when dealing with large observation gaps.

The remainder of the paper is organised as follows. Section 2 sets out the framework for the OCM and establishes the corresponding set of discrete QVIs. A model problem with an explicit solution is also provided to illustrate the setup. We also outline the case of parameter uncertainty in Section 5. In Section 3 we prove a comparison principle for a class of discrete QVIs which subsumes the QVIs obtained in Section 2, as well as outlining the penalty method as a numerical scheme for the QVI. Finally the numerical experiments are presented in detail in Section 4.

1.1 Notation for MDPs and POMDPs

As the goal of the next section is to layout the OCM precisely by formulating the model in terms of a POMDP, we shall quote here some of the standard notation for MDPs, POMDPs and a brief overview of its construction. These are largely taken from [15, Ch 4] and we refer the reader to the references within for further detail.

We will generally be considering Markov decision processes on finite state spaces. Whilst most arguments extend naturally to more general state spaces, we shall focus on the finite setting here to retain a simplified presentation. Let $\mathcal{P}(\mathcal{X})$ denote the space of probability measures over a set \mathcal{X} . If \mathcal{X} is finite, we will also identify $\mathcal{P}(\mathcal{X})$ with the simplex $\Delta_{\mathcal{X}}$.

Definition 1.3. A Markov control model is a tuple $\langle \mathcal{X}, A, p, r \rangle$, where

- \mathcal{X} is the finite *underlying state space*;
- A is the finite *action space*;
- $p : \mathcal{X} \times A \rightarrow \Delta_{\mathcal{X}}$ is the *transition kernel*;
- $r : \mathcal{X} \times A \rightarrow \mathbb{R}$ is the *one-step reward function*.

At each time t , a state $x_t \in \mathcal{X}$ is observed. The controller chooses an action $a_t \in A$ and receives a reward $r(x_t, a_t)$. The system then moves to a new state $x_{t+1} \in \mathcal{X}$ with probability $p(x_{t+1} | x_t, a_t)$ and the process repeats at time $t + 1$. Actions are chosen according to a *policy* $\pi = (\pi_t)_t$, which is a sequence of kernels $\pi_t : H_t \rightarrow A$, where $H_0 := \mathcal{X}$ and $H_t := (\mathcal{X} \times A)^t \times \mathcal{X}$ for $t \geq 1$, known as the *history set* at time t . The set of all policies is denoted by Π . Given an initial state $x_0 \in \mathcal{X}$ and policy $\pi \in \Pi$, by the Ionescu-Tulcea theorem (see [15, Appendix C]), there exists a unique probability measure $\mathbb{P}_{x_0}^{\pi}$ on the canonical sample space $\Omega := H_{\infty} := (\mathcal{X} \times A)^{\infty}$, such that given $\omega = (x_0, a_0, x_1, a_1, \dots) \in \Omega$,

$$\mathbb{P}_{x_0}^{\pi}(\omega) = \delta(x_0) \pi_0(a_0 | x_0) p(x_1 | x_0, a_0) \pi(a_1 | x_0, a_0, x_1) \dots$$

The objective is to maximise an objective function over the set of policies Π , for example, in the finite horizon case,

$$J(\pi, x_0) = \mathbb{E}_{x_0}^{\pi} \left[\sum_{n=0}^N r(x_n, a_n) \right], \quad \pi \in \Pi, \quad x_0 \in \mathcal{X},$$

where $N \in \mathbb{N}$ is the time horizon, and $\mathbb{E}_{x_0}^\pi$ is the expectation under the measure $\mathbb{P}_{x_0}^\pi$. The value function is given by

$$v(x) = \sup_{\pi \in \Pi} J(\pi, x), \quad x \in \mathcal{X}.$$

It is well known that an optimal policy π^* for an MDP is deterministic and Markovian, i.e. there exists deterministic functions $\{\phi_t\}_{t \geq 0}$ such that for $h_t = (x_0, a_0, \dots, x_t) \in H_t$, $\pi_t^*(h_t) = \phi_t(x_t)$, and $v(x) = J(\pi^*, x)$.

In many cases, rather than having full information of the MDP, one instead has access to noisy observations correlated to the underlying MDP. This gives rise to the notion of partially observable Markov decision processes (POMDPs), which can also be described by a given tuple as follows.

Definition 1.4. A partially observable control model is a tuple $\langle \mathcal{X}, \mathcal{O}, A, p, p_0, q, q_0, r \rangle$, where

- \mathcal{X} is the finite *state space*;
- \mathcal{O} is the finite *observation space*;
- A is the finite *action space*;
- $p : \mathcal{X} \times A \rightarrow \Delta_{\mathcal{X}}$ is the *transition kernel*;
- $p_0 \in \Delta_{\mathcal{X}}$ is the *initial distribution*;
- $q : A \times \mathcal{X} \rightarrow \Delta_{\mathcal{O}}$ is the *observation kernel*;
- $q_0 : \mathcal{X} \rightarrow \Delta_{\mathcal{O}}$ is the *initial observation kernel*;
- $r : \mathcal{X} \times A \rightarrow [0, \infty)$ is the *one-step reward function*.

In this setting, given an underlying state $x_t \in \mathcal{X}$, an observation $\bar{x}_t \in \mathcal{O}$ is generated according to the observation kernel $q(\bar{x}_t | a_{t-1}, x_t)$. The controller chooses an action $a_t \in A$ based on their observations, rather than the values of the underlying states. For this, define the *observable history sets*

$$\mathcal{H}_0 := \mathcal{O}, \quad \mathcal{H}_t := (\mathcal{O} \times A)^t \times \mathcal{O}, \quad t \geq 1.$$

A policy for a POMDP is now a sequence of kernels $\pi_t : \mathcal{H}_t \rightarrow A$. Denote the set of policies for the POMDP as Π_{po} . By the Ionescu-Tulcea theorem again, given an initial distribution $p_0 \in \Delta_{\mathcal{X}}$ and policy $\pi \in \Pi_{\text{po}}$, there exists a unique probability measure $\mathbb{P}_{p_0}^\pi$ on the canonical space $\Omega = (\mathcal{X} \times \mathcal{O} \times A)^\infty$ such that for $\omega = (x_0, \bar{x}_0, a_0, x_1, \bar{x}_1, a_1, \dots) \in \Omega$,

$$\mathbb{P}_{p_0}^\pi(\omega) = p_0(x_0) q_0(\bar{x}_0 | x_0) \pi_0(a_0 | \bar{x}_0) p(x_1 | x_0, a_0) q(\bar{x}_1 | x_1, a_0) \pi(x_1 | \bar{x}_0, a_0, \bar{x}_1) \dots$$

The maximisation is now performed over $\pi \in \Pi_{\text{po}}$, with the objective and value function

$$J(\pi, p_0) = \mathbb{E}_{p_0}^\pi \left[\sum_{n=0}^N r(x_n, a_n) \right], \quad v(p) = \sup_{\pi \in \Pi_{\text{po}}} J(\pi, p), \quad \pi \in \Pi_{\text{po}}, \quad p, p_0 \in \Delta_{\mathcal{X}}.$$

Without knowledge of the underlying states, the POMDP is a non-Markovian problem. The standard approach to solve a POMDP is to consider an equivalent (fully observable) problem, known as the *belief MDP*, on the space $Z := \Delta_{\mathcal{X}}$. The Markovian structure is recovered when lifted to the belief MDP, so that classical dynamic programming techniques can be applied. The transition kernel for the *belief state* $z = (z_t)_t$ can be constructed as follows: given the POMDP, construct a kernel $R : Z \times A \rightarrow \mathcal{X} \times \mathcal{O}$ such that

$$R(x, \bar{x} | z, a) = q(\bar{x} | a, x') p(x' | x, a) z(x), \quad (x, \bar{x}) \in \mathcal{X} \times \mathcal{O}, \quad (z, a) \in Z \times A.$$

It can be shown that there exists a kernel $H' : Z \times A \rightarrow \mathcal{X}$, such that R can be disintegrated into

$$R(x, \bar{x} \mid z, a) = H'(x \mid z, a, \bar{x})R'(\bar{x} \mid z, a), \quad (x, \bar{x}) \in \mathcal{X} \times \mathcal{O}, \quad (z, a) \in Z \times A.$$

where R' is the marginal of R on \mathcal{O} . Then, letting $H_{z,a,\bar{x}} = H'(\cdot \mid z, a, \bar{x}) \in Z$, define the kernel $q' : Z \times A \rightarrow Z$ by

$$q'(z' \mid z, a) = \sum_{\bar{x} \in \mathcal{O}} \delta_{H_{z,a,\bar{x}}}(z')R'(\bar{x} \mid z, a), \quad z, z' \in Z, \quad a \in A.$$

Then one takes q' as the transition kernel for the belief state, and construct the initial kernel q'_0 analogously. The belief MDP is then $\langle Z, A, q', q'_0, r_z \rangle$, where $r_z : Z \times A \rightarrow \mathbb{R}$ as $r_z(z, a) = \sum_{x \in \mathcal{X}} r(x, a)z(x)$. The belief state can be interpreted as the conditional distribution of the underlying state x_t , given the observed history $(\bar{x}_0, a_0, \dots, \bar{x}_t)$. Let Π^z be the set of policies for the belief MDP, which are now a sequence of kernels on A , given the history sets $\mathcal{H}_t^z = (Z \times A)^t \times Z$. The objective and value function for the belief MDP are given by

$$J_z(\pi^z, z_0) := \mathbb{E}_{z_0}^{\pi^z} \left[\sum_{n=0}^N r_z(z_n, a_n) \right], \quad v_z(z) = \sup_{\pi^z \in \Pi^z} J_z(\pi^z, z), \quad \pi^z \in \Pi^z, \quad z, z_0 \in Z,$$

where $\mathbb{E}_{z_0}^{\pi^z}$ is the expectation over the canonical space $(Z \times A)^\infty$ under the policy π^z and initial condition $z_0 \in Z$. It can then be shown that policies in Π_{po} are equivalent to policies in Π^z , in the sense that any $h_t \in \mathcal{H}_t$ can be mapped to a corresponding $h_t^z \in \mathcal{H}_t^z$, so that given $\pi^z \in \Pi^z$, one can construct a corresponding $\pi \in \Pi_{\text{po}}$ via

$$\pi(\cdot \mid h_t) := \pi^z(\cdot \mid h_t^z),$$

and the conditional probabilities assigned on the action set A are the same. Moreover, $\pi^* \in \Pi_{\text{po}}$ is optimal for J if and only if $\pi^{*,z} \in \Pi^z$ is optimal for J_z , and it holds that

$$v_z(z) = J_z(\pi^{*,z}, z) = J(\pi^*, z) = v(z), \quad z \in Z = \Delta_{\mathcal{X}}.$$

Thus, when considering a POMDP, it is sufficient to consider its equivalent MDP in the belief state, of which the optimal policy is Markovian with respect to $z = (z_t)_t$.

2 Problem formulation

In the Observation Cost Model, the process evolves sequentially as follows. At each time t , the controller decides if they would like to pay an *observation cost* $c_{\text{obs}} > 0$ to observe the state $x_t \in \mathcal{X}$. If they decide to do so, then the controller applies the action $a_t \in A$ according to some suitable policy π , and receives a reward $r(x_t, a_t)$. If they decide not to observe, then no cost is paid, but the controller cannot change the action value, so that $a_t = a_{t-1}$. We assume that the reward $r(x_t, a_t)$ is collected and ‘locked in’ at time t , but is not observable to the user if x_t is not observable. In both instances, the system moves to a new state $x_{t+1} \in \mathcal{X}$ according to the transition kernel $p(x_{t+1} \mid x_t, a_t)$.

We now formally write down the objective function, and shall make precise the terms appearing within in the rest of this section. In view of the description above, the controller wishes to maximise (for example, in the finite horizon case)

$$\mathbb{E}^\pi \left[\sum_{n=0}^N (r(x_n, a_n) - i_n \cdot c_{\text{obs}}) \right], \quad (2.1)$$

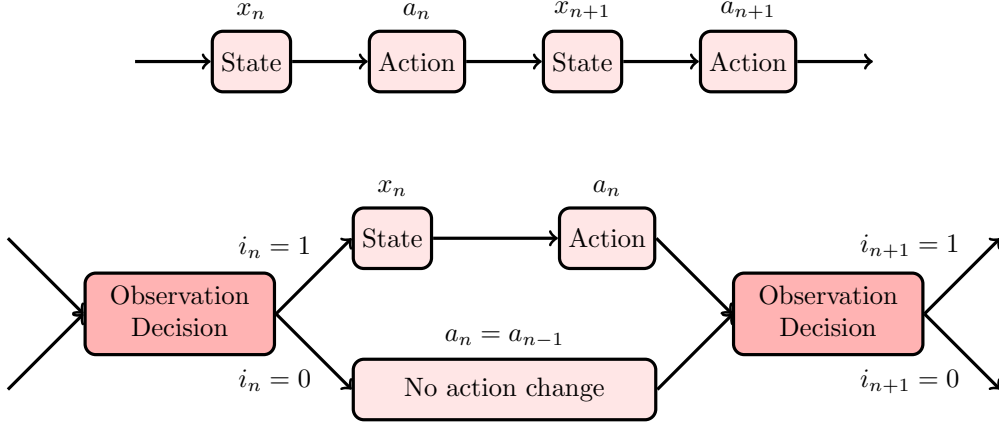


Figure 2.1: *Top: Standard formulation with full observations; Bottom: The inclusion of observation costs leads to an extra decision step.*

where π a suitably admissible control policy to be made precise later. The sequence $i = (i_n)_n$ takes values in $\mathcal{I} := \{0, 1\}$ and will be referred to as the *inspection values*. A value of $i_n = 1$ represents an observation made at time n , so that observation cost c_{obs} is deducted from the total reward in (2.1). Conversely no observations are made if $i_n = 0$. Figure 2.1 illustrates the sequential flow of a standard MDP, compared to that of the OCM.

We now proceed to establish the OCM as a non-standard form of a POMDP, in order to fully make sense of (2.1). A policy should output an action value $a_n \in A$, as well as an inspection value $i_n \in \mathcal{I}$. The observation space is represented by $\mathcal{X} \cup \{\emptyset\}$: either the underlying chain with values in \mathcal{X} is observed, or the dummy variable \emptyset nothing is observed, which represents the case of no observations. A final but subtle point is the difference in the sequential structure of the OCM compared to a POMDP. In the case of a POMDP, first a state x_n is generated, follow by the observation \bar{x}_n , and then the action a_n . Thus a realisation on the canonical space of a POMDP looks like:

$$(x_0, \bar{x}_0, a_0, x_1, \bar{x}_1, a_1, \dots).$$

In the OCM, as depicted in Figure 2.1, the observation occurs *after* the inspection value, after which the action value follows. Thus a realisation of the system of an OCM will instead look like

$$(x_0, i_0, \bar{x}_0, a_0, x_1, i_1, \bar{x}_1, a_1, \dots). \quad (2.2)$$

In order to obtain the sequential structure of ‘state - observation - action’ for the OCM, we augment the sequence (2.2) with fictitious state and observation values, and treat both a_n and i_n as an ‘action’. By suitably augmenting the transition and observation kernels, the OCM takes the form of a POMDP, over the timescale of $\frac{1}{2}\mathbb{N}$. This augmented sequence then takes the form of

$$\begin{aligned} & (x_0, \bar{x}_0, \pi_0, x_{1/2}, \bar{x}_{1/2}, \pi_{1/2}, x_1, \bar{x}_1, \pi_1, \dots) \\ & := (x_0, \bar{x}_0, i_0, x_0, \bar{x}_0, a_0, x_1, \bar{x}_1, i_1, \dots). \end{aligned}$$

This leads us to the following definition.

Definition 2.1. Given an MDP $\langle \mathcal{X}, A, p, r \rangle$, the associated observation cost model (OCM) is defined as the POMDP $\langle \mathcal{X}, \mathcal{X}_{\emptyset}, \mathcal{A}, p, p_0, q, q_0, r_{\text{obs}} \rangle$ (see Definition 1.4), on the time scale $\frac{1}{2}\mathbb{N} = \{0, 1/2, 1, \dots\}$, where

- $\mathcal{X}_\emptyset = \mathcal{X} \cup \{\emptyset\}$ is the *observation space*, with \emptyset a dummy variable representing no observations.
- \mathcal{A} is the *disjoint union* of A and \mathcal{I} , with time dependent admissible sets given by $\mathcal{A}(n) = \mathcal{I}$ and $\mathcal{A}(n + \frac{1}{2}) = A$;
- $p : \mathcal{X} \times \mathcal{A} \rightarrow \Delta_{\mathcal{X}}$ is the transition kernel, with its domain extended to $\mathcal{X} \times \mathcal{A}$ by defining

$$p(\cdot \mid x, i) = \delta_x(\cdot), \quad i \in \mathcal{I}.$$

- $q : \mathcal{A} \times \mathcal{X} \rightarrow \Delta_{\mathcal{X}_\emptyset}$ is the *observation kernel*, given by

$$\begin{aligned} q(\cdot \mid a, x) &= \delta_\emptyset(\cdot), \quad a \in A, \\ q(\cdot \mid i, x) &= i \cdot \delta_x(\cdot) + (1 - i) \delta_\emptyset(\cdot), \quad i \in \mathcal{I}. \end{aligned}$$

- $r_{\text{obs}} : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ is the *one-step reward function* given by

$$\begin{aligned} r_{\text{obs}}(x, a) &= r(x, a), \quad a \in A, \\ r_{\text{obs}}(x, i) &= -i \cdot c_{\text{obs}}, \quad i \in \mathcal{I}. \end{aligned}$$

The kernels p and q above are defined such that transitions of the underlying chain only occurs at the integer steps, and new observations at the half steps. Let us write the observable history sets here as

$$\mathcal{H}_0 = \mathcal{X}_\emptyset, \quad \mathcal{H}_t = (\mathcal{X}_\emptyset \times \mathcal{A})^{2t} \times \mathcal{X}_\emptyset, \quad t \in \frac{1}{2}\mathbb{N}.$$

To be precise, we should only consider *admissible* history sets, that is state-action pairs that satisfy the constraints $\mathcal{A}(n) = \mathcal{I}$ and $\mathcal{A}(n + \frac{1}{2}) = A$, but we shall take this assumption as implicit for ease of notation. A policy π is hence a sequence of kernels $\pi_t : \mathcal{H}_t \rightarrow \mathcal{A}$.

As before, a policy π as defined above and an initial distribution $p_0 \in \Delta_{\mathcal{X}}$ induces a unique measure $\mathbb{P}_{p_0}^\pi$ on the canonical sample space $\Omega = (\mathcal{X} \times \mathcal{X}_\emptyset \times \mathcal{A})^\infty$. With slight abuse of notation, we also write π for the values applied by the policy. Then, the action values $a = (a_n)_n$ and inspection values $i = (i_n)_n$ appearing in (2.1) can be recovered by defining

$$i_n = \pi_n, \quad a_n = \pi_{n+\frac{1}{2}}.$$

As we are assuming in the OCM that actions remain constant between new observations. We will have to consider a smaller class of admissible policies.

Definition 2.2. An **admissible policy** $\pi = (\pi_t)_t$ for the OCM is a sequence of kernels $\pi_t : \mathcal{H}_t \rightarrow \Delta_{\mathcal{A}}$ which satisfies the following: for $n \in \mathbb{N}$, if $h_{n+\frac{1}{2}} = (\bar{x}_0, \pi_0, \dots, \bar{x}_{n-\frac{1}{2}}, \pi_{n-\frac{1}{2}}, \bar{x}_n, \pi_n, \bar{x}_{n+\frac{1}{2}})$ with $\pi_n = i_n = 0$, then

$$\pi_{n+\frac{1}{2}} \left(\cdot \mid h_{n+\frac{1}{2}} \right) = \delta_{\pi_{n-\frac{1}{2}}} = \delta_{a_{n-1}}.$$

The set of admissible policies for the observation cost model is denoted Π_{obs} .

With the above setup, we can give a full meaning to the expression (2.1) by writing

$$J(\pi, p_0) := \mathbb{E}_{p_0}^\pi \left[\sum_{n=0}^{2N} r_{\text{obs}}(x_{n/2}, \pi_{n/2}) \right] = \mathbb{E}_{p_0}^\pi \left[\sum_{n=0}^N (r(x_n, a_n) - i_n \cdot c_{\text{obs}}) \right], \quad \pi \in \Pi_{\text{obs}}, \quad p_0 \in \Delta_{\mathcal{X}}.$$

Remark 2.3. *Regarding the initial distribution p_0 and observation kernel q_0 : for the purposes of this paper, we will assume that an observation (made at some previous time) is always available. This allows a consistent characterisation of the belief state by a finite tuple, which leads to a system of finite-dimensional QVIs in Sections 2.1 and 2.2. Thus, we will only consider initial kernels p_0 that is in the form of some n -step transition probabilities of the kernel p , and the initial observation kernel q_0 will be taken as the Dirac measure $q_0(\cdot \mid x) = \delta_x(\cdot)$.*

Remark 2.4. *The formulation with the half time-steps and the inclusion of fictitious state/ observation variables above are strictly a theoretical construct, such that the OCM can be reframed as a POMDP, and thus allowing us to directly appeal to standard results to formulate dynamic programming. It will be shown later that the half steps become redundant once dynamic programming is established, and the value function will only need to be considered over the integer steps.*

2.1 Finite horizon problem

For the finite horizon problem, let $N \in \mathbb{N}$ be the time horizon. For a policy $\pi \in \Pi_{\text{obs}}$, the objective function is

$$\mathbb{E}_{p_0}^{\pi} \left[\sum_{n=0}^{2N} r_{\text{obs}}(x_{n/2}, \pi_{n/2}) \right] = \mathbb{E}_{p_0}^{\pi} \left[\sum_{n=0}^N (r(x_n, a_n) - i_n \cdot c_{\text{obs}}) \right].$$

With the OCM problem characterised as a POMDP, we consider the belief MDP, with the belief state given by $\mathbb{P}_{p_0}^{\pi}(x_t | h_t)$, where $h_t \in \mathcal{H}_t$. By the Markov property, the belief state is fully characterised by the controller's most recently observed information, in that

$$\mathbb{P}_{p_0}^{\pi}(x_t | h_t) = p^{(\lfloor t \rfloor - k)}(x_t | x_k, a_k),$$

where $p^{(\lfloor t \rfloor - k)}$ is the $(\lfloor t \rfloor - k)$ -step transition kernel of the underlying process X , and $k \in \mathbb{N}$ is the last occurrence such that $i_k = 1$. Note that $k = t$ implies an immediate observation, and the belief state reduces trivially to a Dirac measure at x_t . The belief state therefore has a finite dimensional parametrisation, and we can consider instead the **augmented state** $y = (y_t)_{t \in \frac{1}{2}\mathbb{N}}$:

$$y_t := \begin{cases} (k, x_k, a_k), & \text{if } \mathbb{N} \ni k = \arg \max_{n \leq t} \{i_n = 1\} \neq t, \\ (t, x_t, \emptyset), & \text{otherwise,} \end{cases}$$

where \emptyset also acts as a dummy variable here. The three components of y_t represent the most recent observation time k , with the correspondingly observed state $x_k \in \mathcal{X}$, and applied action $a_k \in A$.

Given this equivalence in representation, we can consider the OCM as an MDP with the tuple $\langle \mathcal{Y}, \mathcal{A}, p_y, r_y \rangle$ on the timescale $\frac{1}{2}\mathbb{N}$, where

- $\mathcal{Y} := \mathbb{N} \times \mathcal{X} \times A_{\emptyset}$ is the *augmented state space*;
- \mathcal{A} is the *disjoint* union of A and \mathcal{I} , with admissible sets $\mathcal{A}(n) = \mathcal{I}$ and $\mathcal{A}(n + \frac{1}{2}) = A$;
- $p_y = (p_{y,t})_{t \in \frac{1}{2}\mathbb{N}}$ is a (time-inhomogeneous) transition kernel on \mathcal{Y} given $\mathcal{Y} \times \mathcal{A}$: for $n \in \mathbb{N}$, $y = (k, x, a)$, $\hat{y} \in \mathcal{Y}$, $i \in \mathcal{I}$, and $a' \in A$,

$$p_{y,n}(\hat{y} | y, i) = i \cdot p^{(n-k)}(\hat{x} | x, a) \mathbb{1}_{\{\hat{y}=(n,\hat{x},\emptyset)\}} + (1-i) \mathbb{1}_{\{\hat{y}=y\}}, \quad (2.3)$$

$$p_{y,n+\frac{1}{2}}(\hat{y} | y, a') = \mathbb{1}_{\{\hat{y}=(k,x,a'), a'=a\}}, \quad (2.4)$$

and for $y = (n, x, \emptyset)$,

$$p_{y,n+\frac{1}{2}}(\hat{y} | y, a') = \mathbb{1}_{\{\hat{y}=(n,x,a')\}}. \quad (2.5)$$

- $r_y = (r_{y,t})_{t \in \frac{1}{2}\mathbb{N}}$ is a time-dependent one-step reward function on $\mathcal{Y} \times \mathcal{A}$: for $y = (k, x, a)$, $i \in \mathcal{I}$, $a' \in A$,

$$\begin{aligned} r_{y,n}(y, i) &= -i \cdot c_{\text{obs}}, \\ r_{y,n+\frac{1}{2}}(y, a') &= \sum_{x' \in \mathcal{X}} r(x', a') p^{(n-k)}(x' | x, a), \end{aligned} \quad (2.6)$$

and for $y = (n, x, \emptyset)$,

$$r_{y, n+\frac{1}{2}}(y, a') = r(x, a'). \quad (2.7)$$

In this augmented problem, define its observable history sets as $\mathcal{H}_t^y = (\mathcal{Y} \times \mathcal{A})^{2t} \times \mathcal{Y}$, $t \in \frac{1}{2}\mathbb{N}$. Policies π^y are then a sequence of kernels $\pi_t^y : \mathcal{Y} \rightarrow \Delta_{\mathcal{A}}$. For the set of admissible policies of this augmented MDP, we will have to consider the corresponding ‘image’ of Π_{obs} . This in turn, is equivalent to imposing a state constraint on the admissible action sets:

$$\mathcal{A}(n + 1/2, (k, x, a)) = \{a\}, \quad n \in \mathbb{N}, \quad a \in A.$$

With the above constraints noted, we shall not distinguish between a policy $\pi \in \Pi_{\text{obs}}$ and its corresponding policy π^y in the augmented MDP, and write π for both.

For the finite horizon problem, the objective function for the augmented state MDP is

$$J(t, y, \pi) = \mathbb{E}_y^\pi \left[\sum_{n=t}^{2N} r_{y, \frac{n}{2}} \left(y_{\frac{n}{2}}, \pi_{\frac{n}{2}} \right) \right], \quad 0 \leq t \leq 2N, \quad y \in \mathcal{Y}, \quad \pi \in \Pi_{\text{obs}},$$

with value function

$$v(t, y) = \sup_{\pi \in \Pi_{\text{obs}}} J(t, y, \pi). \quad (2.8)$$

The proposition below shows the dynamic programming equation in the form of a quasi-variational inequality, for which the value function for the OCM satisfies.

Proposition 2.5. *For $n \in \mathbb{N}$, $y = (k, x, a) \in \mathcal{Y}$, define $v_{a,x}^{n,k} = v(n, y)$ as in (2.8). Then the value function satisfies the following quasi-variational inequality (QVI): for all $0 \leq k < n \leq N-1$, $x \in \mathcal{X}$, and $a \in A$,*

$$\min \left\{ v_{a,x}^{n,k} - v_{a,x}^{n+1,k} - \left(P_a^{n-k} r_a \right)_x, \right. \\ \left. v_{a,x}^{n,k} - \left(P_a^{n-k} \overline{v^{n+1,n} + r} \right)_x + c_{\text{obs}} \right\} = 0, \quad (2.9)$$

with the terminal condition

$$v_{a,x}^{N,k} = \left(P^{N-k} r_a \right)_x,$$

where P_a^n is the n -step transition matrix with constant action a , and

$$(r_a)_x = r(x, a), \quad \bar{r}_x = \max_{a \in A} r(x, a), \\ \left(\overline{v^{n+1,n} + r} \right)_x = \max_{a \in A} (v_{a,x}^{n+1,n} + r_{a,x}).$$

Proof. A standard application of dynamic programming gives us

$$v(t, y) = \sup_{\pi \in \Pi_{\text{obs}}} \left\{ r_{y,t}(y, \pi_t) + \mathbb{E}^\pi \left[v \left(t + \frac{1}{2}, y_{t+\frac{1}{2}} \right) \right] \right\}.$$

Expanding explicitly, for $n \in \mathbb{N}$ and $y = (k, x, a) \in \mathcal{Y}$,

$$v(n, (k, x, a)) = \max \left\{ -c_{\text{obs}} + \max_{a \in A} \mathbb{E}^a \left[v \left(n + \frac{1}{2}, (n, x_n, \emptyset) \right) \right], \quad v \left(n + \frac{1}{2}, (k, x, a) \right) \right\},$$

where \mathbb{E}^a is the expectation taken with respect to a constant $a \in A$. Furthermore

$$\begin{aligned} v\left(n + \frac{1}{2}, (k, x, a)\right) &= \mathbb{E}^a[r(x_n, a)] + v(n+1, (k, x, a)), \\ v\left(n + \frac{1}{2}, (n, x, \emptyset)\right) &= \max_{a \in A} \{r(x, a) + v(n+1, (n, x, a))\} \end{aligned}$$

Thus, by combining the inspection stage and the action stage together, and rearranging the terms accordingly, we obtain the QVI in the desired form. \square

The optimal policy π^* is then Markovian with respect to the augmented state y , i.e. the optimal policy depends on the most recent observation. Given a solution to the QVI (2.9), one can retrieve the optimal policy at time n , by first finding the region where the minimum is achieved, which determines if an inspection is optimal. If an inspection is optimal, one observes the latest state, say x , and the optimal action is given by $\arg \max_{a \in A} (v_{a,x}^{n+1,n} + r_{a,x})$.

Remark 2.6 (On the relation between observation costs and switching costs). *Due to the nature of the piecewise constant controls, the OCM resembles an optimal switching problem, where switching costs $\{g_{ij}\}_{i,j \in A}$ are present, such that g_{ij} is paid each time action i is changed to action j . There are two main differences between the two problems. Firstly, for the OCM, the optimal policy on both inspections and actions have to take into account the lack of updated observations. Secondly, in switching problems, g_{ii} is assumed to be 0 for all $i \in A$, i.e., no cost is paid if no switching occurs. This contrasts to the OCM, where $c_{\text{obs}} > 0$ is paid at each observation, regardless of the action applied afterwards. Indeed, one may decide that no change in action is required upon observation, but nonetheless the observation cost has to be paid upfront in order to arrive at such a conclusion. We can incorporate switching costs $\{g_{ij}\}_{i,j \in A}$ into the OCM by writing the augmented reward in the form (see (2.6, 2.7))*

$$\begin{aligned} r_{y,n+\frac{1}{2}}(y, a') &= \sum_{x' \in \mathcal{X}} r(x', a') p^{(n-k)}(x' | x, a) - g_{a,a'}, & y = (k, x, a), \\ r_{y,n+\frac{1}{2}}(y, a') &= r(x, a') - g_{a,a'}, & y = (n, x, \emptyset). \end{aligned}$$

2.2 Infinite horizon problem

For the discounted infinite horizon problem of the OCM, we will have to consider the appropriate stationary formulations. This can easily be obtained by further considering (n, y_n) as an augmented state. Now recall that the transition kernel of y in (2.3) to (2.5) depends on n and k strictly through the difference $n - k$. Hence, after relabelling, it is sufficient to consider $y = (n, x, a)$ as the augmented state, where here n now represents the time elapsed from the previous observation, rather than the standard linear passage of time. The objective function of this equivalent MDP is

$$J(y, \pi) = \mathbb{E}^\pi \left[\sum_{n=0}^{\infty} \gamma^{\lfloor \frac{n}{2} \rfloor} r_y \left(y_{\frac{n}{2}}, \pi_{\frac{n}{2}} \right) \right], \quad y \in \mathcal{Y}, \pi \in \Pi_{\text{obs}},$$

where $\gamma \in (0, 1)$ is the discount factor, the value function is

$$v(y) = \sup_{\pi \in \Pi_{\text{obs}}} J(y, \pi). \quad (2.10)$$

This gives us the following QVI for the value function, which we shall state here without proof.

Proposition 2.7. For $y = (n, x, a) \in \mathcal{Y}$, define $v_{a,x}^n = v(y)$. Then the value function (2.10) satisfies the following quasi-variational inequality (QVI): for all $n \geq 1$, $x \in \mathcal{X}$, and $a \in A$,

$$\min \left\{ v_{a,x}^n - \gamma v_{a,x}^{n+1} - \left(P_a^n r_a \right)_x, \right. \\ \left. v_{a,x}^n - \left(P_a^n \overline{\gamma v^1 + r} \right)_x + c_{\text{obs}} \right\} = 0, \quad (2.11)$$

where P_a^n is the n -step transition matrix with constant action a , and

$$(r_a)_x = r(x, a), \quad \bar{r}_x = \max_{a \in A} r(x, a) \\ \left(\overline{\gamma v^1 + r} \right)_x = \max_{a \in A} (\gamma v_{a,x}^1 + r_{a,x}).$$

Note that the QVI (2.11) is defined on the infinite domain $\mathbb{N}_{\geq 1} \times \mathcal{X} \times A$. In practice, we will have to truncate this domain for the time variable. A natural boundary condition is to enforce an inspection of the underlying chain after some large time N has elapsed. This is equivalent to further restricting the admissible policies in Π_{obs} to those such that $i_N = 1$.

2.3 Approximating the infinite horizon QVI by truncation

Note that the QVI (2.11) is defined on the infinite domain $\mathbb{N}_{\geq 1} \times \mathcal{X} \times A$. In practice, we will have to truncate this domain for the time variable. A natural boundary condition is to enforce an inspection of the underlying chain after some large time N has elapsed.

We can justify the truncation in the following way. For the original untruncated problem, we can write the reward functional and value function as

$$J(\pi, p) = \mathbb{E}_p^\pi \left[\sum_{n=0}^{\infty} \gamma^n (r(x_n, a_n) - i_n \cdot c_{\text{obs}}) \right], \quad \pi \in \Pi_{\text{obs}}, \quad p \in \Delta \mathcal{X}, \\ v(p) = \sup_{\pi \in \Pi_{\text{obs}}} J(\pi, p).$$

For the truncated problem, enforcing a maximum time for inspection is equivalent to further restricting the set of admissible policies within Π_{obs} . To this end, define $\Pi_{\text{obs}}^N \subset \Pi_{\text{obs}}$ such that any $\pi^N = (a^N, i^N) \in \Pi_{\text{obs}}^N$ satisfies

$$\mathbb{P}(\min\{n > k : i_n^N = 1\} \leq k + N \mid i_k^N = 1) = 1 \quad \text{for all } k \geq 0.$$

Then we can define the truncated value function

$$v^N(p) := \sup_{\pi \in \Pi_{\text{obs}}^N} J(\pi, p).$$

Proposition 2.8. The truncated value functions v^N are monotonically increasing in N , and furthermore $v^N \uparrow v$ pointwise.

Proof. The first part is clear as $\Pi_{\text{obs}}^M \subset \Pi_{\text{obs}}^N \subset \Pi_{\text{obs}}$ for any $M < N$, so that v^N is monotonically increasing and bounded above by v . Therefore v^N converges to a limit v^* and

$$v^*(p) = \lim_{N \rightarrow \infty} v^N(p) \leq v(p).$$

To show the reverse inequality, let $\{\pi_n\} \subset \Pi_{\text{obs}}$ be a maximising sequence for v . As both \mathcal{X} and A are finite, for each $n > 0$, we can choose $N_n \in \mathbb{N}$ such that

$$\left| \sum_{n=N_n}^{\infty} \gamma^n (r(x, a) - c_{\text{obs}}) \right| \leq 2^{-n-1} \quad \text{for all } x \in \mathcal{X} \text{ and } a \in A.$$

Then, choose $\pi_n^N \in \Pi_{\text{obs}}^N$ such that $\pi_n^N(k) = \pi_n(k)$ for all $k \leq N_n$, so that

$$\begin{aligned} J(\pi_n, p) &\leq J(\pi_n^N, p) + 2^{-n} \\ &\leq v^{N_n}(y) + 2^{-n}. \end{aligned}$$

But since $\{\pi_n\}$ is a maximising sequence for v , and $N_n \rightarrow \infty$ by construction, letting $n \rightarrow \infty$ on both sides leads to

$$v(y) \leq \lim_{n \rightarrow \infty} v^{N_n}(y) = v^*(y),$$

as required. \square

2.4 Parametrised time-inhomogeneous actions

Here, we show how the assumption of constant controls between observations can be relaxed to allow for parametrised time-dependent actions. Optimising over open-loop actions in general leads to an unbounded number of parametrising variables for the belief state y . For tractability, we consider instead a set of parametrised time-inhomogeneous open-loop actions. This can be incorporated into the previous framework as follows. Let Θ be a finite parameter set. Let $\{f_\theta\}_{\theta \in \Theta}$ be a parametrised set of open-loop actions, with $f_\theta : \mathbb{N} \rightarrow A$. Then we write the OCM POMDP as $\langle \mathcal{Y}, \mathcal{A}, p_y, r_y \rangle$ on the timescale $\frac{1}{2}\mathbb{N}$, but now with \mathcal{A} being the union of Θ and \mathcal{I} , with admissible sets $\mathcal{A}(n) = \mathcal{I}$ and $\mathcal{A}(n + \frac{1}{2}) = \Theta$. Then the augmented transition kernel $p_y = (p_{y,t})_{t \in \frac{1}{2}\mathbb{N}}$ can be written, for $n \in \mathbb{N}$, $y = (k, x, \theta)$, $\hat{y} \in \mathcal{Y}$, $i \in \mathcal{I} = \{0, 1\}$, and $\theta' \in \Theta$, as

$$\begin{aligned} p_{y,n}(\hat{y} \mid y, i) &= i \cdot p^{(n-k)}(\hat{x} \mid x, f_\theta) \mathbb{1}_{\{\hat{y}=(n,\hat{x},\emptyset)\}} + (1-i) \mathbb{1}_{\{\hat{y}=y\}}, \\ p_{y,n+\frac{1}{2}}(\hat{y} \mid y, \theta') &= \mathbb{1}_{\{\hat{y}=(k,x,\theta'), a'=a\}}, \end{aligned}$$

and for $y = (n, x, \emptyset)$,

$$p_{y,n+\frac{1}{2}}(\hat{y} \mid y, \theta') = \mathbb{1}_{\{\hat{y}=(n,x,\theta')\}},$$

where the term $p^{(n-k)}(\hat{x} \mid x, f_\theta)$ represents the $(n-k)$ -step transition kernel starting from $x \in \mathcal{X}$, applying the successive actions $f_\theta(1), \dots, f_\theta(n-k)$. A similar replacement can be done with the augmented reward r_y . Then the QVI can be rewritten, for the infinite horizon case for example, as

$$\begin{aligned} \min \left\{ v_{\theta,x}^n - \gamma v_{\theta,x}^{n+1} - \left(Q_\theta^n r_{f_\theta(n)} \right)_x, \right. \\ \left. v_{\theta,x}^n - \left(Q_\theta^n \overline{\gamma v^1 + r} \right)_x + c_{\text{obs}} \right\} = 0, \end{aligned}$$

with $Q_\theta^n = \prod_{k=1}^n P_{f_\theta(k)}$.

2.5 Toy problem

To illustrate the framework, we present a model problem involving a two-state Markov chain and give an explicit solution. We assume the following setup:

- the state space $\mathcal{X} = \{0, 1\}$;
- the action space $A = \{0, 1\}$;
- the reward function $r(x, a) = a \cdot x + (1-a)(a-x)$;

– the transition matrix

$$P_a = \begin{bmatrix} 0 & 1 \\ a + p(1-a) & (1-p)(1-a) \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

where $p \in (0, 1)$.

This can be seen as a model for a maintenance problem, to find an optimal interval for inspecting equipment to avoid wear and tear over time. The reward function r gives a reward of 1 when the state and action values are the same, and zero otherwise. If no changes are made to the action, the chain eventually arrives at the absorbing state which does not incur any reward. Figure 2.2 illustrates the chain for the case $p = 0.9$.

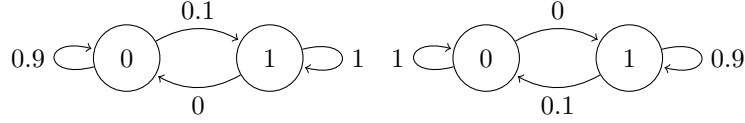


Figure 2.2: Illustration of the two-state Markov chain. Left: $a = 0$; Right: $a = 1$.

Consider the infinite horizon problem. The QVI for $x = 0$ and $a = 0$ is

$$\min \left\{ v_{0,0}^n - \gamma v_{0,0}^{n+1} - \left(P_0^n r_0 \right)_x, \right. \\ \left. v_{0,0}^n - \left(P_0^n \gamma v^1 + r \right)_x + c_{\text{obs}} \right\} = 0. \quad (2.12)$$

Due to the symmetry of the problem, we have $v_{0,0}^n = v_{1,1}^n$. Moreover, it is clear that given the knowledge of x_n , the optimal action is to set $a_n = x_n$. Hence we can write (2.12) as

$$\min \{ v_{0,0}^n - \gamma v_{0,0}^{n+1} - p^n, v_{0,0}^n - \gamma v_{0,0}^1 - 1 + c_{\text{obs}} \} = 0$$

or simply by writing $v(n) = v_{0,0}^n$:

$$v(n) = \max \{ p^n + \gamma v(n+1), 1 - c_{\text{obs}} + \gamma v(1) \}.$$

Let T be the first optimal inspection time (where by convention $T = \infty$ if it is optimal to never inspect). The value function is given recursively by

$$v(n) = \begin{cases} 1 - c_{\text{obs}} + \gamma v(1), & \text{if } n \geq T; \\ p^n + \gamma v(n+1), & \text{otherwise.} \end{cases} \quad (2.13)$$

Solving the above for $v(1)$, we obtain the explicit solution

$$v(1) = \max \left\{ \sup_{m \geq 2} \left(\frac{p \sum_{k=0}^{m-2} (\gamma p)^k + \gamma^{m-1} (1 - c_{\text{obs}})}{1 - \gamma^m} \right), \frac{1 - c_{\text{obs}}}{1 - \gamma} \right\}, \quad (2.14)$$

from which $v(n)$ for $n \geq 1$ can be calculated from (2.13). The first term in (2.14) is a geometric series, where $p \sum_{k=0}^{m-2} (\gamma p)^k + \gamma^{m-1} (1 - c_{\text{obs}})$ is the expected returns across the optimal inspection interval, and γ^m is the discount factor over the whole interval. We can interpret (2.14) as searching for the optimal inspection interval to maximise the sum of the rewards, minus the observation cost.

3 Comparison principle and penalisation

In this section, we consider the well-posedness of solutions to the QVI for the OCM. We establish uniqueness by proving a comparison principle, and existence by constructing arbitrarily close approximations via truncation and penalisation. To this end, we shall recast the QVIs in Section 2 in a more abstract form.

Problem 3.1. Find a function $u : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}^A$, with $u_a(\cdot, \cdot)$ as its components, satisfying

$$\min \{F_a(n, x, u(n, x), u_a), u_a(n, x) - \mathcal{M}u(n, x)\} = 0, \quad n \in \mathbb{N}, x \in \mathcal{X}, a \in A, \quad (3.1)$$

where $F_a : \mathbb{N} \times \mathcal{X} \times \mathbb{R}^A \times \mathbb{R}^{\mathbb{N} \times \mathcal{X}} \rightarrow \mathbb{R}$ satisfies Assumption 3.2, and $\mathcal{M}u : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$ is defined for all $(n, x) \in \mathbb{N} \times \mathcal{X}$ by

$$\mathcal{M}u(n, x) = \left(Q_n \overline{u^1 - c} \right)_x, \quad \left(\overline{u^1 - c} \right)_x = \max_{a \in A} (u_{a,x}^1 - c_{a,x}), \quad (3.2)$$

for a given vector $c \in \mathbb{R}^{\mathcal{X} \times A}$, and $\{Q_n\} \subset \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is a sequence of substochastic matrices, i.e., with non negative row sums bounded by 1.

Assumption 3.2. The functions $\{F_a\}_{a \in A}$ in Problem 3.1 satisfy:

1. There exists $\beta > 0$ such that for any $n \in \mathbb{N}$, $x \in \mathcal{X}$, bounded functions $u, v : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}^A$ such that $u \leq v$, and $r, s \in \mathbb{R}^A$ such that $\theta := r_a - s_a = \max_{j \in A} (r_j - s_j) \geq 0$, then

$$F_a(n, x, r, u_a + \theta) - F_a(n, x, s, v_a) \geq \beta \theta.$$

2. For any $n \in \mathbb{N}$, $a \in A$, $r \in \mathbb{R}^A$ and continuous bounded function $\phi : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$, the functions $x \mapsto F_a(n, x, r, \phi)$ are continuous and bounded. Moreover, the functions $r \mapsto F_a(n, x, r, \phi)$ are uniformly continuous for bounded r , uniformly with respect to $(n, x) \in \mathbb{N} \times \mathcal{X}$.

Assumption 3.2 are monotonicity and regularity assumptions typically used in viscosity solution theory [8], which applies to standard approximation schemes, such as finite differences. The monotonicity condition arises naturally, for example, from the discretisation of QVIs in continuous time involving switching controls. The operator \mathcal{M} , which shall be referred to as the *inspection operator*, is a non-local term in the QVI that couples the solution u across the different action values. In general, the vector c in (3.2) can also depend on n , and the proofs in this section extends naturally to that case.

We now prove a comparison principle for the QVI (3.1), extending a similar result in [28] for finite QVIs and with the identity matrix in place of Q_n . This establishes the uniqueness of solutions to Problem 3.1.

Proposition 3.3. Given Assumption 3.2, suppose $u : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}^A$ (resp. v) is bounded and satisfies for $c > 0$,

$$\min \{F_a(n, x, u(n, x), u_a), u_a(n, x) - \mathcal{M}u(n, x)\} \leq 0 \quad (\text{resp. } \geq 0).$$

Then we have $u \leq v$.

Proof. We write $u_{a,x}^n$ for $u_a(n, x)$, and let $M = \sup_{n,x,a} (u_{a,x}^n - v_{a,x}^n)$. Let (n_k, x_k, a_k) be a maximising sequence, with $M_k = u_{a_k, x_k}^{n_k} - v_{a_k, x_k}^{n_k}$. Suppose that along the whole sequence, $u_{a_k} - \mathcal{M}u \leq 0$. As v is a supersolution, we also have that $v_{a_k} - \mathcal{M}v \geq 0$. Therefore,

$$\begin{aligned} u_{a_k, x_k}^{n_k} - v_{a_k, x_k}^{n_k} &\leq \left(Q_{n_k} \overline{u^1 - c} \right)_{x_k} - \left(Q_{n_k} \overline{v^1 - c} \right)_{x_k} \\ &\leq \left(Q_{n_k} \overline{u^1 - v^1} \right)_{x_k} \\ &\leq \left(\overline{u^1 - v^1} \right)_{x^*} \quad (\text{for some } x^* \in \mathcal{X}). \end{aligned}$$

As the RHS is independent of k , taking the supremum on the LHS leads to

$$M \leq \left(\overline{u^1 - v^1} \right)_{x^*} \leq \max_{a \in A} (u_{a,x^*}^1 - v_{a,x^*}^1).$$

As A is finite, the supremum for M is achieved at some $(1, b_1, y_1)$, $b_1 \in A$, $y_1 \in \mathcal{X}$. Applying the same argument as above, we have

$$u_{b_1, y_1}^1 - v_{b_1, y_1}^1 \leq -c \leq u_{b_2, y_2}^1 - v_{b_2, y_2}^1,$$

for some $(b_2, y_2) \neq (b_1, y_1)$, since by assumption $c \geq 0$. Repeating the same argument leads to successive different indices (b_i, y_i) such that

$$u_{b_1, y_1}^1 - u_{b_n, y_n}^1 \leq \sum_{j=1}^{n-1} (u_{b_j, y_j}^1 - u_{b_{j+1}, y_{j+1}}^1) \leq -(n-1)c < 0. \quad (3.3)$$

But since both \mathcal{X} and A are finite, by the pigeonhole principle, we must eventually have some n such that $(b_1, y_1) = (b_n, y_n)$, which contradicts (3.3). Therefore, we must have some (n_k, x_k, a_k) such that

$$F_{a_k}(n_k, x_k, u(n_k, x_k), u_{a_k}) - F_{a_k}(n_k, x_k, v(n_k, x_k), v_{a_k}) \leq 0. \quad (3.4)$$

Now by construction, $u_{a_k} \leq v_{a_k} + M$. Hence, by taking $r = u(n_k, x_k)$, $s = v(n_k, x_k) + M_k$ and $\theta := \max_{a \in A} (r_a - s_a) = 0$ in the monotonicity condition in Assumption 3.2, we have

$$F_{a_k}(n_k, x_k, v(n_k, x_k) + M, v_{a_k} + M) \leq F_{a_k}(n_k, x_k, u(n_k, x_k), u_{a_k}).$$

Next, as $r \mapsto F_a(n, x, r, \phi)$ is uniformly continuous, it admits a modulus of continuity, which we denote by $\omega(|r - s|)$, so that we have

$$F_{a_k}(n_k, x_k, v(n_k, x_k) + M, v_{a_k} + M) - \omega(M - M_k) \leq F_{a_k}(n_k, x_k, v(n_k, x_k) + M_k, v_{a_k} + M).$$

Combining the above with (3.4), we arrive at

$$\beta M - \omega(M - M_k) \leq F_{a_k}(n_k, x_k, u(n_k, x_k), u_{a_k}) - F_{a_k}(n_k, x_k, v(n_k, x_k), v_{a_k}) \leq 0.$$

Then, as $k \rightarrow \infty$, we arrive at a contradiction, since the modulus of continuity $\omega \rightarrow 0$ and $\beta M > 0$ by assumption. \square

Corollary 3.4. *The solutions to the QVI for the OCM infinite horizon problem (2.11) are unique.*

Proof. Clearly (2.11) can be cast in the form of (3.1), given by

$$\begin{aligned} F_a(n, x, u, \phi) &= u - \gamma \phi(n+1, x) - (P_a^n r_a)_x, \\ Q_n &= \gamma P_a^n, \\ c_{a,x} &= c_{\text{obs}} - \frac{1}{\gamma} r_{a,x}. \end{aligned}$$

Moreover, F_a satisfies Assumption 3.2 with $\beta = 1 - \gamma$. Since a solution to (3.1) is both a subsolution and supersolution, it follows from Proposition 3.3 that any solution must be unique. \square

As alluded to in Section 2, in order to compute the solutions to the QVI (3.1), we have to resort to computing a finite QVI, truncated at some large time. Here we shall also show the comparison principle for the truncated QVI. Let N be the truncated time, $|\mathcal{X}| = L$, $|A| = d$. Then a function $u : \{1, \dots, N\} \times \mathcal{X} \rightarrow \mathbb{R}^A$ can be considered as an element of $\mathbb{R}^{N \times L \times d}$. Therefore, given F_a in Problem 3.1, consider its restriction to $\{1, \dots, N\} \times \mathcal{X} \times \mathbb{R} \times \mathbb{R}^{N \times L \times d}$, and write

$$F_a(n, x, u(n, x), u_a) =: \tilde{F}_a(u)_x^n,$$

i.e., we consider $\{F_a\}_{a \in A}$ as functions from $\mathbb{R}^{N \times L \times d}$ to $\mathbb{R}^{N \times L}$. The monotonicity condition can then be reduced to the form below.

Lemma 3.5. $\tilde{F}_a : \mathbb{R}^{N \times L \times d} \rightarrow \mathbb{R}^{N \times L}$ is a continuous function that satisfies the following property: there exists a constant $\beta > 0$ such that for any $u, v \in \mathbb{R}^{N \times L \times d}$ with $u_{\bar{a}, \bar{l}}^{\bar{n}} - v_{\bar{a}, \bar{l}}^{\bar{n}} = \max_{n, a, l} (u_{a, l}^n - v_{a, l}^n) \geq 0$, we have

$$\tilde{F}_a(u)_i^{\bar{n}} - \tilde{F}_a(v)_i^{\bar{n}} \geq \beta (u_{\bar{a}, \bar{l}}^{\bar{n}} - v_{\bar{a}, \bar{l}}^{\bar{n}}). \quad (3.5)$$

Proof. For the first part, let $u, v \in \mathbb{R}^{N \times L \times d}$ with $\theta := u_{\bar{a}, \bar{l}}^{\bar{n}} - v_{\bar{a}, \bar{l}}^{\bar{n}} = \max_{n, a, l} (u_{a, l}^n - v_{a, l}^n) \geq 0$. Then by definition of θ , $\tilde{u} := u - \theta \leq v$. Applying Assumption 3.2 with $r = u(\bar{n}, \bar{l})$, $s = v(\bar{n}, \bar{l})$ and functions \tilde{u}, v lead to

$$F_{\bar{a}}(\bar{n}, \bar{l}, u(\bar{n}, \bar{l}), \tilde{u}_{\bar{a}} + \theta) - F_{\bar{a}}(\bar{n}, \bar{l}, v(\bar{n}, \bar{l}), v_{\bar{a}}) \geq \beta \theta. \quad (3.6)$$

Noting that $\tilde{u}_{\bar{a}} + \theta = u_{\bar{a}}$, it then follows from (3.6) that $\tilde{F}_a(u)_x^n = F_a(n, x, u(n, x), u_a)$ satisfies (3.5) as required. \square

Thus, the finite truncated QVI can be summarised by the problem below.

Problem 3.6. Find $u = (u_1, \dots, u_d) \in \mathbb{R}^{N \times L \times d}$ such that

$$\min \left\{ \tilde{F}_a(u), u_a - \mathcal{M}u \right\} = 0, \quad a \in \{1, \dots, d\} =: A, \quad (3.7)$$

with time boundary condition

$$u_a(N, \cdot) = \mathcal{M}u(N, \cdot),$$

where $\tilde{F}_a : \mathbb{R}^{N \times L \times d} \rightarrow \mathbb{R}^{N \times L}$ satisfies (3.5), and $\mathcal{M} : \mathbb{R}^{N \times L \times d} \rightarrow \mathbb{R}^{N \times L}$ is defined by

$$(\mathcal{M}u)_i^n = \left(Q_n \overline{u^1 - c} \right)_i, \quad \left(\overline{u^1 - c} \right)_l = \max_{a \in A} (u_{a, l}^1 - c_{a, l}),$$

for a given vector $c \in \mathbb{R}^{L \times d}$ and $\{Q_n\} \subset \mathbb{R}^{L \times L}$ is a sequence of substochastic matrices.

Remark 3.7. For the case with parameter uncertainty, if the measures $\rho_n(d\theta)$ can be parametrised by a finite number of parameters w , this can also be considered as part of the spatial domain. In this case $L = |\mathcal{X}| \cdot |w|$.

When the domain is finite as above, and in the case where the Q_n 's are the identity matrix, then (3.7) reduces to a QVI with interconnected obstacles, see [28] for a more detailed analysis for such classes of QVIs. Naturally, the truncated QVI also satisfies a comparison principle:

Proposition 3.8. Suppose $u = (u_a)_{a \in A}$ (resp. $v = (v_a)_{a \in A}$) satisfies for $c > 0$:

$$\min \left\{ \tilde{F}_a(u), u_a - \mathcal{M}u \right\} \leq 0 \quad (\text{resp. } \geq 0), \quad a \in A;$$

then $u \leq v$.

Proof. Let $M := u_{\bar{a},\bar{l}}^{\bar{n}} - v_{\bar{a},\bar{l}}^{\bar{n}} = \max_{n,a,l}(u_{a,l}^n - v_{a,l}^n)$. The proof follows directly from the arguments in Proposition 3.3. Assuming that the subsolution u satisfies $u_{\bar{a}} \leq \mathcal{M}u$ leads to

$$u_{\bar{a},\bar{l}}^{\bar{n}} - v_{\bar{a},\bar{l}}^{\bar{n}} \leq -c \leq u_{a^*,l^*}^1 - v_{a^*,l^*}^1 \quad \text{for some } a^* \text{ and } l^*,$$

so that the maximum is also achieved at $(1, a^*, l^*)$. Then the same contradiction argument leads to $\tilde{F}_{\bar{a}}(u)_{\bar{l}}^{\bar{n}} \leq 0$, but then since v is a supersolution and $M > 0$, we have by the monotonicity property

$$\beta(u_{\bar{a},\bar{l}}^{\bar{n}} - v_{\bar{a},\bar{l}}^{\bar{n}}) \leq \tilde{F}_{\bar{a}}(u)_{\bar{l}}^{\bar{n}} - \tilde{F}_{\bar{a}}(v)_{\bar{l}}^{\bar{n}} \leq 0$$

which is again a contradiction. Hence we must have $M \leq 0$ as required. \square

We now present a penalty approximation to the QVI (3.7). The motivations behind this are threefold. Firstly, the penalty method holds a crucial advantage over a more traditional policy iteration approach, as using the latter can lead to numerical instabilities or indeed ill-defined iterates, due to the lack of a guarantee of the invertibility of the matrix arising from the intervention operator \mathcal{M} . A simple example of such an instance is given in [29, Section 6], with a more in depth explanation given in [3]. Secondly, for QVIs that arise from the discretisation of diffusions on continuous state spaces, the number of policy iterations is typically large for fine meshes, as exposed in [27], while semismooth Newton iterations for the penalised system are robust under mesh refinement. Lastly, penalisation provides naturally a constructive proof of existence to the solutions of the QVI (3.7). Consider the following penalised problem.

Problem 3.9. Let $\rho \geq 0$ be the penalty parameter. Find $u^\rho = (u_a^\rho)_{a \in A} \in \mathbb{R}^{N \times L \times d}$ such that

$$G_a^\rho(v) := \tilde{F}_a(u^\rho) - \rho \pi(\mathcal{M}u^\rho - u_a^\rho) = 0, \quad a \in A, \quad (3.8)$$

where the penalisation function $\pi : \mathbb{R} \rightarrow \mathbb{R}$ is continuous, non-decreasing with $\pi|_{(-\infty, 0]} = 0$ and $\pi|_{(0, \infty)} > 0$, and is applied elementwise.

Thus in the penalised problem, a penalty ρ is applied whenever the condition $u_a^\rho - \mathcal{M}u^\rho \geq 0$ is violated. As $\rho \uparrow \infty$, the penalised solution should then converge to the solution of the discrete QVI (3.7). We first show below that for each fixed ρ , (3.8) satisfies a comparison principle. This implies uniqueness for Problem 3.9. The argument follows similarly to the approach in [28] and Proposition 3.8.

Proposition 3.10. For any penalty parameter $\rho \geq 0$, if $u^\rho = (u_a^\rho)_{a \in A}$ (resp., $v^\rho = (v_a^\rho)_{a \in A}$) satisfies

$$\tilde{F}_a(u^\rho) - \rho \pi(\mathcal{M}u^\rho - u_a^\rho) \leq 0 \quad (\text{resp., } \geq 0),$$

then $u^\rho \leq v^\rho$.

Proof. Let $M := u_{\bar{a},\bar{l}}^{\rho,\bar{n}} - v_{\bar{a},\bar{l}}^{\rho,\bar{n}} = \max_{n,a,l}(u_{a,l}^{\rho,n} - v_{a,l}^{\rho,n})$. Suppose for a contradiction that $M > 0$. From the previous proposition, we have that

$$\left(Q_{\bar{n}} \overline{u^{\rho,1} - c} \right)_{\bar{l}} - \left(Q_{\bar{n}} \overline{v^{\rho,1} - c} \right)_{\bar{l}} \leq u_{\bar{a},\bar{l}}^{\rho,\bar{n}} - v_{\bar{a},\bar{l}}^{\rho,\bar{n}}.$$

As π is non-decreasing,

$$\pi \left(\left(Q_{\bar{n}} \overline{u^{\rho,1} - c} \right)_{\bar{l}} - u_{\bar{a},\bar{l}}^{\rho,\bar{n}} \right) \leq \pi \left(\left(Q_{\bar{n}} \overline{v^{\rho,1} - c} \right)_{\bar{l}} - v_{\bar{a},\bar{l}}^{\rho,\bar{n}} \right). \quad (3.9)$$

As u^ρ and v^ρ are respectively sub and super solutions of (3.8), we have

$$\begin{aligned} & \tilde{F}_{\bar{a}}(u^\rho)_{\bar{l}}^{\bar{n}} - \rho \pi \left(\left(\mathcal{M}u^\rho \right)_{\bar{l}}^{\bar{n}} - u_{\bar{a},\bar{l}}^{\rho,\bar{n}} \right) - \left(\tilde{F}_{\bar{a}}(v^\rho)_{\bar{l}}^{\bar{n}} - \rho \pi \left(\left(\mathcal{M}v^\rho \right)_{\bar{l}}^{\bar{n}} - v_{\bar{a},\bar{l}}^{\rho,\bar{n}} \right) \right) \\ & \leq \tilde{F}_{\bar{a}}(u^\rho)_{\bar{l}}^{\bar{n}} - \tilde{F}_{\bar{a}}(v^\rho)_{\bar{l}}^{\bar{n}} - \rho \left(\pi \left(\left(Q_{\bar{n}} \overline{u^{\rho,1} - c} \right)_{\bar{l}} - u_{\bar{a},\bar{l}}^{\rho,\bar{n}} \right) - \pi \left(\left(Q_{\bar{n}} \overline{v^{\rho,1} - c} \right)_{\bar{l}} - v_{\bar{a},\bar{l}}^{\rho,\bar{n}} \right) \right) \\ & \leq 0 \end{aligned}$$

Hence by (3.9), $\tilde{F}_{\bar{a}}(u^\rho)_{\bar{l}}^{\bar{n}} - \tilde{F}_{\bar{a}}(v^\rho)_{\bar{l}}^{\bar{n}} \leq 0$. The monotonicity assumption of F then leads to a contradiction, so that $M \leq 0$ as required. \square

Next, we show that the penalised solutions are uniformly bounded above, which is crucial for the convergence towards the unpenalised problem.

Lemma 3.11. *For $\rho \geq 0$ and $c \geq 0$, let u^ρ be the solution to (3.8). Then $\|u^\rho\| \leq \|F(0)\|/\beta$.*

Proof. This is a direct adaptation of [28, Lemma 2.3], where the same estimate is shown for discrete switching problems, relying on the monotonicity condition and the non-negativity of the penalty function. \square

We are now ready to show the well-posedness of the QVI (3.1), which arises from the infinite horizon OCM (2.11). In order to approximate the truncated problem, we require the following additional assumption.

Assumption 3.12. For each $(n, x) \in \mathbb{N} \times \mathcal{X}$, there is a finite set $\Omega_{n,x}$ such that for every $r \in \mathbb{R}^A$, $F_a(n, x, r, \phi)$ only depends on the values of ϕ on $\Omega_{n,x}$, i.e.

$$F_a(n, x, r, \phi) = f_a(n, x, r, \phi|_{\Omega_{n,x}})$$

for some function f_a .

Theorem 3.13. *Given Assumptions 3.2 and 3.12, the QVI (3.1) is well-posed. In particular, let $u^{\rho,N}$ denote the solution to the penalised problem (3.8) with penalty parameter ρ and truncation time N . Then, for fixed $c \geq 0$, $u^N = \lim_{\rho \rightarrow \infty} u^{\rho,N}$ exists and $u = \lim_{N \rightarrow \infty} u^N$ exists (here consider u^N as functions from $\mathbb{N} \times \mathcal{X}$ to \mathbb{R}^A by trivially defining $u_a^N(n, x) = 0$ for $n > N$). In both cases the convergence is monotone from below. Moreover, if $c > 0$ then u^N and u solves (3.7) and (3.1) respectively.*

Proof. The existence of penalised solutions $u^{\rho,N}$ to (3.8) can be established via a mollification argument, applying [28, Theorem 2.5] to $G^{\rho,N}$. Uniqueness of the penalised equation is given by Proposition 3.10. The proof of the convergence of $u^{\rho,N} \uparrow u^N$, where u^N solves (3.7) mirrors the argument in [28, Theorem 2.6]: by Proposition 3.10, for each N , $u^{\rho,N}$ is monotone in ρ , and together with the bound Lemma 3.11, this gives the desired convergence. The continuity of the operators in the QVI is then used to show that u^N solves (3.7). We shall defer the reader to the relevant theorems cited above for the precise details.

For the convergence of u^N to some function u , we see also that u^N is monotone in N : given a solution u^{N+1} of (3.7), we have $u^{N+1}(N, \cdot) \geq \mathcal{M}u^{N+1}(N, \cdot)$, so that $u^{N+1} \geq u^N$ by Proposition 3.8. By Lemma 3.11, u^N is also bounded above, independent of N . Hence $u^N \uparrow u$.

We now show that u solves the QVI (3.1). As $F(n, x, r, \phi)$ is only assumed to be continuous in r in the untruncated case, we cannot directly appeal to continuity as before. We proceed as follows. Fix $n \in \mathbb{N}$, $x \in \mathcal{X}$, $a \in A$ and $\varepsilon > 0$. As \mathcal{M} is continuous, we have $u_a - \mathcal{M}u \geq 0$. By Assumption 3.12 and the pointwise convergence of u^N , we can choose a sufficiently large N' such that

$$\begin{aligned} (u_a - u_a^{N'})|_{\Omega_{n,x}} &\leq \varepsilon, \\ \varepsilon' = u_a(n, x) - u_a^{N'}(n, x) &\leq \varepsilon. \end{aligned}$$

Moreover, we can find $\tilde{N} \geq N'$ such that for all $b \neq a \in A$,

$$u_b(n, x) - u_b^{N'}(n, x) \leq \varepsilon',$$

so that we can define \tilde{u} by $\tilde{u}_a = u_a^{N'}(n, x)$ and $\tilde{u}_b = u_b^{\tilde{N}}(n, x)$ for $b \neq a$. Therefore,

$$u_a(n, x) - \tilde{u}_a = \max_{b \in A} (u_b(n, x) - \tilde{u}_b) = \varepsilon'.$$

Then, we apply the monotonicity condition in Assumption 3.2, with $r = u(n, x) + (\varepsilon - \varepsilon')$, $s = \tilde{u}$, noting also that the functions $u_a - \varepsilon \leq u_a^{\tilde{N}}$ on $\Omega_{n,x}$. This gives

$$F_a(n, x, u(n, x) + \varepsilon - \varepsilon', u_a) - F_a(n, x, \tilde{u}, u_a^{\tilde{N}}) \geq \beta\varepsilon.$$

By the uniform continuity of $r \mapsto F_a(n, x, r, \phi)$, we have

$$\begin{aligned} F_a(n, x, u(n, x) + \varepsilon - \varepsilon', u_a) &\geq F_a(n, x, u^{\tilde{N}}(n, x), u_a^{\tilde{N}}) - \omega(|u_a^{\tilde{N}}(n, x) - u_a^{N'}(n, x)|) + \beta\varepsilon \\ &\geq -\omega(|u_a^{\tilde{N}}(n, x) - u_a^{N'}(n, x)|) + \beta\varepsilon, \end{aligned}$$

where the inequality follows from the fact that $u^{\tilde{N}}$ satisfies the QVI (3.7), and ω is the modulus of continuity. Finally, by the convergence of $u^{\tilde{N}}$ to u , and the uniform continuity of $r \mapsto F_a(n, x, r, \phi)$ again, we have $F_a(n, x, u(n, x), u_a) \geq 0$.

We now have $F_a(n, x, u(n, x), u_a) \geq 0$ and $u_a - \mathcal{M}u \geq 0$. Suppose $u_a - \mathcal{M}u > 0$. We show that $F_a(n, x, u(n, x), u_a) = 0$ to conclude. By the uniform continuity of $r \mapsto F_a(n, x, r, \phi)$, for any $\varepsilon > 0$ there exists $N' \geq 0$ such that for all $N \geq N'$,

$$\begin{aligned} 0 \leq F_a(n, x, u(n, x), u_a) &< F_a(n, x, u^N(n, x), u_a) + \varepsilon \\ &\leq F_a(n, x, u^N(n, x), u_a^N) + \varepsilon, \end{aligned}$$

where the second inequality follows from the monotonicity condition. By the continuity of \mathcal{M} , $u_a^N - \mathcal{M}u^N > 0$ for sufficiently large N , and therefore $F_a(n, x, u^N(n, x), u_a^N) = 0$. But as ε is arbitrary, we have $F_a(n, x, u(n, x), u_a) = 0$ as required. \square

Thus, we have a straightforward computation scheme to solve for the OCM. We first set up the discrete QVIs arising from the problem and choose a suitable time for truncation, then the equations are approximated by the penalised problem. The solution of the penalised problem can then be in turn approximated iteratively with semismooth Newton methods [36]. Formally speaking, starting with an initialisation $v^{(0)}$ to the penalised problem

$$G^\rho(v) = \tilde{F}(v) - \rho \pi(\mathcal{M}v - v) = 0,$$

we obtain the next iterate by solving for

$$v^{(k+1)} = v^{(k)} - \mathcal{L}^\rho(v^{(k)})^{-1} G^\rho(v^{(k)}),$$

where \mathcal{L}^ρ denotes the generalised derivative of the function G^ρ . This shall be the numerical scheme that we adopt for our numerical experiments in the next section.

4 Numerical experiments

In this section, we apply our observation cost framework to three numerical experiments. Sections 4.1 and 4.2 analyse two infinite horizon problems. For these examples, we examine the numerical performance of the penalty method and Newton iterations, as well as the effects of the observation cost on the qualitative behavior of the solutions. For the penalised equations, we will employ the penalty function $\pi(x) = x^+$ as in [28]. Section 5.1 considers the parameter uncertainty formulation over a finite horizon. The solutions are obtained through backwards recursion from the terminal conditions. We examine the impact that the extra parameter uncertainty has on the optimal trajectories.

4.1 Random walk with drift

Consider an integer-valued random walk whose drift depends on the action space $A = \{+1, -1\}$. The probability of each step is parametrised by θ . Specifically, for any $x \in \mathcal{X} = \mathbb{N}$,

$$\begin{aligned} p(x+1 \mid x, +1) &= \theta, & p(x-1 \mid x, +1) &= 1-\theta; \\ p(x+1 \mid x, -1) &= 1-\theta, & p(x-1 \mid x, -1) &= \theta. \end{aligned} \quad (4.1)$$

We also adopt the following reward function:

$$r(x, a) = \frac{1}{|x| + 1}.$$

The mass of this reward function r is concentrated around the origin, so naturally, the optimal action is one that reverts the process back towards the origin.

For this example, we consider the infinite horizon problem. Recall that the discrete QVI (2.11) reads: for all $n \geq 0$, $x \in \mathcal{X}$, and $a \in A$,

$$\min \left\{ v_{a,x}^n - \gamma v_{a,x}^{n+1} - \left(P_a^n r_a \right)_x, v_{a,x}^n - \left(P_a^n \gamma v^1 + r \right)_x + c_{\text{obs}} \right\} = 0, \quad (4.2)$$

Note that there exists a path from x to y over m units of time if and only if $m \geq |y - x|$ and $m \equiv y \pmod{2}$. If \mathcal{S}_m^x denotes the set of states that can be reached from x after m units of time, then for a constant action, the n -step transition probabilities are given by

$$\begin{aligned} p^{(n)}(x' \mid x, +1) &= \begin{cases} \binom{n}{k} \theta^k (1-\theta)^{n-k} & , x' \in \mathcal{S}_n^x; \\ 0 & , x' \notin \mathcal{S}_n^x, \end{cases} \\ p^{(n)}(x' \mid x, -1) &= \begin{cases} \binom{n}{k} \theta^{n-k} (1-\theta)^k & , x' \in \mathcal{S}_n^x; \\ 0 & , x' \notin \mathcal{S}_n^x, \end{cases} \end{aligned}$$

where $k = (n + x' - x)/2$. For this problem, we shall also incorporate a switching cost g whenever the drift changes (see Remark 2.6). Hence, in full, the QVI reads:

$$\begin{aligned} \min \left\{ v_{+1,x}^n - \gamma v_{+1,x}^{n+1} - \sum_{x' \in \mathcal{S}_n^x} \frac{1}{|x'| + 1} \binom{n}{k} \theta^k (1-\theta)^{n-k}, \right. \\ \left. v_{+1,x}^n - \gamma \sum_{x' \in \mathcal{S}_n^x} \binom{n}{k} \theta^k (1-\theta)^{n-k} \left(\frac{1}{|x'| + 1} + \max\{v_{+1,x'}^1, v_{-1,x'}^1 - g\} \right) + c_{\text{obs}} \right\} = 0, \\ \min \left\{ v_{-1,x}^n - \gamma v_{-1,x}^{n+1} - \sum_{x' \in \mathcal{S}_n^x} \frac{1}{|x'| + 1} \binom{n}{k} \theta^{n-k} (1-\theta)^k, \right. \\ \left. v_{-1,x}^n - \gamma \sum_{x' \in \mathcal{S}_n^x} \binom{n}{k} \theta^{n-k} (1-\theta)^k \left(\frac{1}{|x'| + 1} + \max\{v_{+1,x'}^1 - g, v_{-1,x'}^1\} \right) + c_{\text{obs}} \right\} = 0. \end{aligned} \quad (4.3)$$

To close the system to ensure a unique solution, we enforce the following time and spatial boundary conditions. We impose a reflecting boundary at $x = \pm L$, where L is suitably large. In particular,

$$\begin{aligned} p(L \mid L, +1) &= \theta, & p(L-1 \mid L, +1) &= 1-\theta, \\ p(L \mid L, -1) &= 1-\theta, & p(L-1 \mid L, -1) &= \theta, \\ p(-L \mid -L, +1) &= 1-\theta, & p(-L+1 \mid -L, +1) &= \theta, \\ p(-L \mid -L, -1) &= \theta, & p(-L+1 \mid -L, -1) &= 1-\theta, \end{aligned} \quad (4.4)$$

so that the QVI (4.2) for the states $x = \pm L$ will use the transition probabilities (4.4) instead.

For the time boundary, we enforce an observation at some large time $N > 0$. The terminal condition then reads (for $-L < x < L$):

$$\begin{cases} v_{+1,x}^N - \gamma \sum_{x' \in \mathcal{S}_N^x} \binom{N}{k'} \theta^k (1-\theta)^{N-k'} \left(\frac{1}{|x'|+1} + \max\{v_{+1,x'}^1, v_{-1,x'}^1 - g\} \right) + c_{\text{obs}} = 0, \\ v_{-1,x}^N - \gamma \sum_{x' \in \mathcal{S}_N^x} \binom{N}{k'} \theta^{N-k'} (1-\theta)^{k'} \left(\frac{1}{|x'|+1} + \max\{v_{+1,x'}^1 - g, v_{-1,x'}^1\} \right) + c_{\text{obs}} = 0. \end{cases} \quad (4.5)$$

where $k' = (N + x' - x)/2$. The analogous equations hold for the spatial boundary $x = \pm L$, but with the transition probabilities (4.4). These terminal conditions can be interpreted as the largest possible interval between two observations.

We now proceed to solve the penalised problem for the system (4.3), with boundary conditions (4.4) and (4.5), through the use of semismooth Newton methods. To initialise the iteration, we solve for the uncoupled system

$$\begin{cases} v_{+1,x}^n - \gamma v_{+1,x}^{n+1} - \sum_{x' \in \mathcal{S}_n^x} \frac{1}{|x'|+1} \binom{n}{k} \theta^k (1-\theta)^{n-k} = 0, \\ v_{-1,x}^n - \gamma v_{-1,x}^{n+1} - \sum_{x' \in \mathcal{S}_n^x} \frac{1}{|x'|+1} \binom{n}{k} \theta^{n-k} (1-\theta)^k = 0, \\ 0 \leq n < N, \quad -L < x < L, \end{cases} \quad (4.6)$$

with the spatial boundary transition probabilities (4.4) and uncoupled time boundary condition

$$\begin{cases} v_{+1,x}^N - \gamma \sum_{x' \in \mathcal{S}_N^x} \binom{N}{k'} \theta^{k'} (1-\theta)^{N-k'} \left(\frac{1}{|x'|+1} + v_{+1,x'}^1 \right) + c_{\text{obs}} = 0, \\ v_{-1,x}^N - \gamma \sum_{x' \in \mathcal{S}_N^x} \binom{N}{k'} \theta^{N-k'} (1-\theta)^{k'} \left(\frac{1}{|x'|+1} + v_{-1,x'}^1 \right) + c_{\text{obs}} = 0, \\ -L < x < L. \end{cases}$$

The system (4.6) corresponds to the penalised equation with penalty parameter $\rho = 0$. The uncoupled time boundary condition is equivalent to enforcing an observation but with no switching (i.e., assuming that $\bar{v} = v_a$ in each equation for v_a). The iteration terminates once a relative tolerance threshold of 10^{-8} is reached.

We investigate the numerical performance of our described methods for the case $\theta = 0.75$, $\gamma = 0.99$, $L = 50$ and $N = 500$, across different cost parameters c_{obs} . Computations are performed using MATLAB R2019b. The numerical solutions are shown in Table 4.1, for the case of zero switching cost. Row (a) shows that the number of Newton iterations required to reach the tolerance threshold is independent from the size of the penalty parameter ρ . Fewer iterations are required for more extreme values of c_{obs} , but the overall number of iterations remains low across different observation costs. Row (b) shows the increments $\|v^\rho - v^{2\rho}\|_\infty$. The values suggests a first-order convergence of the penalisation error with respect to the penalty parameter ρ , which is in line with the analogous theoretical results in [28, Theorem 3.9, 4.2]. Similar experiments with non-zero switching costs also demonstrate a first-order convergence.

We now examine the qualitative behaviour of the solution. It is clear that if the chain is observed to be at a positive state, then the control should be switched to $a = -1$ for a negative drift and vice versa. Figure 4.2 shows the optimal waiting time for the next inspection, across the state space of $[-50, 50]$. To demonstrate the extent to which the waiting time is due to the

| | ρ | 10^3 | 2×10^3 | 4×10^3 | 8×10^3 | 16×10^3 | 32×10^3 |
|------------------------|--------|-----------|-----------------|-----------------|-----------------|------------------|------------------|
| $c_{\text{obs}} = 0$ | (a) | 2 | 2 | 2 | 2 | 2 | 2 |
| | (b) | 0.0063278 | 0.0031650 | 0.0015828 | 0.0007915 | 0.0003957 | 0.0001979 |
| $c_{\text{obs}} = 1/8$ | (a) | 5 | 5 | 5 | 5 | 5 | 5 |
| | (b) | 0.0048459 | 0.0024240 | 0.0012123 | 0.0006062 | 0.0003031 | 0.0001516 |
| $c_{\text{obs}} = 1/4$ | (a) | 6 | 6 | 6 | 6 | 6 | 6 |
| | (b) | 0.0033831 | 0.0016926 | 0.0008466 | 0.0004234 | 0.0002117 | 0.0001059 |
| $c_{\text{obs}} = 1/2$ | (a) | 6 | 6 | 6 | 6 | 6 | 6 |
| | (b) | 0.0015376 | 0.0007691 | 0.0003846 | 0.0001923 | 0.0000962 | 0.0000481 |
| $c_{\text{obs}} = 1$ | (a) | 7 | 7 | 7 | 7 | 7 | 7 |
| | (b) | 0.0006210 | 0.0003105 | 0.0001553 | 0.0000776 | 0.0000388 | 0.0000194 |
| $c_{\text{obs}} = 2$ | (a) | 8 | 8 | 8 | 8 | 8 | 8 |
| | (b) | 0.0002077 | 0.0001038 | 0.0000519 | 0.0000260 | 0.0000130 | 0.0000065 |
| $c_{\text{obs}} = 4$ | (a) | 7 | 7 | 7 | 7 | 7 | 7 |
| | (b) | 0.0000852 | 0.0000426 | 0.0000213 | 0.0000157 | 0.0000053 | 0.0000027 |
| $c_{\text{obs}} = 6$ | (a) | 6 | 6 | 6 | 6 | 6 | 6 |
| | (b) | 0.0000307 | 0.0000154 | 0.0000077 | 0.0000038 | 0.0000019 | 0.0000010 |

Table 4.1: Numerical results for the random walk with drift problem with switching cost $g = 0$. Line (a): number of Newton iterations to reach the relative tolerance threshold of $1e^{-8}$. Line (b): the increment sizes $\|v^\rho - v^{2\rho}\|_\infty$.

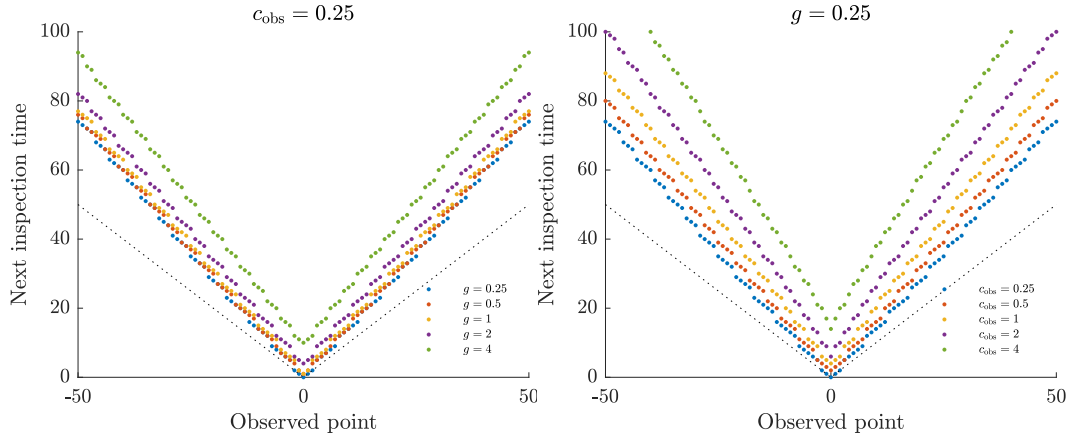


Figure 4.2: The optimal waiting time before the next inspection, across the states $[-50, 50]$. Left: fixed observation cost $c_{\text{obs}} = 0.25$ against variable switching costs. Right: fixed switching cost $g = 0.25$ against variable observation costs.

observation cost versus the switching cost, we vary both parameters and observe the optimal behaviour. The black dotted line represents the base case of zero observation and switching cost, in which case there is no need to observe until it is possible for the walk to cross the origin again. On the left plot, the observation cost c_{obs} is fixed whilst the switching cost g gradually increases, and vice versa on the right. We observe empirically that the observation cost contributes more significantly to the waiting times than the switching cost.

4.2 Extension of an HIV-treatment model

In this subsection, we implement our formulation of the OCM to an HIV-treatment scheduling problem in [34]. There, the authors modelled the problem with a continuous-time MDP with observation costs, but does not include the time elapsed variable in dynamic programming. This can be interpreted as an implicit assumption that the observer is given the state of the underlying process at initialisation. We shall implement a discretised version of their model under our formulation with the time elapsed variable. As alluded to in the introduction, this allows in addition initial conditions that are outdated or sub-optimal relative to the objective. We demonstrate the qualitative difference in the optimal policies when varying the initial conditions, whilst replicating the results in the original paper when the initial conditions coincide. We also examine the numerical performance of the penalty method when applied to the system of QVIs for this larger system, compared to that in Section 4.1.

We now proceed to describe the original problem in [34]. A continuous-time MDP is used to model virus levels of HIV-positive patients over time. With two types of treatment available, the action space is $A = \{0, 1, 2\}$ (where 0 represents no treatment given). Four virus strains are considered: WT denotes the wild type (susceptible to both treatments), R1 and R2 denotes strains that are each resistant to Treatment 1 and Treatment 2 respectively, and HR denotes the strain that is highly resistant to both. The level of each strain is represented by the states ‘none’ (0), ‘low’ (l), ‘medium’ (m), and ‘high’ (h). Therefore, the state space for the Markov chain is $\mathcal{X} = \{0, l, m, h\}^4 \cup \{*\}$, where the asterisk represents patient death. Note in particular that $*$ is an absorbing state. The goal in the original model is to then minimise a cost functional $J : \mathcal{X} \times A \rightarrow \mathbb{R}$ of the form:

$$J(x, \alpha) = \mathbb{E} \left[\sum_{j=0}^{\infty} \left(\int_{\tau_j}^{\tau_{j+1}} e^{-\gamma s} c(X_s, \iota(X_{\tau_j})) ds + e^{-\gamma \tau_{j+1}} c_{\text{obs}} \right) \right],$$

where $\{\tau_j\}_{j=0}^{\infty}$ are the observation times, and the cost function $c : \mathcal{X} \times A \rightarrow \mathbb{R}$ is a linear combination of the productivity loss resulting from each patient’s condition and their received treatment.

To adapt the model above for our formulation, we first discretise the MDP, taking each step to represent one day. We then take the model parameters from the original article [34, Section 3], which provides the transition rate matrices $\{Q_a\}_{a \in A}$ and the cost function $c(x, a)$. The transition matrices $\{P_a\}_{a \in A}$ are then given by $P_a = e^{Q_a}$ (as the time unit in [34] is one day). For illustration purposes, a sparse plot of the transition matrix P_0 is shown in Figure 4.3. As we are considering maximisation problems in this article, we take $r = -c$ for the reward function. We can now formulate our problem in terms of the following QVI:

$$\min \left\{ v_{a,x}^n - \gamma v_{a,x}^{n+1} + \left(e^{nQ_a} c_a \right)_x, v_{a,x}^n - \left(e^{nQ_a} \overline{\gamma v^1 + c} \right)_x + c_{\text{obs}} \right\} = 0.$$

We now follow the same procedure in Section 4.1 to obtain a numerical solution. Note that for this problem, the spatial domain is finite and we also have a natural spatial boundary arising from the absorbing death state $*$, that is, for all $n \geq 0$ and $a \in A$,

$$v_{a,*}^n = \sum_{k=0}^{\infty} l \gamma^k = \frac{l}{1 - \gamma},$$

where l is a constant representing the average GDP loss due to patient death, the value of which is taken from the parameters in [33, 34]. A time boundary is once again enforced at some large time $N > 0$, which can be interpreted as a mandatory observation at time N . Explicitly, this reads

$$v_{a,x}^N - \left(e^{NQ_a} \overline{\gamma v^1 + c} \right)_x + c_{\text{obs}} = 0, \quad x \in \mathcal{X} \setminus \{*\}, \quad a \in A.$$

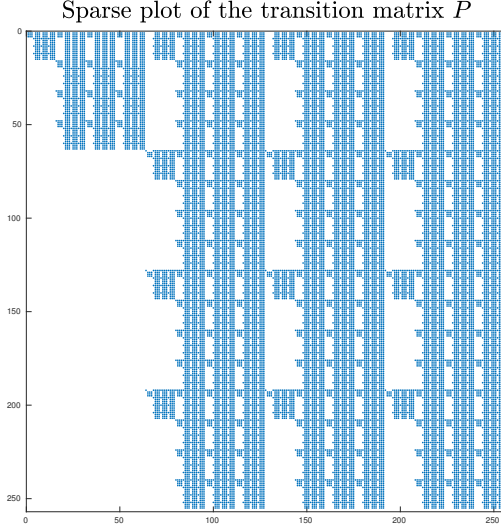


Figure 4.3: *Sparsity pattern of the transition matrix P_0 (the pattern is the same across all control states). The state space is encoded as $\{1, \dots, 256\}$, by considering the state vectors $[WT, R1, R2, HR]$ as a base-4 string in reverse order (for example, $[h, 0, l, l]$ corresponds to 83). The death state $*$ is represented by 256.*

We now solve the associated penalised problem with semismooth Newton methods. As in Section 4.1, we choose the initial guess to be the solution to the penalised problem with $\rho = 0$, with uncoupled time boundary conditions. The iterations terminate once a relative tolerance threshold of 10^{-8} is reached. The numerical experiments are performed on MATLAB R2019b.

Table 4.5 shows the numerical solution for different values of the truncation time N and c_{obs} across different penalty parameters ρ . Row (a) shows that the number of iterations remains constant with respect to ρ , much like the random walk experiment in Section 4.1. For this problem, the number of Newton iterations required to reach the $1\text{e-}8$ threshold is higher at approximately 20 iterations. However, we find that convergence to the optimal policy is typically achieved within the first 2 iterations. This is depicted in Figure 4.4, which graphs the first two iterates as well as the final solution for the value function. Row (b) in Table 4.5 shows the successive increments $\|v^\rho - v^{2\rho}\|_\infty$ between doubling penalty values. Reassuringly, for this more complicated system, we still see a clear first-order convergence of the penalisation error with respect to the penalty parameter ρ . Even for small values of ρ , the successive increments were within $O(1)$ (in comparison to the magnitude of the solution which is of $O(10^6)$). This shows that the penalty approximation is very effective for small penalty parameters, and that it works well when extended to the class of QVIs that we introduced in Section 3.

We now analyse the behaviour of the value function when plotted as a function against time. The top-left graph of Figure 4.6 depicts an instance where the patient is under a stable condition. Here the observation region is $[15, N]$. There are limited benefits of frequently paying a high observation cost when it is unlikely that the patient’s condition will deteriorate over a short period of time. On the other hand, the top-right graph has an observation region of $[0, 53]$. The mathematical intuition behind this is that beyond the observation region, the MDP is expected to enter the absorbing state $*$ with high probability, and the negative reward associated with this absorbing state outweighs any potential benefits of paying the observation cost c_{obs} for

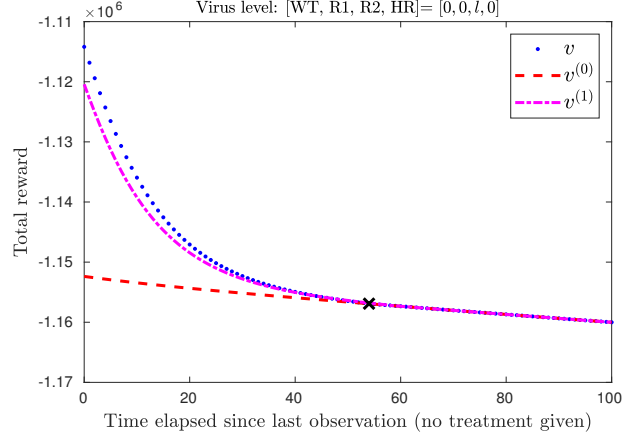


Figure 4.4: Convergence of the Newton iterates towards the solution. The lines show the graphs of $n \mapsto v_{0,4}^n$ for the initial guess $v^{(0)}$, first iterate $v^{(1)}$ and true solution v , where the state $[WT, R1, R2, HR] = [0, 0, l, 0]$ is encoded as 4 in base 4. The cross indicates the boundary between the observation regions.

| | ρ | 10^3 | 2×10^3 | 4×10^3 | 8×10^3 | 16×10^3 | 32×10^3 |
|---------------------------------|--------|--------|-----------------|-----------------|-----------------|------------------|------------------|
| $N = 150, c_{\text{obs}} = 200$ | (a) | 18 | 18 | 18 | 18 | 18 | 18 |
| | (b) | 1.6141 | 0.8071 | 0.4036 | 0.2018 | 0.1009 | 0.0504 |
| $N = 150, c_{\text{obs}} = 400$ | (a) | 21 | 21 | 21 | 21 | 21 | 21 |
| | (b) | 1.5147 | 0.7577 | 0.3790 | 0.1895 | 0.0948 | 0.0474 |
| $N = 150, c_{\text{obs}} = 800$ | (a) | 20 | 20 | 20 | 20 | 20 | 20 |
| | (b) | 1.4087 | 0.7047 | 0.3524 | 0.1762 | 0.0881 | 0.0441 |
| $N = 300, c_{\text{obs}} = 200$ | (a) | 20 | 20 | 20 | 20 | 20 | 20 |
| | (b) | 1.6122 | 0.8061 | 0.4031 | 0.2015 | 0.1008 | 0.0504 |
| $N = 300, c_{\text{obs}} = 400$ | (a) | 19 | 19 | 19 | 19 | 19 | 19 |
| | (b) | 1.5131 | 0.7569 | 0.3785 | 0.1893 | 0.0947 | 0.0473 |
| $N = 300, c_{\text{obs}} = 800$ | (a) | 20 | 20 | 20 | 20 | 20 | 20 |
| | (b) | 1.4102 | 0.7055 | 0.3528 | 0.1764 | 0.0882 | 0.0441 |
| $N = 600, c_{\text{obs}} = 200$ | (a) | 19 | 19 | 19 | 19 | 19 | 19 |
| | (b) | 1.6111 | 0.8056 | 0.4028 | 0.2014 | 0.1007 | 0.0504 |
| $N = 600, c_{\text{obs}} = 400$ | (a) | 17 | 17 | 17 | 17 | 17 | 17 |
| | (b) | 1.5114 | 0.7561 | 0.3781 | 0.1891 | 0.0945 | 0.0473 |
| $N = 600, c_{\text{obs}} = 800$ | (a) | 18 | 18 | 18 | 18 | 18 | 18 |
| | (b) | 1.4065 | 0.7036 | 0.3519 | 0.1760 | 0.0880 | 0.0440 |

Table 4.5: Numerical results for the HIV-treatment problem. Line (a): number of Newton iterations. Line (b): the increments $\|v^\rho - v^{2\rho}\|$.

information. In the original model in [34], one determines the optimal policy based on an immediate observation in hand. Putting this in the context of our formulation, this amounts to fixing an initial condition in the form of $y = (1, x, a) \in \mathcal{Y}$, and looking forward ahead in time to find the first observation time. This overlooks situations such as that occurring in the top-right graph of Figure 4.6: by initialising at the time origin, one immediately ‘loops back’ for an immediate observation, and therefore does not see the effect of the passage of time on the optimal observation policy.

To examine the behaviour around the decision boundaries, we plot the central finite difference terms $(v_{a,x}^{n+1} - v_{a,x}^{n-1})/2\Delta n$ in the bottom row of Figure 4.6, underneath their respective graphs of the value function. If we consider the plots as a discretisation of a continuous value function, we see that there is much bigger variation within the observation region. Critically, there is non-smoothness across the boundary in the bottom-left graph. This suggests that the solution in continuous-time is C^2 in time within each decision region, but only C^1 across the boundary. This is in line with theoretical results on the regularity of viscosity solutions in optimal stopping and switching problems [24, Chapter 5], which is a potential direction for future analysis.

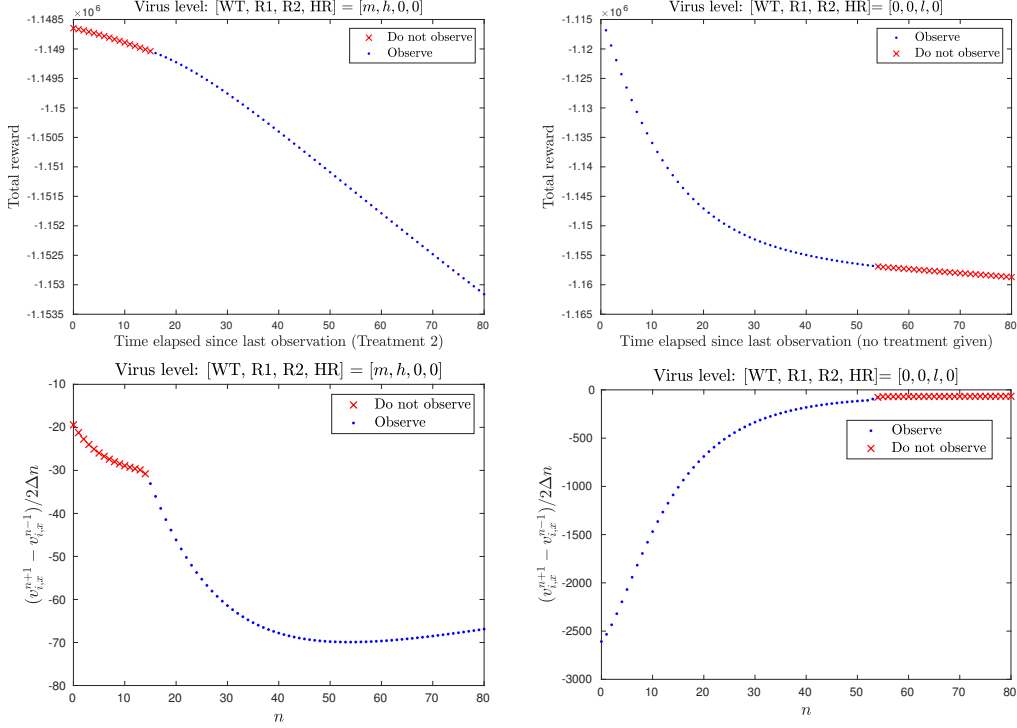


Figure 4.6: The value function exhibits two qualitatively different decay modes depending on the starting states x . Left: a stable condition with the correct treatment. Right: a worse condition with no treatment. The top row shows the mappings $n \mapsto v_{i,x}^n$. The bottom row plots the corresponding central finite difference terms.

5 Observation cost with parameter uncertainty

For this final section, we consider an extension of the OCM with parameter uncertainty in the dynamics of the Markov chain. We shall adopt the approach of Bayesian adaptive control. Suppose that the transition kernel p now depends on an unknown parameter $\theta \in \Theta$, where Θ denotes a finite parameter space. We write $p_\theta(\cdot \mid x, a)$ for a fixed value of θ . To consider a Markov system for the problem, we take $\mathcal{X} \times \Theta$ as our underlying space, with transition kernel

$$p((x', \theta') \mid (x, \theta), a) := \mathbb{1}_{\{\theta=\theta'\}} p_\theta(x' \mid x, a), \quad (x, \theta), (x', \theta') \in \mathcal{X} \times \Theta, \quad a \in A.$$

For each value of $\theta \in \Theta$, we associate a probability measure $p_0^\theta \in \Delta_{\mathcal{X}}$, so that the initial distribution is given by

$$\mathbf{p}_0(x, \theta) = \sum_{\theta \in \Theta} p_0^\theta(x) \rho_0(\theta)$$

for some $\rho_0 \in \Delta_\Theta$. Given the half-step construction as laid out in the beginning of this section, the OCM with parameter uncertainty can be written as a POMDP $\langle \mathcal{X} \times \Theta, \mathcal{X}_\emptyset, \mathcal{A}, \mathbf{p}, \mathbf{q}, \mathbf{q}_0, r_{\text{obs}} \rangle$ over the timescale $\frac{1}{2}\mathbb{N}$, where the domain of the transition kernel \mathbf{p} is extended to $(\mathcal{X} \times \Theta) \times \mathcal{A}$ by defining

$$\mathbf{p}(\cdot \mid x, \theta, i) = \delta_{(x, \theta)}(\cdot), \quad i \in \mathcal{I},$$

and the observation kernel \mathbf{q} is now defined on $\mathcal{A} \times (\mathcal{X} \times \Theta)$, with

$$\begin{aligned} \mathbf{q}(\cdot \mid a, x, \theta) &= \delta_\emptyset(\cdot), \quad a \in A, \\ \mathbf{q}(\cdot \mid i, x, \theta) &= i \cdot \delta_x(\cdot) + (1 - i) \delta_\emptyset(\cdot), \quad i \in \mathcal{I}. \end{aligned}$$

The initial observation kernel \mathbf{q}_0 will be taken as $\mathbf{q}_0(\cdot \mid x, \theta) = \delta_x(\cdot)$ (see Remark 2.3). The set of admissible policies, denoted by Π_{obs}^Θ in this case, can be established analogously as in Definition 2.2. Denote the canonical measure here by $\mathbb{P}_{\mathbf{p}_0}^\pi$, under which θ can be considered as a constant process, i.e. $\theta_{n+1} \equiv \theta_n$ with $\theta_0 \sim \rho_0 \in \Delta_\Theta$, and ρ_0 can be interpreted as a prior estimate for θ .

When considering the belief MDP for this problem, the observable sequence at time t remains as previously,

$$h_t = (\bar{x}_0, \pi_0, \dots, \bar{x}_{t-1/2}, \pi_{t-1/2}, \bar{x}_t) \in \mathcal{H}_t.$$

Then, the belief state $\mathbb{P}_{\mathbf{p}_0}^\pi(x_t, \theta_t \mid h_t)$ can be decomposed as

$$\mathbb{P}_{\mathbf{p}_0}^\pi(x_t, \theta_t \mid h_t) = \sum_{\theta \in \Theta} \mathbb{P}_{\mathbf{p}_0}^\pi(x_t \mid \theta_t, h_t) \mathbb{P}_{\mathbf{p}_0}^\pi(\theta_t \mid h_t). \quad (5.1)$$

For a fixed value of θ , $\mathbb{P}_{\mathbf{p}_0}^\pi(x_t \mid \theta_t, h_t)$ has a finite dimensional characterisation by the Markov property. As in the previous section, we denote this characterisation by $y = (y_t)_t$, which is a tuple given by

$$y_t := \begin{cases} (k, x_k, a_k), & \text{if } k = \arg \max_{n \leq t} \{i_n = 1\} \neq t, \\ (t, x_t, \emptyset), & \text{otherwise.} \end{cases}$$

The second term on the right hand side of (5.1), $\mathbb{P}_{\mathbf{p}_0}^\pi(\theta_t \mid h_t)$ can be interpreted as the posterior distribution of θ at time t . Denote this term by $\rho_t(\theta_t)$. Note that ρ_t can be computed online via the classical Bayes' Theorem,

$$\rho_t(\theta) = \frac{\mathbb{P}_{\mathbf{p}_0}^\pi(\bar{x}_t \mid \theta, h_{t-\frac{1}{2}}, \pi_{t-\frac{1}{2}})}{\sum_{\theta' \in \Theta} \mathbb{P}_{\mathbf{p}_0}^\pi(\bar{x}_t \mid \theta', h_{t-\frac{1}{2}}, \pi_{t-\frac{1}{2}}) \rho_{t-\frac{1}{2}}(\theta')} \rho_{t-\frac{1}{2}}(\theta). \quad (5.2)$$

In the OCM, observations only occur at the half steps, therefore we have in fact $\rho_n \equiv \rho_{n-1/2}$ for $n \in \mathbb{N}$. Thus, the update (5.2) can be reduced to

$$\begin{aligned} \rho_n(\theta) &= \frac{\mathbb{P}_{\mathbf{p}_0}^\pi(\bar{x}_{n-\frac{1}{2}} \mid \theta, y_{n-1}, i_{n-1})}{\sum_{\theta' \in \Theta} \mathbb{P}_{\mathbf{p}_0}^\pi(\bar{x}_{n-\frac{1}{2}} \mid \theta', y_{n-1}, i_{n-1}) \rho_{n-1}(\theta')} \rho_{n-1}(\theta) \\ &=: U(\rho_{n-1}, y_{n-1}, i_{n-1}), \quad n \in \mathbb{N}. \end{aligned} \quad (5.3)$$

The belief state at time t can now be represented by $y_t \in \mathcal{Y}$ and $\rho_t \in \Delta_\Theta$, with its transitions given by the kernel $\mathbf{p}' = (\mathbf{p}'_t)_{t \in \frac{1}{2}\mathbb{N}}$ on $\mathcal{Y} \times \Delta_\Theta$, given $(\mathcal{Y} \times \Delta_\Theta) \times \mathcal{A}$:

$$\begin{aligned} \mathbf{p}'_n(y', \rho' \mid y, \rho, i) &= i \cdot \sum_{\theta \in \Theta} p_\theta^{(n-k)}(\hat{x} \mid x, a) \rho(\theta) \mathbb{1}_{\{y'=(n, \hat{x}, \emptyset), \rho'=U(\rho, y, i)\}} \\ &\quad + (1-i) \mathbb{1}_{\{\hat{y}=y, \rho'=\rho\}}, & y &= (k, x, a), \\ \mathbf{p}'_{n+\frac{1}{2}}(y', \rho' \mid y, \rho, a') &= \mathbb{1}_{\{y'=(k, x, a'), a'=a, \rho'=\rho\}}, & y &= (k, x, a), \\ \mathbf{p}'_{n+\frac{1}{2}}(y', \rho' \mid y, \rho, a') &= \mathbb{1}_{\{\hat{y}=(n, x, a'), \rho'=\rho\}}, & y &= (n, x, \emptyset). \end{aligned}$$

As before, we shall not distinguish between policies for the POMDP and policies for the belief state MDP. For the finite horizon problem, let $y = (k, x, a) \in \mathcal{Y}$, $\rho \in \Delta_\Theta$, and consider

$$J(t, y, \rho, \pi) = \mathbb{E}_{y, \rho}^\pi \left[\sum_{n=t}^{2N} r_{n/2}^\rho \left(y_{n/2}, \pi_{n/2} \right) \right], \quad 0 \leq t \leq 2N, \quad \pi \in \Pi_{\text{obs}}, \quad (5.4)$$

where for $y = (k, x, a)$, $i \in \mathcal{I}$, $a' \in A$, $n \in \mathbb{N}$,

$$\begin{aligned} r_n^\rho(y, i) &= -i \cdot c_{\text{obs}}, \\ r_{n+\frac{1}{2}}^\rho(y, a') &= \sum_{\theta \in \Theta} \sum_{x' \in \mathcal{X}} r(x', a') p_\theta^{(n-k)}(x' \mid x, a) \rho(\theta), \end{aligned}$$

and for $y = (n, x, \emptyset)$,

$$r_{n+\frac{1}{2}}^\rho(y, a') = r(x, a').$$

The value function is

$$v(t, y, \rho) = \sup_{\pi \in \Pi_{\text{obs}}} J(t, y, \rho, \pi). \quad (5.5)$$

As in the previous case without parameter uncertainty, the dynamic programming equation can be reduced to only the integer time steps. We state the optimality equation for the observation cost model under parameter uncertainty below.

Proposition 5.1. *For $y = (k, x, a) \in \mathcal{Y}$ and $\rho \in \Delta_\Theta$, the value function (5.5) satisfies the following equation:*

$$\begin{aligned} v(n, (k, x, a), \rho) &= \max \left\{ v(n+1, (k, x, a), \rho) + \sum_{\substack{\theta \in \Theta \\ x' \in \mathcal{X}}} p_\theta^{(n-k)}(x' \mid x, a) r(x', a) \rho(\theta), \right. \\ &\quad \left. \sum_{\substack{\theta \in \Theta \\ x' \in \mathcal{X}}} p_\theta^{(n-k)}(x' \mid x, a) \rho(\theta) \left[\max_{a' \in A} \left(v(n+1, (n, x', a'), \rho') + r(x', a') \right) \right] - c_{\text{obs}} \right\}, \end{aligned} \quad (5.6)$$

where $\rho' = U(\rho, y, 1)$ as in (5.3).

For the infinite horizon case, a similar stationary argument leads us to the objective function and value function:

$$\begin{aligned} J(y, \rho, \pi) &= \mathbb{E}^\pi \left[\sum_{n=0}^{\infty} \gamma^{\lfloor \frac{n}{2} \rfloor} r^\rho \left(y_{n/2}, \pi_{n/2} \right) \right], \quad y \in \mathcal{Y}, \quad \rho \in \Delta_\Theta, \quad \pi \in \Pi_{\text{obs}}, \\ v(y, \rho) &= \sup_{\pi \in \Pi_{\text{obs}}} J(y, \rho, \pi). \end{aligned} \quad (5.7)$$

Proposition 5.2. For $y = (n, x, a) \in \mathcal{Y}$ and $\rho \in \Delta_\Theta$, the value function (5.7) satisfies the following equation:

$$v((n, x, a), \rho) = \max \left\{ \gamma v((n+1, x, a), \rho) + \sum_{\substack{\theta \in \Theta \\ x' \in \mathcal{X}}} p_\theta^{(n)}(x' | x, a) r(x', a) \rho(\theta), \right. \\ \left. \sum_{\substack{\theta \in \Theta \\ x' \in \mathcal{X}}} p_\theta^{(n)}(x' | x, a) \rho(\theta) \left[\max_{a' \in A} \left(\gamma v((1, x', a'), \rho') + r(x', a') \right) \right] - c_{\text{obs}} \right\}, \quad (5.8)$$

where $\rho' = U(\rho, y, 1)$ as in (5.3).

It is worth noting that both (5.6) and (5.8) are MDPs over the augmented space $\mathcal{Y} \times \Delta_\Theta$. The inclusion of the simplex Δ_Θ makes the MDP non-discrete. For computation, one would have to approximate the solution, either via computing a discrete MDP on a finite grid on $\mathcal{Y} \times \Delta_\Theta$, or via functional approximation methods such as the use of neural networks on larger scale problems. We refer the reader to the textbook [21] and survey paper [31] for a comprehensive review of choosing appropriate approximating grids. Then, given a finite grid $\mathbb{G} = \{s_1, \dots, s_G\}$ on the simplex Δ_Θ , one can define the approximating transition kernels by

$$\mathbf{p}_\mathbb{G}(y', s_i | y, s_j, \pi) = \frac{\mathbf{p}'(y', s_i | y, s_j, a)}{\sum_{i=1}^G \mathbf{p}'(y', s_i | y, s_j, a)}, \quad y, y' \in \mathcal{Y}, \quad s_i, s_j \in \mathbb{G}, \quad \pi \in \mathcal{A},$$

so that one solves the approximating finite MDP on $\mathcal{Y} \times \mathbb{G}$ instead. For the QVIs resulting from the observation cost problems, we can solve the MDPs by penalisation, as detailed in Section 3.

5.1 Random walk with drift with parameter uncertainty

In this subsection, we consider a random walk with drift, as set up in Section 4.1, but with the additional assumption that the true value of the drift parameter θ is unknown to the user. To avoid complications with boundary conditions and infinite domains, we shall only consider the finite horizon problem. Recall that for a fixed value of θ and constant action, the n -step transition probabilities are given by

$$p_\theta^{(n)}(x' | x, +1) = \binom{n}{k} \theta^k (1-\theta)^{n-k}, \quad p_\theta^{(n)}(x' | x, -1) = \binom{n}{k} \theta^{n-k} (1-\theta)^k, \quad x' \in \mathcal{S}_n^x,$$

where \mathcal{S}_n^x is the set of states that can be reached from x after n units of time, and $k = \frac{1}{2}(n + x' - x)$. We choose the prior from a family of beta distributions to obtain conjugacy in the parameter distributions. The posterior can be updated as follows. Suppose the prior $\rho_0 \sim \text{Beta}(\alpha, \beta)$ and the next observation occurs at time n at a state $x' \in \mathcal{X}$. Then a standard calculation shows that

$$\int_{\Theta} p_\theta^{(n)}(x' | x, +1) \rho_0(d\theta) = g(k | n, \alpha, \beta), \quad \int_{\Theta} p_\theta^{(n)}(x' | x, -1) \rho_0(d\theta) = g(n-k | n, \alpha, \beta),$$

where

$$g(k | n, \alpha, \beta) = \binom{n}{k} \frac{B(k + \alpha, n - k + \beta)}{B(\alpha, \beta)}, \quad k = \frac{1}{2}(n + x' - x),$$

and g is the probability mass function of the Beta-binomial distribution, $B(\alpha, \beta)$ is the Beta function. The posterior distribution is then given by

$$\rho_n \sim \begin{cases} \text{Beta}(\alpha + k, \beta + n - k), & a_0 = +1, \\ \text{Beta}(\alpha + n - k, \beta + k), & a_0 = -1. \end{cases}$$

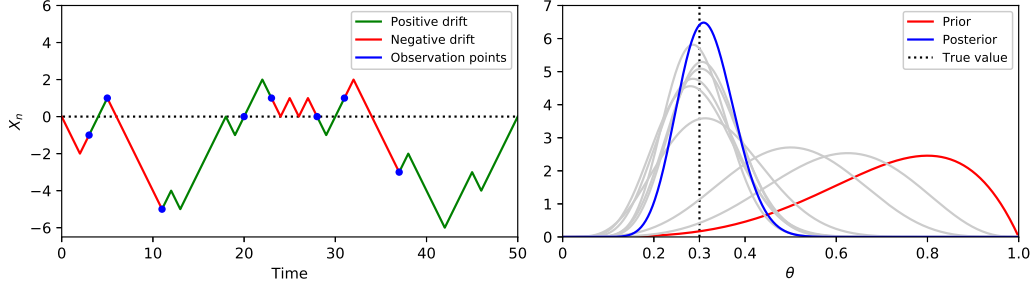


Figure 5.1: *Left: sample realisation of the controlled random walk along the optimal trajectory. Right: prior and posterior distribution of θ ; the grey lines indicate ‘intermediate posteriors’ obtained from earlier observations.*

The distributions $\{\rho_n\}$ can then be described by a finite number of values over a finite time horizon. Hence we write $v(n, (k, x, a), (\alpha, \beta))$ for $v(n, y, \rho)$. Let us first consider the same reward function $r(x, a) = r(x) = \frac{1}{|x|+1}$ as in the previous section. We can write the QVI (5.6) as

$$\begin{aligned} & v(n, (k, x, a), (\alpha, \beta)) \\ &= \max \left\{ v(n+1, (k, x, a), (\alpha, \beta)) + \sum_{x' \in \mathcal{X}} g'(x', x | n-k, \alpha, \beta) r(x'), \right. \\ & \quad \left. \sum_{x' \in \mathcal{X}} g'(x', x | n-k, \alpha, \beta) \left[\max_{a' \in A} \left(v(n+1, (n, x', a'), (\alpha', \beta')) + r(x') \right) \right] - c_{\text{obs}} \right\}, \end{aligned}$$

where we define $g'(x', x | n, \alpha, \beta) := g((n-x'+x)/2 | k, \alpha, \beta)$. The terminal conditions are

$$v(N, (N-k, x, a), (\alpha, \beta)) = \sum_{x' \in \mathcal{S}_k^x} g'(x', x | k, \alpha, \beta) r(x'), \quad k < N.$$

For our experiment, we set the true value of $\theta = 0.3$ and a time horizon of $N = 50$. Figure 5.1 illustrates a sample realisation of an optimal trajectory, given a prior of $\text{Beta}(5, 2)$, as well as the evolution of the estimate over θ over time.

We consider three different choices of (α, β) for the prior ρ_0 as well as varying the observation cost. For each parameter combination we compute the optimal policy and compare their respective performances across 5000 sampled trajectories. A typical criteria of measuring the performance of the policy is to examine its regret, defined as

$$\text{reg}(N, \pi) = J_0^{\theta^*}(\pi^*) - J^N(\pi),$$

where $J_0^{\theta^*}(\pi^*) = \mathbb{E}^{\pi^*}[\sum r(x, a)]$ is the reward functional with no observation cost, with known parameter θ^* and optimal policy π^* , and $J^N(\pi)$ is the reward functional (5.4) under a policy $\pi \in \Pi_{\text{obs}}$. In this case we consider π to be the optimal policy under observation costs and parameter uncertainty. The regret is therefore the cumulative sum of the suboptimal gap from the optimal policy. For the observation cost problem, the control between observations are constant and therefore suboptimal in general, as such we do not expect the regret to achieve asymptotically sublinear growth. Instead, we consider the following alternative criteria:

$$\text{reg}_{c_{\text{obs}}}(N, \pi) = J_{c_{\text{obs}}}^{\theta^*}(\pi^*) - J^N(\pi),$$

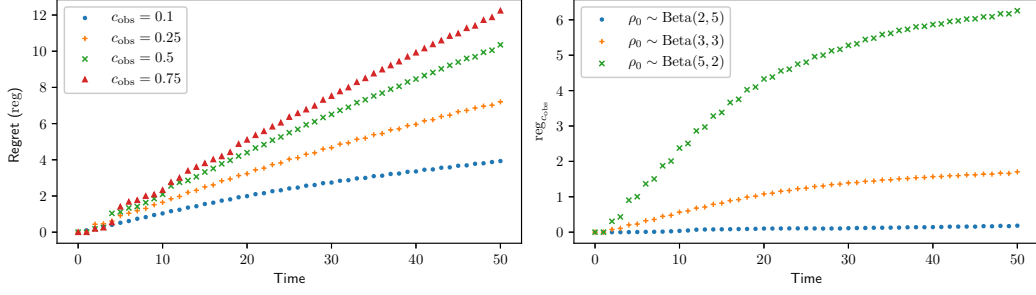


Figure 5.2: Left: regret over time for $\rho_0 \sim \text{Beta}(3, 3)$ for different values of c_{obs} . Right: the growth of $\text{reg}_{c_{\text{obs}}}$ for fixed $c_{\text{obs}} = 0.1$ and different initial priors ρ_0 .

where here $J_{c_{\text{obs}}}^{\theta^*}$ is the reward functional with known parameter θ^* and observation cost c_{obs} , so that $\text{reg}_{c_{\text{obs}}}$ measures the contribution of the regret that arises from parameter uncertainty. On the left of Figure 5.2, we show the overall regret for varying the observation cost for a fixed prior, and on the right, $\text{reg}_{c_{\text{obs}}}$ is plotted with a fixed observation cost of $c_{\text{obs}} = 0.1$ across different initial priors ρ_0 . As expected, the regret is generally higher when the prior estimate ρ_0 is less accurate, or when a larger c_{obs} value is used. Moreover the regret grows in a rather linear fashion. However, when examining the graph involving $\text{reg}_{c_{\text{obs}}}$ on the right side, we empirically observe sublinear growth. This can be interpreted as a gradual learning of the unknown parameters, despite the fact that observations only arrive in intervals. The results suggests that $\text{reg}_{c_{\text{obs}}}$ can be used as an alternative notion to capture the learning rate in problems involving observation costs, which we see as a possible direction for future analysis.

To demonstrate the effects of observation cost and prior estimates on the number of observations, we consider an alternative reward function, given by

$$r(x, a) = r(x) = \begin{cases} 2 & x = 0 \\ -1 & x = \pm 2 \\ 0 & \text{otherwise} \end{cases}$$

In the absence of observation cost and parameter uncertainty, the controller aims to keep the process at the origin as often as possible, whilst avoiding the penalising boundary at $x = \pm 2$. Table 5.3 lists the performance of the optimal policies under each combination of observation cost and prior estimate. As the true value of $\theta = 0.3$, a prior of $\rho_0 \sim \text{Beta}(2, 5)$ acts a good estimate, and $\rho_0 \sim \text{Beta}(5, 2)$ acts as a poor estimate. In general, we see that the value of the observation cost c_{obs} has a more dominating effect on the resulting optimal policies and rewards obtained, as seen in the big drop-off in the number of observations when $c_{\text{obs}} = 0.75$ in row (a), at which each observation comes at the cost of a significant proportion of the potential reward. Its effect on the sub-optimality is compounded with a bad prior estimate, with a negative reward and a 95% credible interval width of 0.5 in the extreme case in the bottom-right entry of Table 5.3.

| | | $c_{\text{obs}} = 0.1$ | $c_{\text{obs}} = 0.25$ | $c_{\text{obs}} = 0.5$ | $c_{\text{obs}} = 0.75$ |
|---------------------------------|-----|------------------------|-------------------------|------------------------|-------------------------|
| $\rho_0 \sim \text{Beta}(2, 5)$ | (a) | 22.48 | 22.2 | 21.2 | 17.55 |
| | (b) | 20.622 | 17.15 | 11.26 | 6.0375 |
| | (c) | 0.2341 | 0.2360 | 0.2455 | 0.2844 |
| $\rho_0 \sim \text{Beta}(3, 3)$ | (a) | 21.4 | 20.97 | 18.36 | 11.27 |
| | (b) | 17.99 | 14.6475 | 8.92 | 2.5775 |
| | (c) | 0.2437 | 0.2459 | 0.2696 | 0.3624 |
| $\rho_0 \sim \text{Beta}(5, 2)$ | (a) | 19.22 | 17.3 | 11.21 | 3.34 |
| | (b) | 10.628 | 7.55 | 1.825 | -0.835 |
| | (c) | 0.2488 | 0.2583 | 0.3302 | 0.5034 |

Table 5.3: Numerical results for the parameter uncertainty problem. Line (a): average number of observations. Line (b): average profit ($N = 50$). Line (c): average credible interval width (HDI 95%).

Acknowledgments.

The authors would like to thank Prof. Dr. Dirk Becherer (Humboldt University of Berlin) for his insightful suggestions during discussion, as well as the two anonymous referees for their feedback. Jonathan Tam is supported by the EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation (EP/S023925/1).

References

- [1] R. F. Anderson and A. Friedman. Optimal inspections in a stochastic control problem with costly observations. *Math. Oper. Res.*, 2(2):155–190, 1977.
- [2] R. F. Anderson and A. Friedman. Optimal inspections in a stochastic control problem with costly observations, II. *Math. Oper. Res.*, 3(1):67–81, 1978.
- [3] P. Azimzadeh and P. A. Forsyth. Weakly chained matrices, policy iteration, and impulse control. *SIAM J. Numer. Anal.*, 54(3):1341–1364, 2016.
- [4] E. Bayraktar, E. Ekström, and J. Guo. Disorder detection with costly observations. *J. Appl. Probab.*, 59(2):338–349, 2022.
- [5] E. Bayraktar and R. Kravitz. Quickest detection with discretely controlled observations. *Sequential Anal.*, 34(1):77–133, 2015.
- [6] C. Bellinger, R. Coles, M. Crowley, and I. Tamblyn. Active measure reinforcement learning for observation cost minimization. *arXiv:2005.12697*, 2020.
- [7] C. Bellinger, A. Drozdoyuk, M. Crowley, and I. Tamblyn. Balancing information with observation costs in deep reinforcement learning. In *Proceedings of the 35th Canadian Conference on Artificial Intelligence*. CAIAC, 2022.
- [8] A. Briani, F. Camilli, and H. Zidani. Approximation Schemes for Monotone Systems of Nonlinear Second Order Partial Differential Equations: Convergence Result and Error Estimate. *Differential Equations and Applications*, 4:297–317, 2012.
- [9] C. Cooper and N. Hahi. An optimal stochastic control problem with observation cost. *IEEE T. Automat. Contr.*, 16(2):185–189, 1971.
- [10] R. C. Dalang and A. N. Shiryaev. A quickest detection problem with an observation cost. *Ann. Appl. Probab.*, 25(3):1475–1512, 2015.
- [11] H. Dyrssen and E. Ekström. Sequential testing of a Wiener process with costly observations. *Sequential Anal.*, 37(1):47–58, 2018.

- [12] P. A. Forsyth and K. R. Vetzal. Quadratic convergence for valuing american options using a penalty method. *SIAM J. Sci. Comput.*, 23(6):2095–2122, 2002.
- [13] N. Guo and V. Kostina. Optimal causal rate-constrained sampling for a class of continuous markov processes. *IEEE T. Inform. Theory*, 67(12):7876–7890, 2021.
- [14] B. Hajek, K. Mitzel, and S. Yang. Paging and registration in cellular networks: Jointly optimal policies and an iterative algorithm. *IEEE T. Inform. Theory*, 54:608 – 622, 2008.
- [15] O. Hernández-Lerma. *Adaptive Markov Control Processes*. Applied mathematical sciences. Springer-Verlag, 1989.
- [16] C. Huang and S. Wang. A power penalty approach to a nonlinear complementarity problem. *Oper. Res. Lett.*, 38(1):72–76, 2010.
- [17] Y. Huang and Q. Zhu. Self-triggered Markov decision processes. In *Proc. 60th IEEE CDC*, pages 4507–4514, 2021.
- [18] K. Ito and K. Kunisch. Parabolic variational inequalities: The lagrange multiplier approach. *J. Math. Pures Appl.*, 85(3):415–449, 2006.
- [19] D. Krueger, J. Leike, O. Evans, and J. Salvatier. Active reinforcement learning: Observing rewards at a cost. *arXiv:2011.06709*, 2020.
- [20] H. Kushner. On the optimum timing of observations for linear control systems with unknown initial state. *IEEE T. Automat. Contr.*, 9(2):144–150, 1964.
- [21] H. J. Kushner and P. G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, Berlin, Heidelberg, 1992.
- [22] L. Meier, J. Peschon, and R. Dressler. Optimal control of measurement subsystems. *IEEE T. Automat. Contr.*, 12(5):528–536, 1967.
- [23] A. Nayyar, T. Başar, D. Teneketzis, and V. V. Veeravalli. Optimal strategies for communication and remote estimation with an energy harvesting sensor. *IEEE T. Automat. Contr.*, 58(9):2246–2260, 2013.
- [24] H. Pham. *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Springer Publishing Company, Incorporated, 1st edition, 2009.
- [25] R. Reis. Inattentive consumers. *J. Monet. Econ.*, 53(8):1761–1800, 2006.
- [26] R. Reis. Inattentive producers. *Rev. Econ. Stud.*, 73(3):793–821, 2006.
- [27] C. Reisinger and J. H. Witte. On the use of policy iteration as an easy way of pricing American options. *SIAM Jo. Financ. Math.*, 3(1):459–478, 2012.
- [28] C. Reisinger and Y. Zhang. A penalty scheme for monotone systems with interconnected obstacles: Convergence and error estimates. *SIAM J. Numer. Anal.*, 57(4):1625–1648, 2019.
- [29] C. Reisinger and Y. Zhang. Error estimates of penalty schemes for quasi-variational inequalities arising from impulse control problems. *SIAM Journal on Control and Optimization*, 58(1):243–276, 2020.
- [30] C. Reisinger and Y. Zhang. A penalty scheme and policy iteration for nonlocal hjb variational inequalities with monotone nonlinearities. *Comput. Math. Appl.*, 93:199–213, 2021.
- [31] J. Rust. Numerical dynamic programming in economics. *Handbook of computational economics*, 1:619–729, 1996.
- [32] V. Tzoumas, L. Carlone, G. J. Pappas, and A. Jadbabaie. LQG control and sensing co-design. *IEEE T. Automat. Contr.*, 66(4):1468–1483, 2020.
- [33] S. Winkelmann. *Markov Decision Processes with Information Costs*. PhD thesis, Freie Universität Berlin, Berlin, 2013.

- [34] S. Winkelmann, C. Schütte, and M. v. Kleist. Markov control processes with rare state observation: Theory and application to treatment scheduling in HIV-1. *Commun. Math. Sci.*, 12(5):859–877, 2014.
- [35] J. H. Witte and C. Reisinger. A penalty method for the numerical solution of Hamilton–Jacobi–Bellman (HJB) equations in finance. *SIAM J. Numer. Anal.*, 49(1):213–231, 2011.
- [36] J. H. Witte and C. Reisinger. Penalty methods for the solution of discrete HJB equations—continuous control and obstacle problems. *SIAM J. Numer. Anal.*, 50(2):595–625, 2012.
- [37] W. Wu and A. Arapostathis. Optimal sensor querying: General Markovian and LQG models with controlled observations. *IEEE T. Automat. Contr.*, 53(6):1392–1405, 2008.
- [38] H. Yoshioka and M. Tsujimura. Analysis and computation of an optimality equation arising in an impulse control problem with discrete and costly observations. *J. Comput. Appl. Math.*, 366:112399, 2020.
- [39] H. Yoshioka, M. Tsujimura, K. Hamagami, and Y. Yoshioka. A hybrid stochastic river environmental restoration modeling with discrete and costly observations. *Optimal Control Appl. Methods*, 41(6):1964–1994, 2020.
- [40] H. Yoshioka, Y. Yaegashi, M. Tsujimura, and Y. Yoshioka. Cost-efficient monitoring of continuous-time stochastic processes based on discrete observations. *Appl. Stoch. Models Bus. Ind.*, 37(1):113–138, 2021.
- [41] H. Yoshioka, Y. Yoshioka, Y. Yaegashi, T. Tanaka, M. Horinouchi, and F. Aranishi. Analysis and computation of a discrete costly observation model for growth estimation and management of biological resources. *Comput. Math. Appl.*, 79(4):1072–1093, 2020.